

Azure Provider

The Azure Provider can be used to configure infrastructure in Microsoft Azure (<https://azure.microsoft.com/en-us/>) using the Azure Resource Manager API's. Documentation regarding the Data Sources (</docs/configuration/data-sources.html>) and Resources (</docs/configuration/resources.html>) supported by the Azure Provider can be found in the navigation to the left.

Interested in the provider's latest features, or want to make sure you're up to date? Check out the [changelog](https://github.com/terraform-providers/terraform-provider-azurerm/blob/master/CHANGELOG.md) (<https://github.com/terraform-providers/terraform-provider-azurerm/blob/master/CHANGELOG.md>) for version information and release notes.

Authenticating to Azure

Terraform supports a number of different methods for authenticating to Azure:

- Authenticating to Azure using the Azure CLI ([/docs/providers/azurerm/auth/azure_cli.html](/docs/providers/azurerm/auth/azure_cli.html))
- Authenticating to Azure using Managed Service Identity ([/docs/providers/azurerm/auth/managed_service_identity.html](/docs/providers/azurerm/auth/managed_service_identity.html))
- Authenticating to Azure using a Service Principal and a Client Certificate ([/docs/providers/azurerm/auth/service_principal_client_certificate.html](/docs/providers/azurerm/auth/service_principal_client_certificate.html))
- Authenticating to Azure using a Service Principal and a Client Secret ([/docs/providers/azurerm/auth/service_principal_client_secret.html](/docs/providers/azurerm/auth/service_principal_client_secret.html))

We recommend using either a Service Principal or Managed Service Identity when running Terraform non-interactively (such as when running Terraform in a CI server) - and authenticating using the Azure CLI when running Terraform locally.

Example Usage

```
# Configure the Azure Provider
provider "azurerm" {
  # whilst the `version` attribute is optional, we recommend pinning to a given version of the Provider
  version = "=1.36.0"
}

# Create a resource group
resource "azurerm_resource_group" "test" {
  name      = "production"
  location  = "West US"
}

# Create a virtual network within the resource group
resource "azurerm_virtual_network" "test" {
  name                = "production-network"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  address_space       = ["10.0.0.0/16"]
}
```

Features and Bug Requests

The Azure provider's bugs and feature requests can be found in the [GitHub repo issues](https://github.com/terraform-providers/terraform-provider-azurerm/issues) (<https://github.com/terraform-providers/terraform-provider-azurerm/issues>). Please avoid "me too" or "+1" comments. Instead, use a thumbs up reaction (<https://blog.github.com/2016-03-10-add-reactions-to-pull-requests-issues-and-comments/>) on enhancement requests. Provider maintainers will often prioritize work based on the number of thumbs on an issue.

Community input is appreciated on outstanding issues! We love to hear what use cases you have for new features, and want to provide the best possible experience for you using the Azure provider.

If you have a bug or feature request without an existing issue

- if an existing resource or field is working in an unexpected way, file a bug (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/new?template=bug.md>).
- if you'd like the provider to support a new resource or field, file an enhancement/feature request (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/new?template=enhancement.md>).

The provider maintainers will often use the assignee field on an issue to mark who is working on it.

- An issue assigned to an individual maintainer indicates that maintainer is working on the issue
- If you're interested in working on an issue please leave a comment in that issue

If you have configuration questions, or general questions about using the provider, try checking out:

- Terraform's community resources (<https://www.terraform.io/docs/extend/community/index.html>)
- HashiCorp support (<https://support.hashicorp.com>) for Terraform Enterprise customers

Argument Reference

The following arguments are supported:

- `client_id` - (Optional) The Client ID which should be used. This can also be sourced from the `ARM_CLIENT_ID` Environment Variable.
- `environment` - (Optional) The Cloud Environment which should be used. Possible values are `public`, `usgovernment`, `german`, and `china`. Defaults to `public`. This can also be sourced from the `ARM_ENVIRONMENT` environment variable.
- `subscription_id` - (Optional) The Subscription ID which should be used. This can also be sourced from the `ARM_SUBSCRIPTION_ID` Environment Variable.
- `tenant_id` - (Optional) The Tenant ID which should be used. This can also be sourced from the `ARM_TENANT_ID` Environment Variable.

When authenticating as a Service Principal using a Client Certificate, the following fields can be set:

- `client_certificate_password` - (Optional) The password associated with the Client Certificate. This can also be sourced from the `ARM_CLIENT_CERTIFICATE_PASSWORD` Environment Variable.

- `client_certificate_path` - (Optional) The path to the Client Certificate associated with the Service Principal which should be used. This can also be sourced from the `ARM_CLIENT_CERTIFICATE_PATH` Environment Variable.

More information on how to configure a Service Principal using a Client Certificate can be found in this guide ([/docs/providers/azurerm/auth/service_principal_client_certificate.html](/docs/providers/azurerm/auth/service_principal_client_certificate.html)).

When authenticating as a Service Principal using a Client Secret, the following fields can be set:

- `client_secret` - (Optional) The Client Secret which should be used. This can also be sourced from the `ARM_CLIENT_SECRET` Environment Variable.

More information on how to configure a Service Principal using a Client Secret can be found in this guide ([/docs/providers/azurerm/auth/service_principal_client_secret.html](/docs/providers/azurerm/auth/service_principal_client_secret.html)).

When authenticating using Managed Service Identity, the following fields can be set:

- `msi_endpoint` - (Optional) The path to a custom endpoint for Managed Service Identity - in most circumstances this should be detected automatically. This can also be sourced from the `ARM_MSI_ENDPOINT` Environment Variable.
- `use_msi` - (Optional) Should Managed Service Identity be used for Authentication? This can also be sourced from the `ARM_USE_MSI` Environment Variable. Defaults to `false`.

More information on how to configure a Service Principal using Managed Service Identity can be found in this guide ([/docs/providers/azurerm/auth/managed_service_identity.html](/docs/providers/azurerm/auth/managed_service_identity.html)).

For some advanced scenarios, such as where more granular permissions are necessary - the following properties can be set:

- `disable_terraform_partner_id` - (Optional) Disable sending the Terraform Partner ID if a custom `partner_id` isn't specified, which allows Microsoft to better understand the usage of Terraform. The Partner ID does not give HashiCorp any direct access to usage information. This can also be sourced from the `ARM_DISABLE_TERRAFORM_PARTNER_ID` environment variable. Defaults to `false`.
- `partner_id` - (Optional) A GUID/UUID that is registered (<https://docs.microsoft.com/azure/marketplace/azure-partner-customer-usage-attribution#register-guids-and-offers>) with Microsoft to facilitate partner resource usage attribution. This can also be sourced from the `ARM_PARTNER_ID` Environment Variable.
- `skip_credentials_validation` - (Optional) Should the AzureRM Provider skip verifying the credentials being used are valid? This can also be sourced from the `ARM_SKIP_CREDENTIALS_VALIDATION` Environment Variable. Defaults to `false`.
- `skip_provider_registration` - (Optional) Should the AzureRM Provider skip registering the Resource Providers it supports? This can also be sourced from the `ARM_SKIP_PROVIDER_REGISTRATION` Environment Variable. Defaults to `false`.

By default, Terraform will attempt to register any Resource Providers that it supports, even if they're not used in your configurations to be able to display more helpful error messages. If you're running in an environment with restricted permissions, or wish to manage Resource Provider Registration outside of Terraform you may wish to disable this flag; however please note that the error messages returned from Azure may be confusing as a result (example: `API version 2019-01-01 was not found for Microsoft.Foo`).

It's also possible to use multiple Provider blocks within a single Terraform configuration, for example to work with resources

across multiple Subscriptions - more information can be found in the documentation for Providers (<https://www.terraform.io/docs/configuration/providers.html#multiple-provider-instances>).

Azure Provider: Authenticating using the Azure CLI

Terraform supports a number of different methods for authenticating to Azure:

- Authenticating to Azure using the Azure CLI (which is covered in this guide)
- Authenticating to Azure using Managed Service Identity
(/docs/providers/azurerm/auth/managed_service_identity.html)
- Authenticating to Azure using a Service Principal and a Client Certificate
(/docs/providers/azurerm/auth/service_principal_client_certificate.html)
- Authenticating to Azure using a Service Principal and a Client Secret
(/docs/providers/azurerm/auth/service_principal_client_secret.html)

We recommend using either a Service Principal or Managed Service Identity when running Terraform non-interactively (such as when running Terraform in a CI server) - and authenticating using the Azure CLI when running Terraform locally.

Important Notes about Authenticating using the Azure CLI

- Prior to version 1.20 the AzureRM Provider used a different method of authorizing via the Azure CLI where credentials reset after an hour - as such we'd recommend upgrading to version 1.20 or later of the AzureRM Provider.
- Terraform only supports authenticating using the `az` CLI (and this must be available on your PATH) - authenticating using the older `azure` CLI or PowerShell Cmdlets is not supported.
- Authenticating via the Azure CLI is only supported when using a User Account. If you're using a Service Principal (for example via `az login --service-principal`) you should instead authenticate via the Service Principal directly (either using a Client Secret (/docs/providers/azurerm/auth/service_principal_client_secret.html) or a Client Certificate (/docs/providers/azurerm/auth/service_principal_client_certificate.html)).

Logging into the Azure CLI

Note: If you're using the **China**, **German** or **Government** Azure Clouds - you'll need to first configure the Azure CLI to work with that Cloud. You can do this by running:

```
$ az cloud set --name AzureChinaCloud|AzureGermanCloud|AzureUSGovernment
```

Firstly, login to the Azure CLI using:

```
$ az login
```

Once logged in - it's possible to list the Subscriptions associated with the account via:

```
$ az account list
```

The output (similar to below) will display one or more Subscriptions - with the `id` field being the `subscription_id` field referenced above.

```
[
  {
    "cloudName": "AzureCloud",
    "id": "000000000-0000-0000-0000-000000000000",
    "isDefault": true,
    "name": "PAYG Subscription",
    "state": "Enabled",
    "tenantId": "000000000-0000-0000-0000-000000000000",
    "user": {
      "name": "user@example.com",
      "type": "user"
    }
  }
]
```

Should you have more than one Subscription, you can specify the Subscription to use via the following command:

```
$ az account set --subscription="SUBSCRIPTION_ID"
```

Configuring Azure CLI authentication in Terraform

Now that we're logged into the Azure CLI - we can configure Terraform to use these credentials.

To configure Terraform to use the Default Subscription defined in the Azure CLI - we can use the following Provider block:

```
provider "azurerm" {
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being
  # used
  version = "=1.36.0"
}
```

More information on the fields supported in the Provider block can be found here (</docs/providers/azurerm/index.html#argument-reference>).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Azure CLI to authenticate.

It's also possible to configure Terraform to use a specific Subscription - for example:

```
provider "azurerm" {  
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being  
  used  
  version = "=1.36.0"  
  
  subscription_id = "00000000-0000-0000-0000-000000000000"  
}
```

More information on the fields supported in the Provider block can be found here
(</docs/providers/azurerm/index.html#argument-reference>).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Azure CLI to authenticate.

If you're looking to use Terraform across Tenants - it's possible to do this by configuring the Tenant ID field in the Provider block, as shown below:

```
provider "azurerm" {  
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being  
  used  
  version = "=1.36.0"  
  
  subscription_id = "00000000-0000-0000-0000-000000000000"  
  tenant_id      = "11111111-1111-1111-1111-111111111111"  
}
```

More information on the fields supported in the Provider block can be found here
(</docs/providers/azurerm/index.html#argument-reference>).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Azure CLI to authenticate.

Azure Provider: Authenticating using managed identities for Azure resources

Terraform supports a number of different methods for authenticating to Azure:

- Authenticating to Azure using the Azure CLI ([/docs/providers/azurerm/auth/azure_cli.html](/docs/providers/azurerm/auth/azure_cli.html))
- Authenticating to Azure using Managed Identity (covered in this guide)
- Authenticating to Azure using a Service Principal and a Client Certificate ([/docs/providers/azurerm/auth/service_principal_client_certificate.html](/docs/providers/azurerm/auth/service_principal_client_certificate.html))
- Authenticating to Azure using a Service Principal and a Client Secret ([/docs/providers/azurerm/auth/service_principal_client_secret.html](/docs/providers/azurerm/auth/service_principal_client_secret.html))

We recommend using a service principal or a managed identity when running Terraform non-interactively (such as when running Terraform in a CI/CD pipeline), and authenticating using the Azure CLI when running Terraform locally.

What is a managed identity?

Managed identities for Azure resources (<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>) can be used to authenticate to services that support Azure Active Directory (Azure AD) authentication. There are two types of managed identities: system-assigned and user-assigned. This article is based on system-assigned managed identities.

Managed identities work in conjunction with Azure Resource Manager (ARM), Azure AD, and the Azure Instance Metadata Service (IMDS). Azure resources that support managed identities expose an internal IMDS endpoint that the client can use to request an access token. No credentials are stored on the VM, and the only additional information needed to bootstrap the Terraform connection to Azure is the subscription ID and tenant ID.

Azure AD creates an AD identity when you configure an Azure resource to use a system-assigned managed identity. The configuration process is described in more detail, below. Azure AD then creates a service principal to represent the resource for role-based access control (RBAC) and access control (IAM). The lifecycle of a system-assigned identity is tied to the resource it is enabled for: it is created when the resource is created and it is automatically removed when the resource is deleted.

Before you can use the managed identity, it has to be configured. There are two steps:

1. Assign a role for the identity, associating it with the subscription that will be used to run Terraform. This step gives the identity permission to access Azure Resource Manager (ARM) resources.
2. Configure access control for one or more Azure resources. For example, if you use a key vault and a storage account, you will need to configure the vault and container separately.

Before you can create a resource with a managed identity and then assign an RBAC role, your account needs sufficient permissions. You need to be a member of the account **Owner** role, or have **Contributor** plus **User Access Administrator** roles.

Not all Azure services support managed identities, and availability varies by region. Configuration details vary slightly among services. For more information, see [Services that support managed identities for Azure resources](https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/services-support-managed-identities) (<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/services-support-managed-identities>).

Configuring a VM to use a system-assigned managed identity

The (simplified) Terraform configuration below provisions a virtual machine with a system-assigned managed identity, and then grants the Contributor role to the identity.

```
data "azurerm_subscription" "current" {}

resource "azurerm_virtual_machine" "test" {
  # ...

  identity {
    type = "SystemAssigned"
  }
}

data "azurerm_builtin_role_definition" "contributor" {
  name = "Contributor"
}

resource "azurerm_role_assignment" "test" {
  name                = "${azurerm_virtual_machine.test.name}"
  scope               = "${data.azurerm_subscription.primary.id}"
  role_definition_id = "${data.azurerm_subscription.subscription.id}${data.azurerm_builtin_role_definition.contributor.id}"
  principal_id        = "${lookup(azurerm_virtual_machine.test.identity[0], "principal_id")}"
}
```

Configuring Terraform to use a managed identity

At this point we assume that managed identity is configured on the resource (e.g. virtual machine) being used - and that permissions have been assigned via Azure's Identity and Access Management system.

Terraform can be configured to use managed identity for authentication in one of two ways: using environment variables, or by defining the fields within the provider block.

Configuring with environment variables

Setting the `ARM_USE_MSI` environment variable to `true` tells Terraform to use a managed identity. In addition to a properly-configured management identity, Terraform needs to know the subscription ID and tenant ID to identify the full context for the Azure provider.

```
$ export ARM_USE_MSI=true
$ export ARM_SUBSCRIPTION_ID=159f2485-xxxx-xxxx-xxxx-xxxxxxxxxxxxx
$ export ARM_TENANT_ID=72f988bf-xxxx-xxxx-xxxx-xxxxxxxxxxxxx
```

A provider block is *technically* optional when using environment variables. Even so, we recommend defining a provider block so that you can pin or constrain the version of the provider being used:

```
provider "azurerm" {
  version = "~> 1.23"
}
```

Configuring with the provider block

It's also possible to configure a managed identity within the provider block:

```
provider "azurerm" {
  version = "~> 1.23"

  use_msi = true

  # ...
}
```

If you intend to configure a remote backend in the provider block, put `use_msi` outside of the backend block:

```
provider "azurerm" {
  version = "~> 1.23"
  use_msi = true

  backend "azurerm" {
    storage_account_name = "abcd1234"
    container_name       = "tfstate"
    key                  = "prod.terraform.tfstate"
    subscription_id      = "00000000-0000-0000-0000-000000000000"
    tenant_id            = "00000000-0000-0000-0000-000000000000"
  }
}
```

More information on the fields supported in the provider block can be found here (</docs/providers/azurerm/index.html#argument-reference>).

Azure Provider: Authenticating using a Service Principal with a Client Certificate

Terraform supports a number of different methods for authenticating to Azure:

- Authenticating to Azure using the Azure CLI ([/docs/providers/azurerm/auth/azure_cli.html](/docs/providers/azurerm/auth/azure_cli.html))
- Authenticating to Azure using Managed Service Identity ([/docs/providers/azurerm/auth/managed_service_identity.html](/docs/providers/azurerm/auth/managed_service_identity.html))
- Authenticating to Azure using a Service Principal and a Client Certificate (which is covered in this guide)
- Authenticating to Azure using a Service Principal and a Client Secret ([/docs/providers/azurerm/auth/service_principal_client_secret.html](/docs/providers/azurerm/auth/service_principal_client_secret.html))

We recommend using either a Service Principal or Managed Service Identity when running Terraform non-interactively (such as when running Terraform in a CI server) - and authenticating using the Azure CLI when running Terraform locally.

Creating a Service Principal

A Service Principal is an application within Azure Active Directory which can be used as a means of authentication, either using a Client Secret ([/docs/providers/azurerm/auth/service_principal_client_secret.html](/docs/providers/azurerm/auth/service_principal_client_secret.html)) or a Client Certificate (which is documented in this guide) and can be created through the Azure Portal.

This guide will cover how to generate a client certificate, how to create a Service Principal and then how to assign the Client Certificate to the Service Principal so that it can be used for authentication. Once that's done finally we're going to grant the Service Principal permission to manage resources in the Subscription - to do this we're going to assign `Contributor` rights to the Subscription - however it's possible to assign other permissions (<https://azure.microsoft.com/en-gb/documentation/articles/role-based-access-built-in-roles/>) depending on your configuration.

Generating a Client Certificate

Firstly we need to create a certificate which can be used for authentication. To do that we're going to generate a Certificate Signing Request (also known as a CSR) using `openssl` (this can also be achieved using PowerShell, however that's outside the scope of this document):

```
$ openssl req -newkey rsa:4096 -nodes -keyout "service-principal.key" -out "service-principal.csr"
```

During the generation of the certificate you'll be prompted for various bits of information required for the certificate signing request - at least one item has to be specified for this to complete.

We can now sign that Certificate Signing Request, in this example we're going to self-sign this certificate using the Key we just generated; however it's also possible to do this using a Certificate Authority. In order to do that we're again going to use `openssl` :

```
$ openssl x509 -signkey "service-principal.key" -in "service-principal.csr" -req -days 365 -out "service-principal.crt"
```

Finally we can generate a PFX file which can be used to authenticate with Azure:

```
$ openssl pkcs12 -export -out "service-principal.pfx" -inkey "service-principal.key" -in "service-principal.crt"
```

Now that we've generated a certificate, we can create the Azure Active Directory application.

Creating the Service Principal

We're going to create the Service Principal in the Azure Portal - to do this navigate to the **Azure Active Directory** overview (https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/Overview) within the Azure Portal - then select the **App Registration** blade

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) and click **Endpoints** at the top of the **App Registration** blade. A list of URIs will be displayed and you need to locate the URI for **OAUTH 2.0 AUTHORIZATION ENDPOINT** which contains a GUID. This GUID is your Tenant ID (the `tenant_id` field mentioned above).

Next, navigate back to the **App Registration** blade

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) - from here we'll create the Application in Azure Active Directory. To do this click **New application registration** at the top to add a new Application within Azure Active Directory. On this page, set the following values then press **Create**:

- **Name** - this is a friendly identifier and can be anything (e.g. "Terraform")
- **Application Type** - this should be set to "Web app / API"
- **Sign-on URL** - this can be anything, providing it's a valid URI (e.g. <https://terra.form> (`https://terra.form`))

At this point the newly created Azure Active Directory application should be visible on-screen - if it's not, navigate to the **App Registration** blade

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) and select the Azure Active Directory application. At the top of this page, the "Application ID" GUID is the `client_id` you'll need.

Assigning the Client Certificate to the Service Principal

To associate the public portion of the Client Certificate (the `*.crt` file) with the Azure Active Directory Application - to do this select **Settings** and then **Keys**. This screen displays the Passwords (Client Secrets) and Public Keys (Client Certificates) which are associated with this Azure Active Directory Application.

The Public Key associated with the generated Certificate can be uploaded by selecting **Upload Public Key**, selecting the file which should be uploaded (in the example above, this'd be `service-principal.crt`) - and then hitting **Save**.

Allowing the Service Principal to manage the Subscription

Now that we've created the Application within Azure Active Directory and assigned the certificate we're using for authentication, we can now grant the Application permissions to manage the Subscription. To do this, navigate to the **Subscriptions** blade within the Azure Portal (https://portal.azure.com/#blade/Microsoft_Azure_Billing/SubscriptionsBlade), select the Subscription you wish to use, then click **Access Control (IAM)** and finally **Add role assignment**.

Firstly, specify a Role which grants the appropriate permissions needed for the Service Principal (for example, **Contributor** will grant Read/Write on all resources in the Subscription). More information about the built in roles can be found here (<https://azure.microsoft.com/en-gb/documentation/articles/role-based-access-built-in-roles/>).

Secondly, search for and select the name of the Application created in Azure Active Directory to assign it this role - then press **Save**.

At this point the newly created Azure Active Directory Application should be associated with the Certificate that we generated earlier (which can be used as a Client Certificate) - and should have permissions to the Azure Subscription.

Configuring the Service Principal in Terraform

As we've obtained the credentials for this Service Principal - it's possible to configure them in a few different ways.

When storing the credentials as Environment Variables, for example:

```
$ export ARM_CLIENT_ID="00000000-0000-0000-0000-000000000000"
$ export ARM_CLIENT_CERTIFICATE_PATH="/path/to/my/client/certificate.pfx"
$ export ARM_CLIENT_CERTIFICATE_PASSWORD="Pa55w0rd123"
$ export ARM_SUBSCRIPTION_ID="00000000-0000-0000-0000-000000000000"
$ export ARM_TENANT_ID="00000000-0000-0000-0000-000000000000"
```

The following Provider block can be specified - where **1.21.0** is the version of the Azure Provider that you'd like to use:

```
provider "azurerm" {
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being
  # used
  version = "=1.36.0"
}
```

More information on the fields supported in the Provider block can be found here (</docs/providers/azurerm/index.html#argument-reference>).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Service Principal to authenticate.

It's also possible to configure these variables either in-line or from using variables in Terraform (as the `client_certificate_path` and `client_certificate_password` are in this example), like so:

NOTE: We'd recommend not defining these variables in-line since they could easily be checked into Source Control.

```
variable "client_certificate_path" {}
variable "client_certificate_password" {}

provider "azurerm" {
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being
  used
  version = "=1.36.0"

  subscription_id      = "00000000-0000-0000-0000-000000000000"
  client_id             = "00000000-0000-0000-0000-000000000000"
  client_certificate_path = "${var.client_certificate_path}"
  client_certificate_password = "${var.client_certificate_password}"
  tenant_id             = "00000000-0000-0000-0000-000000000000"
}
```

More information on the fields supported in the Provider block can be found here (</docs/providers/azurerm/index.html#argument-reference>).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Service Principal to authenticate.

Azure Provider: Authenticating using a Service Principal with a Client Secret

Terraform supports a number of different methods for authenticating to Azure:

- Authenticating to Azure using the Azure CLI ([/docs/providers/azurerm/auth/azure_cli.html](/docs/providers/azurerm/auth/azure_cli.html))
- Authenticating to Azure using Managed Service Identity ([/docs/providers/azurerm/auth/managed_service_identity.html](/docs/providers/azurerm/auth/managed_service_identity.html))
- Authenticating to Azure using a Service Principal and a Client Certificate ([/docs/providers/azurerm/auth/service_principal_client_certificate.html](/docs/providers/azurerm/auth/service_principal_client_certificate.html))
- Authenticating to Azure using a Service Principal and a Client Secret (which is covered in this guide)

We recommend using either a Service Principal or Managed Service Identity when running Terraform non-interactively (such as when running Terraform in a CI server) - and authenticating using the Azure CLI when running Terraform locally.

Creating a Service Principal

A Service Principal is an application within Azure Active Directory whose authentication tokens can be used as the `client_id`, `client_secret`, and `tenant_id` fields needed by Terraform (`subscription_id` can be independently recovered from your Azure account details).

It's possible to complete this task in either the Azure CLI or in the Azure Portal - in both we'll create a Service Principal which has `Contributor` rights to the subscription. It's also possible to assign other rights (<https://azure.microsoft.com/en-gb/documentation/articles/role-based-access-built-in-roles/>) depending on your configuration.

Creating a Service Principal using the Azure CLI

Note: If you're using the **China**, **German** or **Government** Azure Clouds - you'll need to first configure the Azure CLI to work with that Cloud. You can do this by running:

```
$ az cloud set --name AzureChinaCloud|AzureGermanCloud|AzureUSGovernment
```

Firstly, login to the Azure CLI using:

```
$ az login
```

Once logged in - it's possible to list the Subscriptions associated with the account via:

```
$ az account list
```

The output (similar to below) will display one or more Subscriptions - with the `id` field being the `subscription_id` field referenced above.

```
[
  {
    "cloudName": "AzureCloud",
    "id": "000000000-0000-0000-0000-000000000000",
    "isDefault": true,
    "name": "PAYG Subscription",
    "state": "Enabled",
    "tenantId": "000000000-0000-0000-0000-000000000000",
    "user": {
      "name": "user@example.com",
      "type": "user"
    }
  }
]
```

Should you have more than one Subscription, you can specify the Subscription to use via the following command:

```
$ az account set --subscription="SUBSCRIPTION_ID"
```

We can now create the Service Principal which will have permissions to manage resources in the specified Subscription using the following command:

```
$ az ad sp create-for-rbac --role="Contributor" --scopes="/subscriptions/SUBSCRIPTION_ID"
```

This command will output 5 values:

```
{
  "appId": "000000000-0000-0000-0000-000000000000",
  "displayName": "azure-cli-2017-06-05-10-41-15",
  "name": "http://azure-cli-2017-06-05-10-41-15",
  "password": "0000-0000-0000-0000-000000000000",
  "tenant": "000000000-0000-0000-0000-000000000000"
}
```

These values map to the Terraform variables like so:

- `appId` is the `client_id` defined above.
- `password` is the `client_secret` defined above.
- `tenant` is the `tenant_id` defined above.

Finally, it's possible to test these values work as expected by first logging in:

```
$ az login --service-principal -u CLIENT_ID -p CLIENT_SECRET --tenant TENANT_ID
```


Once logged in as the Service Principal - we should be able to list the VM sizes by specifying an Azure region, for example here we use the `West US` region:

```
$ az vm list-sizes --location westus
```

Note: If you're using the **China**, **German** or **Government** Azure Clouds - you will need to switch `westus` out for another region. You can find which regions are available by running:

```
$ az account list-locations
```

Finally, since we're logged into the Azure CLI as a Service Principal we recommend logging out of the Azure CLI (but you can instead log in using your user account):

```
$ az logout
```

Information on how to configure the Provider block using the newly created Service Principal credentials can be found below.

Creating a Service Principal in the Azure Portal

There are three tasks necessary to create a Service Principal using the Azure Portal (<https://portal.azure.com>):

1. Create an Application in Azure Active Directory (which acts as a Service Principal)
2. Generating a Client Secret for the Azure Active Directory Application (which can be used for authentication)
3. Grant the Application access to manage resources in your Azure Subscription

1. Creating an Application in Azure Active Directory

Firstly navigate to the **Azure Active Directory** overview

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/Overview) within the Azure Portal - then select the **App Registration** blade

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) and click **Endpoints** at the top of the **App Registration** blade. A list of URIs will be displayed and you need to locate the URI for **OAUTH 2.0 AUTHORIZATION ENDPOINT** which contains a GUID. This GUID is your Tenant ID (the `tenant_id` field mentioned above).

Next, navigate back to the **App Registration** blade

(https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) - from here we'll create the Application in Azure Active Directory. To do this click **New application registration** at the top to add a new Application within Azure Active Directory. On this page, set the following values then press **Create**:

- **Name** - this is a friendly identifier and can be anything (e.g. "Terraform")
- **Application Type** - this should be set to "Web app / API"

- **Sign-on URL** - this can be anything, providing it's a valid URI (e.g. <https://terra.form> (<https://terra.form>))

At this point the newly created Azure Active Directory application should be visible on-screen - if it's not, navigate to the **App Registration** blade (https://portal.azure.com/#blade/Microsoft_AAD_IAM/ActiveDirectoryMenuBlade/RegisteredApps/RegisteredApps/Overview) and select the Azure Active Directory application. At the top of this page, the "Application ID" GUID is the `client_id` you'll need.

2. Generating a Client Secret for the Azure Active Directory Application

Now that the Azure Active Directory Application exists we can create a Client Secret which can be used for authentication - to do this select **Settings** and then **Keys**. This screen displays the Passwords (Client Secrets) and Public Keys (Client Certificates) which are associated with this Azure Active Directory Application.

On this screen we can generate a new Password by entering a Description and selecting an Expiry Date, and then pressing **Save**. Once the Password has been generated it will be displayed on screen - *the Password is only displayed once so be sure to copy it now* (otherwise you will need to regenerate a new key). This newly generated Password is the `client_secret` you will need.

3. Granting the Application access to manage resources in your Azure Subscription

Once the Application exists in Azure Active Directory - we can grant it permissions to modify resources in the Subscription. To do this, navigate to the **Subscriptions** blade within the Azure Portal (https://portal.azure.com/#blade/Microsoft_Azure_Billing/SubscriptionsBlade), then select the Subscription you wish to use, then click **Access Control (IAM)**, and finally **Add role assignment**.

Firstly, specify a Role which grants the appropriate permissions needed for the Service Principal (for example, **Contributor** will grant Read/Write on all resources in the Subscription). There's more information about the built in roles available here (<https://azure.microsoft.com/en-gb/documentation/articles/role-based-access-built-in-roles/>).

Secondly, search for and select the name of the Application created in Azure Active Directory to assign it this role - then press **Save**.

Configuring the Service Principal in Terraform

As we've obtained the credentials for this Service Principal - it's possible to configure them in a few different ways.

When storing the credentials as Environment Variables, for example:

```
$ export ARM_CLIENT_ID="00000000-0000-0000-0000-000000000000"
$ export ARM_CLIENT_SECRET="00000000-0000-0000-0000-000000000000"
$ export ARM_SUBSCRIPTION_ID="00000000-0000-0000-0000-000000000000"
$ export ARM_TENANT_ID="00000000-0000-0000-0000-000000000000"
```

The following Provider block can be specified - where `1.34.0` is the version of the Azure Provider that you'd like to use:

```
provider "azurerm" {  
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being  
  used  
  version = "=1.36.0"  
}
```

More information on the fields supported in the Provider block can be found here
([docs/providers/azurerm/index.html#argument-reference](/docs/providers/azurerm/index.html#argument-reference)).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Service Principal to authenticate.

It's also possible to configure these variables either in-line or from using variables in Terraform (as the `client_secret` is in this example), like so:

NOTE: We'd recommend not defining these variables in-line since they could easily be checked into Source Control.

```
variable "client_secret" {}  
  
provider "azurerm" {  
  # Whilst version is optional, we /strongly recommend/ using it to pin the version of the Provider being  
  used  
  version = "=1.36.0"  
  
  subscription_id = "00000000-0000-0000-0000-000000000000"  
  client_id       = "00000000-0000-0000-0000-000000000000"  
  client_secret   = "${var.client_secret}"  
  tenant_id       = "00000000-0000-0000-0000-000000000000"  
}
```

More information on the fields supported in the Provider block can be found here
([docs/providers/azurerm/index.html#argument-reference](/docs/providers/azurerm/index.html#argument-reference)).

At this point running either `terraform plan` or `terraform apply` should allow Terraform to run using the Service Principal to authenticate.

Data Source: azurerm_api_management_api

Use this data source to access information about an existing API Management API.

Example Usage

```
data "azurerm_api_management_api" "test" {
  name                = "search-api"
  api_management_name = "search-api-management"
  resource_group_name = "search-service"
  revision            = "2"
}

output "api_management_api_id" {
  value = "${data.azurerm_api_management_api.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the API Management API.
- `api_management_name` - (Required) The name of the API Management Service in which the API Management API exists.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists.
- `revision` - (Required) The Revision of the API Management API.

Attributes Reference

- `id` - The ID of the API Management API.
- `description` - A description of the API Management API, which may include HTML formatting tags.
- `display_name` - The display name of the API.
- `is_current` - Is this the current API Revision?
- `is_online` - Is this API Revision online/accessible via the Gateway?
- `path` - The Path for this API Management API.
- `protocols` - A list of protocols the operations in this API can be invoked.
- `service_url` - Absolute URL of the backend service implementing this API.
- `soap_pass_through` - Should this API expose a SOAP frontend, rather than a HTTP frontend?
- `subscription_key_parameter_names` - A `subscription_key_parameter_names` block as documented below.

- `version` - The Version number of this API, if this API is versioned.
 - `version_set_id` - The ID of the Version Set which this API is associated with.
-

A `subscription_key_parameter_names` block exports the following:

- `header` - The name of the HTTP Header which should be used for the Subscription Key.
 - `query` - The name of the QueryString parameter which should be used for the Subscription Key.
-

A `wsdl_selector` block exports the following:

- `service_name` - The name of service to import from WSDL.
- `endpoint_name` - The name of endpoint (port) to import from WSDL.

Data Source: azurerm_api_management_group

Use this data source to access information about an existing API Management Group.

Example Usage

```
data "azurerm_api_management_group" "test" {
  name                = "my-group"
  api_management_name = "example-apim"
  resource_group_name = "search-service"
}

output "group_type" {
  value = "${data.azurerm_api_management_group.test.type}"
}
```

Argument Reference

- `api_management_name` - (Required) The Name of the API Management Service in which this Group exists.
- `name` - (Required) The Name of the API Management Group.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists.

Attributes Reference

- `id` - The ID of the API Management Group.
- `display_name` - The display name of this API Management Group.
- `description` - The description of this API Management Group.
- `external_id` - The identifier of the external Group.
- `type` - The type of this API Management Group, such as `custom` or `external`.

Data Source: azurerm_api_management

Use this data source to access information about an existing API Management Service.

Example Usage

```
data "azurerm_api_management" "test" {
  name                = "search-api"
  resource_group_name = "search-service"
}

output "api_management_id" {
  value = "${data.azurerm_api_management.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the API Management service.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists.

Attributes Reference

- `id` - The ID of the API Management Service.
- `additional_location` - One or more `additional_location` blocks as defined below
- `location` - The Azure location where the API Management Service exists.
- `gateway_url` - The URL for the API Management Service's Gateway.
- `gateway_regional_url` - The URL for the Gateway in the Default Region.
- `hostname_configuration` - A `hostname_configuration` block as defined below.
- `management_api_url` - The URL for the Management API.
- `notification_sender_email` - The email address from which the notification will be sent.
- `portal_url` - The URL of the Publisher Portal.
- `public_ip_addresses` - The Public IP addresses of the API Management Service.
- `publisher_name` - The name of the Publisher/Company of the API Management Service.
- `publisher_email` - The email of Publisher/Company of the API Management Service.
- `scm_url` - The SCM (Source Code Management) endpoint.
- `sku` - A `sku` block as documented below.

- `tags` - A mapping of tags assigned to the resource.
-

A `additional_location` block exports the following:

- `location` - The location name of the additional region among Azure Data center regions.
 - `gateway_regional_url` - Gateway URL of the API Management service in the Region.
 - `public_ip_addresses` - Public Static Load Balanced IP addresses of the API Management service in the additional location. Available only for Basic, Standard and Premium SKU.
-

A `hostname_configuration` block exports the following:

- `management` - One or more `management` blocks as documented below.
 - `portal` - One or more `portal` blocks as documented below.
 - `proxy` - One or more `proxy` blocks as documented below.
 - `scm` - One or more `scm` blocks as documented below.
-

A `management` block exports the following:

- `host_name` - The Hostname used for the Management API.
 - `key_vault_id` - The ID of the Key Vault Secret which contains the SSL Certificate.
 - `negotiate_client_certificate` - Is Client Certificate Negotiation enabled?
-

A `portal` block exports the following:

- `host_name` - The Hostname used for the Portal.
 - `key_vault_id` - The ID of the Key Vault Secret which contains the SSL Certificate.
 - `negotiate_client_certificate` - Is Client Certificate Negotiation enabled?
-

A `proxy` block exports the following:

- `default_ssl_binding` - Is this the default SSL Binding?
 - `host_name` - The Hostname used for the Proxy.
 - `key_vault_id` - The ID of the Key Vault Secret which contains the SSL Certificate.
 - `negotiate_client_certificate` - Is Client Certificate Negotiation enabled?
-

A `scm` block exports the following:

- `host_name` - The Hostname used for the SCM URL.
- `key_vault_id` - The ID of the Key Vault Secret which contains the SSL Certificate.
- `negotiate_client_certificate` - Is Client Certificate Negotiation enabled?

A `sku` block exports the following:

- `name` - Specifies the plan's pricing tier.
- `capacity` - Specifies the number of units associated with this API Management service.

Data Source: azurerm_api_management_product

Use this data source to access information about an existing API Management Product.

Example Usage

```
data "azurerm_api_management_product" "test" {
  product_id           = "my-product"
  api_management_name = "example-apim"
  resource_group_name = "search-service"
}

output "product_terms" {
  value = "${data.azurerm_api_management_product.test.terms}"
}
```

Argument Reference

- `api_management_name` - (Required) The Name of the API Management Service in which this Product exists.
- `product_id` - (Required) The Identifier for the API Management Product.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists.

Attributes Reference

- `id` - The ID of the API Management Product.
- `approval_required` - Do subscribers need to be approved prior to being able to use the Product?
- `display_name` - The Display Name for this API Management Product.
- `published` - Is this Product Published?
- `subscription_required` - Is a Subscription required to access API's included in this Product?
- `description` - The description of this Product, which may include HTML formatting tags.
- `subscriptions_limit` - The number of subscriptions a user can have to this Product at the same time.
- `terms` - Any Terms and Conditions for this Product, which must be accepted by Developers before they can begin the Subscription process.

Data Source: azurerm_api_management_user

Use this data source to access information about an existing API Management User.

Example Usage

```
data "azurerm_api_management_user" "test" {
  user_id          = "my-user"
  api_management_name = "example-apim"
  resource_group_name = "search-service"
}

output "notes" {
  value = "${data.azurerm_api_management_user.test.notes}"
}
```

Argument Reference

- `api_management_name` - (Required) The Name of the API Management Service in which this User exists.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists.
- `user_id` - (Required) The Identifier for the User.

Attributes Reference

- `id` - The ID of the API Management User.
- `first_name` - The First Name for the User.
- `last_name` - The Last Name for the User.
- `email` - The Email Address used for this User.
- `note` - Any notes about this User.
- `state` - The current state of this User, for example `active`, `blocked` or `pending`.

Data Source: azurerm_application_insights

Use this data source to access information about an existing Application Insights component.

Example Usage

```
data "azurerm_application_insights" "test" {
  name                = "production"
  resource_group_name = "networking"
}

output "application_insights_instrumentation_key" {
  value = "${data.azurerm_application_insights.test.instrumentation_key}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Application Insights component.
- `resource_group_name` - (Required) Specifies the name of the resource group the Application Insights component is located in.

Attributes Reference

- `id` - The ID of the Virtual Machine.
- `app_id` - The App ID associated with this Application Insights component.
- `application_type` - The type of the component.
- `instrumentation_key` - The instrumentation key of the Application Insights component.
- `location` - The Azure location where the component exists.
- `tags` - Tags applied to the component.

Data Source: azurerm_application_security_group

Use this data source to access information about an existing Application Security Group.

Example Usage

```
data "azurerm_application_security_group" "test" {
  name                = "tf-appsecuritygroup"
  resource_group_name = "my-resource-group"
}

output "application_security_group_id" {
  value = "${data.azurerm_application_security_group.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - The name of the Application Security Group.
- `resource_group_name` - The name of the resource group in which the Application Security Group exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Security Group.
- `location` - The supported Azure location where the Application Security Group exists.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_app_service_certificate

Use this data source to access information about an App Service certificate.

Example Usage

```
data "azurerm_app_service_certificate" "example" {
  name                = "example-app-service-certificate"
  resource_group_name = "example-rg"
}

output "app_service_certificate_id" {
  value = "${data.azurerm_app_service_certificate.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the certificate.
- `resource_group_name` - (Required) The name of the resource group in which to create the certificate.

Attributes Reference

The following attributes are exported:

- `id` - The App Service certificate ID.
- `friendly_name` - The friendly name of the certificate.
- `subject_name` - The subject name of the certificate.
- `host_names` - List of host names the certificate applies to.
- `issuer` - The name of the certificate issuer.
- `issue_date` - The issue date for the certificate.
- `expiration_date` - The expiration date for the certificate.
- `thumbprint` - The thumbprint for the certificate.

Data Source: azurerm_app_service_certificate_order

Use this data source to access information about an existing App Service Certificate Order.

Example Usage

```
data "azurerm_app_service_certificate_order" "test" {
  name                = "example-cert-order"
  resource_group_name = "example-resources"
}

output "certificate_order_id" {
  value = "${data.azurerm_app_service_certificate_order.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the App Service.
- `resource_group_name` - (Required) The Name of the Resource Group where the App Service exists.

Attributes Reference

- `id` - The ID of the App Service.
- `location` - The Azure location where the App Service exists.
- `auto_renew` - true if the certificate should be automatically renewed when it expires; otherwise, false.
- `certificates` - State of the Key Vault secret. A `certificates` block as defined below.
- `csr` - Last CSR that was created for this order.
- `distinguished_name` - The Distinguished Name for the App Service Certificate Order.
- `key_size` - Certificate key size.
- `product_type` - Certificate product type, such as `Standard` or `Wildcard`.
- `validity_in_years` - Duration in years (must be between 1 and 3).
- `domain_verification_token` - Domain verification token.
- `status` - Current order status.
- `expiration_time` - Certificate expiration time.
- `is_private_key_external` - Whether the private key is external or not.

- `app_service_certificate_not_renewable_reasons` - Reasons why App Service Certificate is not renewable at the current moment.
 - `signed_certificate_thumbprint` - Certificate thumbprint for signed certificate.
 - `root_thumbprint` - Certificate thumbprint for root certificate.
 - `intermediate_thumbprint` - Certificate thumbprint intermediate certificate.
 - `tags` - A mapping of tags to assign to the resource.
-

- `certificate_name` - The name of the App Service Certificate.
- `key_vault_id` - Key Vault resource Id.
- `key_vault_secret_name` - Key Vault secret name.
- `provisioning_state` - Status of the Key Vault secret.

Data Source: azurerm_app_service

Use this data source to access information about an existing App Service.

Example Usage

```
data "azurerm_app_service" "test" {
  name                = "search-app-service"
  resource_group_name = "search-service"
}

output "app_service_id" {
  value = "${data.azurerm_app_service.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the App Service.
- `resource_group_name` - (Required) The Name of the Resource Group where the App Service exists.

Attributes Reference

- `id` - The ID of the App Service.
- `location` - The Azure location where the App Service exists.
- `app_service_plan_id` - The ID of the App Service Plan within which the App Service exists.
- `app_settings` - A key-value pair of App Settings for the App Service.
- `connection_string` - An `connection_string` block as defined below.
- `client_affinity_enabled` - Does the App Service send session affinity cookies, which route client requests in the same session to the same instance?
- `enabled` - Is the App Service Enabled?
- `https_only` - Can the App Service only be accessed via HTTPS?
- `client_cert_enabled` - Does the App Service require client certificates for incoming requests?
- `site_config` - A `site_config` block as defined below.
- `tags` - A mapping of tags to assign to the resource.
- `default_site_hostname` - The Default Hostname associated with the App Service - such as `mysite.azurewebsites.net`
- `outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12`

- `possible_outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12,52.143.43.17` - not all of which are necessarily in use. Superset of `outbound_ip_addresses`.
-

`connection_string` supports the following:

- `name` - The name of the Connection String.
 - `type` - The type of the Connection String.
 - `value` - The value for the Connection String.
-

A `cors` block exports the following:

- `allowed_origins` - A list of origins which are able to make cross-origin calls.
 - `support_credentials` - Are credentials supported?
-

A `ip_restriction` block exports the following:

- `ip_address` - The IP Address used for this IP Restriction.
 - `subnet_mask` - The Subnet mask used for this IP Restriction.
-

`site_config` supports the following:

- `always_on` - Is the app be loaded at all times?
- `app_command_line` - App command line to launch.
- `cors` - A `cors` block as defined above.
- `default_documents` - The ordering of default documents to load, if an address isn't specified.
- `dotnet_framework_version` - The version of the .net framework's CLR used in this App Service.
- `http2_enabled` - Is HTTP2 Enabled on this App Service?
- `ftp_state` - State of FTP / FTPS service for this AppService.
- `ip_restriction` - One or more `ip_restriction` blocks as defined above.
- `java_version` - The version of Java in use.
- `java_container` - The Java Container in use.
- `java_container_version` - The version of the Java Container in use.
- `linux_fx_version` - Linux App Framework and version for the AppService.
- `windows_fx_version` - Windows Container Docker Image for the AppService.
- `local_mysql_enabled` - Is "MySQL In App" Enabled? This runs a local MySQL instance with your app and shares resources from the App Service plan.

- `managed_pipeline_mode` - The Managed Pipeline Mode used in this App Service.
- `min_tls_version` - The minimum supported TLS version for this App Service.
- `php_version` - The version of PHP used in this App Service.
- `python_version` - The version of Python used in this App Service.
- `remote_debugging_enabled` - Is Remote Debugging Enabled in this App Service?
- `remote_debugging_version` - Which version of Visual Studio is the Remote Debugger compatible with?
- `scm_type` - The type of Source Control enabled for this App Service.
- `use_32_bit_worker_process` - Does the App Service run in 32 bit mode, rather than 64 bit mode?
- `websockets_enabled` - Are WebSockets enabled for this App Service?
- `virtual_network_name` - The name of the Virtual Network which this App Service is attached to.

Data Source: azurerm_app_service_plan

Use this data source to access information about an existing App Service Plan (formerly known as a `Server Farm`).

Example Usage

```
data "azurerm_app_service_plan" "test" {
  name                = "search-app-service-plan"
  resource_group_name = "search-service"
}

output "app_service_plan_id" {
  value = "${data.azurerm_app_service_plan.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the App Service Plan.
- `resource_group_name` - (Required) The Name of the Resource Group where the App Service Plan exists.

Attributes Reference

- `id` - The ID of the App Service Plan.
- `location` - The Azure location where the App Service Plan exists
- `kind` - The Operating System type of the App Service Plan
- `sku` - A `sku` block as documented below.
- `properties` - A `properties` block as documented below.
- `tags` - A mapping of tags assigned to the resource.
- `maximum_elastic_worker_count` - The maximum number of total workers allowed for this ElasticScaleEnabled App Service Plan.
- `is_xenon` - A flag that indicates if it's a xenon plan (support for Windows Container)
- `maximum_number_of_workers` - The maximum number of workers supported with the App Service Plan's sku.

A `sku` block supports the following:

- `tier` - Specifies the plan's pricing tier.
- `size` - Specifies the plan's instance size.
- `capacity` - Specifies the number of workers associated with this App Service Plan.

A `properties` block supports the following:

- `app_service_environment_id` - The ID of the App Service Environment where the App Service Plan is located.
- `maximum_number_of_workers` - Maximum number of instances that can be assigned to this App Service plan.
- `reserved` - Is this App Service Plan Reserved ?
- `per_site_scaling` - Can Apps assigned to this App Service Plan be scaled independently?

Data Source:

azurerm_automation_account_registration_info

Use this data source to access information about an existing Automation Account Registration Information.

Example Usage

```
data "azurerm_automation_account_registration_info" "test" {
  name           = "automation-account"
  resource_group_name = "automation-resource-group"
}

output "automation_account_id" {
  value = "${data.azurerm_automation_account_registration_info.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the Automation Account.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the Automation Account exists.

Attributes Reference

- `id` - The ID of the Automation Account
- `primary_key` - The primary key for the Automation Account Registration information
- `secondary_key` - The primary key for the Automation Account Registration information
- `endpoint` - The Assigned Automation Account Registration endpoint

Data Source: azurerm_automation_variable_bool

Use this data source to access information about an existing Automation Bool Variable.

Example Usage

```
data "azurerm_automation_variable_bool" "example" {
  name                = "tfex-example-var"
  resource_group_name = "tfex-example-rg"
  automation_account_name = "tfex-example-account"
}

output "variable_id" {
  value = "${data.azurerm_automation_variable_bool.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable.
- `resource_group_name` - (Required) The Name of the Resource Group where the automation account exists.
- `automation_account_name` - (Required) The name of the automation account in which the Automation Variable exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.
- `description` - The description of the Automation Variable.
- `encrypted` - Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - The value of the Automation Variable as a `boolean`.

Data Source: azurerm_automation_variable_datetime

Use this data source to access information about an existing Automation Datetime Variable.

Example Usage

```
data "azurerm_automation_variable_datetime" "example" {
  name                = "tfex-example-var"
  resource_group_name = "tfex-example-rg"
  automation_account_name = "tfex-example-account"
}

output "variable_id" {
  value = "${data.azurerm_automation_variable_datetime.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable.
- `resource_group_name` - (Required) The Name of the Resource Group where the automation account exists.
- `automation_account_name` - (Required) The name of the automation account in which the Automation Variable exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.
- `description` - The description of the Automation Variable.
- `encrypted` - Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - The value of the Automation Variable in the RFC3339 Section 5.6 Internet Date/Time Format (<https://tools.ietf.org/html/rfc3339#section-5.6>).

Data Source: azurerm_automation_variable_int

Use this data source to access information about an existing Automation Int Variable.

Example Usage

```
data "azurerm_automation_variable_int" "example" {
  name                = "tfex-example-var"
  resource_group_name = "tfex-example-rg"
  automation_account_name = "tfex-example-account"
}

output "variable_id" {
  value = "${data.azurerm_automation_variable_int.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable.
- `resource_group_name` - (Required) The Name of the Resource Group where the automation account exists.
- `automation_account_name` - (Required) The name of the automation account in which the Automation Variable exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.
- `description` - The description of the Automation Variable.
- `encrypted` - Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - The value of the Automation Variable as a `integer`.

Data Source: azurerm_automation_variable_string

Use this data source to access information about an existing Automation String Variable.

Example Usage

```
data "azurerm_automation_variable_string" "example" {
  name                = "tfex-example-var"
  resource_group_name = "tfex-example-rg"
  automation_account_name = "tfex-example-account"
}

output "variable_id" {
  value = "${data.azurerm_automation_variable_string.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable.
- `resource_group_name` - (Required) The Name of the Resource Group where the automation account exists.
- `automation_account_name` - (Required) The name of the automation account in which the Automation Variable exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.
- `description` - The description of the Automation Variable.
- `encrypted` - Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - The value of the Automation Variable as a `string`.

Data Source: azurerm_availability_set

Use this data source to access information about an existing Availability Set.

Example Usage

```
data "azurerm_availability_set" "test" {
  name                = "tf-appsecuritygroup"
  resource_group_name = "my-resource-group"
}

output "availability_set_id" {
  value = "${data.azurerm_availability_set.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - The name of the Availability Set.
- `resource_group_name` - The name of the resource group in which the Availability Set exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Availability Set.
- `location` - The supported Azure location where the Availability Set exists.
- `managed` - Whether the availability set is managed or not.
- `platform_fault_domain_count` - The number of fault domains that are used.
- `platform_update_domain_count` - The number of update domains that are used.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_azuread_application

Use this data source to access information about an existing Application within Azure Active Directory.

NOTE: The Azure Active Directory resources have been split out into a new AzureAD Provider (<http://terraform.io/docs/providers/azuread/index.html>) - as such the AzureAD resources within the AzureRM Provider are deprecated and will be removed in the next major version (2.0). Information on how to migrate from the existing resources to the new AzureAD Provider can be found here (</docs/providers/azurerm/guides/migrating-to-azuread.html>).

NOTE: If you're authenticating using a Service Principal then it must have permissions to both `Read` and `write` all applications and `Sign in and read user profile` within the Windows Azure Active Directory API.

Example Usage

```
data "azurerm_azuread_application" "test" {
  name = "My First AzureAD Application"
}

output "azure_active_directory_object_id" {
  value = "${data.azurerm_azuread_application.test.id}"
}
```

Argument Reference

- `object_id` - (Optional) Specifies the Object ID of the Application within Azure Active Directory.
- `name` - (Optional) Specifies the name of the Application within Azure Active Directory.

NOTE: Either an `object_id` or `name` must be specified.

Attributes Reference

- `id` - the Object ID of the Azure Active Directory Application.
- `application_id` - the Application ID of the Azure Active Directory Application.
- `available_to_other_tenants` - Is this Azure AD Application available to other tenants?
- `identifier_uris` - A list of user-defined URI(s) that uniquely identify a Web application within it's Azure AD tenant, or within a verified custom domain if the application is multi-tenant.
- `oauth2_allow_implicit_flow` - Does this Azure AD Application allow OAuth2.0 implicit flow tokens?

- `object_id` - the Object ID of the Azure Active Directory Application.
- `reply_urls` - A list of URLs that user tokens are sent to for sign in, or the redirect URIs that OAuth 2.0 authorization codes and access tokens are sent to.

Data Source: azurerm_azuread_service_principal

Gets information about an existing Service Principal associated with an Application within Azure Active Directory.

NOTE: The Azure Active Directory resources have been split out into a new AzureAD Provider (<http://terraform.io/docs/providers/azuread/index.html>) - as such the AzureAD resources within the AzureRM Provider are deprecated and will be removed in the next major version (2.0). Information on how to migrate from the existing resources to the new AzureAD Provider can be found here (</docs/providers/azurerm/guides/migrating-to-azuread.html>).

NOTE: If you're authenticating using a Service Principal then it must have permissions to both `Read` and `write` all applications and `Sign in and read user profile` within the Windows Azure Active Directory API.

Example Usage (by Application Display Name)

```
data "azurerm_azuread_service_principal" "test" {
  display_name = "my-awesome-application"
}
```

Example Usage (by Application ID)

```
data "azurerm_azuread_service_principal" "test" {
  application_id = "00000000-0000-0000-0000-000000000000"
}
```

Example Usage (by Object ID)

```
data "azurerm_azuread_service_principal" "test" {
  object_id = "00000000-0000-0000-0000-000000000000"
}
```

Argument Reference

The following arguments are supported:

- `application_id` - (Optional) The ID of the Azure AD Application for which to create a Service Principal.
- `object_id` - (Optional) The ID of the Azure AD Service Principal.
- `display_name` - (Optional) The Display Name of the Azure AD Application associated with this Service Principal.

NOTE: At least one of `application_id`, `display_name` or `object_id` must be specified.

Attributes Reference

The following attributes are exported:

- `id` - The Object ID for the Service Principal.

Data Source: azurerm_batch_account

Use this data source to access information about an existing Batch Account.

Example Usage

```
data "azurerm_batch_account" "test" {
  name                = "testbatchaccount"
  resource_group_name = "test"
}

output "pool_allocation_mode" {
  value = "${data.azurerm_batch_account.test.pool_allocation_mode}"
}
```

Argument Reference

- `name` - (Required) The name of the Batch account.
- `resource_group_name` - (Required) The Name of the Resource Group where this Batch account exists.

Attributes Reference

The following attributes are exported:

- `id` - The Batch account ID.
- `name` - The Batch account name.
- `location` - The Azure Region in which this Batch account exists.
- `pool_allocation_mode` - The pool allocation mode configured for this Batch account.
- `storage_account_id` - The ID of the Storage Account used for this Batch account.
- `primary_access_key` - The Batch account primary access key.
- `secondary_access_key` - The Batch account secondary access key.
- `account_endpoint` - The account endpoint used to interact with the Batch service.
- `key_vault_reference` - The `key_vault_reference` block that describes the Azure KeyVault reference to use when deploying the Azure Batch account using the `UserSubscription` pool allocation mode.
- `tags` - A map of tags assigned to the Batch account.

NOTE: Primary and secondary access keys are only available when `pool_allocation_mode` is set to `BatchService`. See documentation (<https://docs.microsoft.com/en-us/azure/batch/batch-api-basics>) for more information.

A `key_vault_reference` block have the following properties:

- `id` - The Azure identifier of the Azure KeyVault reference.
 - `url` - The HTTPS URL of the Azure KeyVault reference.
-

Data Source: azurerm_batch_certificate

Use this data source to access information about an existing certificate in a Batch Account.

Example Usage

```
data "azurerm_batch_certificate" "example" {
  name                = "SHA1-42C107874FD0E4A9583292A2F1098E8FE4B2EDDA"
  account_name        = "examplebatchaccount"
  resource_group_name = "example"
}

output "thumbprint" {
  value = "${data.azurerm_batch_certificate.example.thumbprint}"
}
```

Argument Reference

- `name` - (Required) The name of the Batch certificate.
- `account_name` - (Required) The name of the Batch account.
- `resource_group_name` - (Required) The Name of the Resource Group where this Batch account exists.

Attributes Reference

The following attributes are exported:

- `id` - The Batch certificate ID.
- `public_data` - The public key of the certificate.
- `format` - The format of the certificate, such as `Cer` or `Pfx`.
- `thumbprint` - The thumbprint of the certificate.
- `thumbprint_algorithm` - The algorithm of the certificate thumbprint.

Data source: azurerm_batch_pool

Use this data source to access information about an existing Batch pool

Example Usage

```
data "azurerm_batch_pool" "test" {
  name           = "testbatchpool"
  account_name   = "testbatchaccount"
  resource_group_name = "test"
}
```

Attributes Reference

The following attributes are exported:

- `id` - The Batch pool ID.
- `name` - The name of the Batch pool.
- `account_name` - The name of the Batch account.
- `node_agent_sku_id` - The Sku of the node agents in the Batch pool.
- `vm_size` - The size of the VM created in the Batch pool.
- `fixed_scale` - A `fixed_scale` block that describes the scale settings when using fixed scale.
- `auto_scale` - A `auto_scale` block that describes the scale settings when using auto scale.
- `storage_image_reference` - The reference of the storage image used by the nodes in the Batch pool.
- `start_task` - A `start_task` block that describes the start task settings for the Batch pool.
- `max_tasks_per_node` - The maximum number of tasks that can run concurrently on a single compute node in the pool.
- `certificate` - One or more `certificate` blocks that describe the certificates installed on each compute node in the pool.
- `container_configuration` - The container configuration used in the pool's VMs.

A `fixed_scale` block exports the following:

- `target_dedicated_nodes` - The number of nodes in the Batch pool.
 - `target_low_priority_nodes` - The number of low priority nodes in the Batch pool.
 - `resize_timeout` - The timeout for resize operations.
-

A `auto_scale` block exports the following:

- `evaluation_interval` - The interval to wait before evaluating if the pool needs to be scaled.
 - `formula` - The autoscale formula that needs to be used for scaling the Batch pool.
-

A `start_task` block exports the following:

- `command_line` - The command line executed by the start task.
 - `max_task_retry_count` - The number of retry count.
 - `wait_for_success` - A flag that indicates if the Batch pool should wait for the start task to be completed.
 - `environment` - A map of strings (key,value) that represents the environment variables to set in the start task.
 - `user_identity` - A `user_identity` block that describes the user identity under which the start task runs.
 - `resource_file` - (Optional) One or more `resource_file` blocks that describe the files to be downloaded to a compute node.
-

A `user_identity` block exports the following:

- `user_name` - The username to be used by the Batch pool start task.
 - `auto_user` - A `auto_user` block that describes the user identity under which the start task runs.
-

A `auto_user` block exports the following:

- `elevation_level` - The elevation level of the user identity under which the start task runs.
 - `scope` - The scope of the user identity under which the start task runs.
-

A `certificate` block exports the following:

- `id` - The fully qualified ID of the certificate installed on the pool.
- `store_location` - The location of the certificate store on the compute node into which the certificate is installed, either `CurrentUser` or `LocalMachine`.

NOTE: This property is applicable only for pools configured with Windows nodes (that is, created with `cloudServiceConfiguration`, or with `virtualMachineConfiguration` using a Windows image reference). For Linux compute nodes, the certificates are stored in a directory inside the task working directory and an environment variable `AZ_BATCH_CERTIFICATES_DIR` is supplied to the task to query for this location. For certificates with visibility of 'remoteUser', a 'certs' directory is created in the user's home directory (e.g., `/home/{user-name}/certs`) and certificates are placed in that directory.

- `store_name` - The name of the certificate store on the compute node into which the certificate is installed.

NOTE: This property is applicable only for pools configured with Windows nodes (that is, created with `cloudServiceConfiguration`, or with `virtualMachineConfiguration` using a Windows image reference).

- `visibility` - Which user accounts on the compute node have access to the private data of the certificate.
-

A `resource_file` block exports the following:

- `auto_storage_container_name` - The storage container name in the auto storage account.
 - `blob_prefix` - The blob prefix used when downloading blobs from an Azure Storage container.
 - `file_mode` - The file permission mode attribute represented as a string in octal format (e.g. `"0644"`).
 - `file_path` - The location on the compute node to which to download the file, relative to the task's working directory. If the `http_url` property is specified, the `file_path` is required and describes the path which the file will be downloaded to, including the filename. Otherwise, if the `auto_storage_container_name` or `storage_container_url` property is specified.
 - `http_url` - The URL of the file to download. If the URL is Azure Blob Storage, it must be readable using anonymous access.
 - `storage_container_url` - The URL of the blob container within Azure Blob Storage.
-

A `container_configuration` block exports the following:

- `type` - The type of container configuration.
 - `container_registries` - Additional container registries from which container images can be pulled by the pool's VMs.
-

A `container_registries` block exports the following:

- `registry_server` - The container registry URL. The default is `"docker.io"`.
- `user_name` - The user name to log into the registry server.
- `password` - The password to log into the registry server.

Data Source: azurerm_builtin_role_definition

Use this data source to access information about a built-in Role Definition. To access information about a custom Role Definition, please see the `azurerm_role_definition` data source ([/docs/providers/azurerm/d/role_definition.html](/docs/providers/azurerm/d/role_definition.html)) instead.

NOTE: This data source has been deprecated in favour of `azurerm_role_definition` that now can look up role definitions by name. As such this data source will be removed in version 2.0 of the AzureRM Provider.

Example Usage

```
data "azurerm_builtin_role_definition" "contributor" {
  name = "Contributor"
}

output "contributor_role_definition_id" {
  value = "${data.azurerm_builtin_role_definition.contributor.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the built-in Role Definition. Possible values are: `Contributor`, `Owner`, `Reader` and `VirtualMachineContributor`.

Attributes Reference

- `id` - the ID of the built-in Role Definition.
- `description` - the Description of the built-in Role.
- `type` - the Type of the Role.
- `permissions` - a `permissions` block as documented below.
- `assignable_scopes` - One or more assignable scopes for this Role Definition, such as `/subscriptions/0b1f6471-1bf0-4dda-aec3-11112223333`, `/subscriptions/0b1f6471-1bf0-4dda-aec3-11112223333/resourceGroups/myGroup`, or `/subscriptions/0b1f6471-1bf0-4dda-aec3-11112223333/resourceGroups/myGroup/providers/Microsoft.Compute/virtualMachines/myVM`.

A `permissions` block contains:

- `actions` - a list of actions supported by this role
- `data_actions` - a list of data actions supported by this role
- `not_actions` - a list of actions which are denied by this role

- `not_data_actions` - a list of data actions which are denied by this role

Data Source: azurerm_cdn_profile

Use this data source to access information about an existing CDN Profile.

Example Usage

```
data "azurerm_cdn_profile" "test" {
  name                = "myfirstcdnprofile"
  resource_group_name = "example-resources"
}

output "cdn_profile_id" {
  value = "${data.azurerm_cdn_profile.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the CDN Profile.
- `resource_group_name` - (Required) The name of the resource group in which the CDN Profile exists.

Attributes Reference

- `location` - The Azure Region where the resource exists.
- `sku` - The pricing related information of current CDN profile.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_client_config

Use this data source to access the configuration of the AzureRM provider.

Example Usage

```
data "azurerm_client_config" "current" {}

output "account_id" {
  value = "${data.azurerm_client_config.current.service_principal_application_id}"
}
```

Argument Reference

There are no arguments available for this data source.

Attributes Reference

- `client_id` is set to the Azure Client ID (Application Object ID).
- `tenant_id` is set to the Azure Tenant ID.
- `subscription_id` is set to the Azure Subscription ID.
- `object_id` is set to the Azure Object ID.

Note: the following fields are only available when authenticating via a Service Principal (as opposed to using the Azure CLI) and have been deprecated:

- `service_principal_application_id` is the Service Principal Application ID (same as `client_id`).
- `service_principal_object_id` is the Service Principal Object ID (now available via `object_id`).

Note: To better understand "application" and "service principal", please read Application and service principal objects in Azure Active Directory (<https://docs.microsoft.com/en-us/azure/active-directory/develop/active-directory-application-objects>).

Data Source: azurerm_container_registry

Use this data source to access information about an existing Container Registry.

Example Usage

```
data "azurerm_container_registry" "test" {
  name                = "testacr"
  resource_group_name = "test"
}

output "login_server" {
  value = "${data.azurerm_container_registry.test.login_server}"
}
```

Argument Reference

- `name` - (Required) The name of the Container Registry.
- `resource_group_name` - (Required) The Name of the Resource Group where this Container Registry exists.

Attributes Reference

The following attributes are exported:

- `id` - The Container Registry ID.
- `login_server` - The URL that can be used to log into the container registry.
- `admin_username` - The Username associated with the Container Registry Admin account - if the admin account is enabled.
- `admin_password` - The Password associated with the Container Registry Admin account - if the admin account is enabled.
- `location` - The Azure Region in which this Container Registry exists.
- `admin_enabled` - Is the Administrator account enabled for this Container Registry.
- `sku` - The SKU of this Container Registry, such as `Basic`.
- `storage_account_id` - The ID of the Storage Account used for this Container Registry. This is only returned for `Classic` SKU's.
- `tags` - A map of tags assigned to the Container Registry.

Data Source: azurerm_cosmosdb_account

Use this data source to access information about an existing CosmosDB (formally DocumentDB) Account.

Example Usage

```
data "azurerm_cosmosdb_account" "test" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

output "cosmosdb_account_endpoint" {
  value = "${data.azurerm_cosmosdb_account.jobs.endpoint}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the CosmosDB Account.
- `resource_group_name` - (Required) Specifies the name of the resource group in which the CosmosDB Account resides.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the CosmosDB Account.
- `location` - The Azure location where the resource exists.
- `tags` - A mapping of tags assigned to the resource.
- `offer_type` - The Offer Type to used by this CosmosDB Account.
- `kind` - The Kind of the CosmosDB account.
- `ip_range_filter` - The current IP Filter for this CosmosDB account
- `enable_automatic_failover` - If automatic failover is enabled for this CosmosDB Account.
- `capabilities` - Capabilities enabled on this Cosmos DB account.
- `is_virtual_network_filter_enabled` - If virtual network filtering is enabled for this Cosmos DB account.
- `virtual_network_rule` - Subnets that are allowed to access this CosmosDB account.
- `enable_multiple_write_locations` - If multi-master is enabled for this Cosmos DB account.

`consistency_policy` The current consistency Settings for this CosmosDB account with the following properties:

- `consistency_level` - The Consistency Level used by this CosmosDB Account.
- `max_interval_in_seconds` - The amount of staleness (in seconds) tolerated when the consistency level is Bounded Staleness.
- `max_staleness_prefix` - The number of stale requests tolerated when the consistency level is Bounded Staleness.

`geo_location` The geographic locations data is replicated to with the following properties:

- `id` - The ID of the location.
- `location` - The name of the Azure region hosting replicated data.
- `priority` - The locations fail over priority.

`virtual_network_rule` The virtual network subnets allowed to access this Cosmos DB account with the following properties:

- `id` - The ID of the virtual network subnet.
- `endpoint` - The endpoint used to connect to the CosmosDB account.
- `read_endpoints` - A list of read endpoints available for this CosmosDB account.
- `write_endpoints` - A list of write endpoints available for this CosmosDB account.
- `primary_master_key` - The Primary master key for the CosmosDB Account.
- `secondary_master_key` - The Secondary master key for the CosmosDB Account.
- `primary_readonly_master_key` - The Primary read-only master Key for the CosmosDB Account.
- `secondary_readonly_master_key` - The Secondary read-only master key for the CosmosDB Account.

Data Source: azurerm_data_factory

Use this data source to access information about an existing Azure Data Factory (Version 2).

Example Usage

```
data "azurerm_data_factory" "example" {
  name                = "${azurerm_data_factory.example.name}"
  resource_group_name = "${azurerm_data_factory.example.resource_group_name}"
}

output "data_factory_id" {
  value = "${azurerm_data_factory.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Factory to retrieve information about.
- `resource_group_name` - (Required) The name of the resource group where the Data Factory exists.

Attributes Reference

The following attributes are exported:

- `id` - The Data Factory ID.
- `location` - The Azure location where the resource exists.
- `github_configuration` - A `github_configuration` block as defined below.
- `identity` - An `identity` block as defined below.
- `vsts_configuration` - A `vsts_configuration` block as defined below.

* `tags` - A mapping of tags assigned to the resource.

A `github_configuration` block exports the following:

- `account_name` - The GitHub account name.
- `branch_name` - The branch of the repository to get code from.

- `git_url` - The GitHub Enterprise host name.
 - `repository_name` - The name of the git repository.
 - `root_folder` - The root folder within the repository.
-

An `identity` block exports the following:

- `principal_id` - The ID of the Principal (Client) in Azure Active Directory.
 - `tenant_id` - The ID of the Azure Active Directory Tenant.
 - `type` - The identity type of the Data Factory.
-

A `vsts_configuration` block exports the following:

- `account_name` - The VSTS account name.
- `branch_name` - The branch of the repository to get code from.
- `project_name` - The name of the VSTS project.
- `repository_name` - The name of the git repository.
- `root_folder` - The root folder within the repository.
- `tenant_id` - The Tenant ID associated with the VSTS account.

Data Source: azurerm_data_lake_store

Use this data source to access information about an existing Data Lake Store.

Example Usage

```
data "azurerm_data_lake_store" "test" {
  name                = "testdatalake"
  resource_group_name = "testdatalake"
}

output "data_lake_store_id" {
  value = "${data.azurerm_data_lake_store.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the Data Lake Store.
- `resource_group_name` - (Required) The Name of the Resource Group where the Data Lake Store exists.

Attributes Reference

- `id` - The ID of the Data Lake Store.
- `encryption_state` - the Encryption State of this Data Lake Store Account, such as `Enabled` or `Disabled`.
- `encryption_type` - the Encryption Type used for this Data Lake Store Account.
- `firewall_allow_azure_ips` - are Azure Service IP's allowed through the firewall?
- `firewall_state` - the state of the firewall, such as `Enabled` or `Disabled`.
- `tier` - Current monthly commitment tier for the account.
- `tags` - A mapping of tags to assign to the Data Lake Store.

Data Source: azurerm_dev_test_lab

Use this data source to access information about an existing Dev Test Lab.

Example Usage

```
data "azurerm_dev_test_lab" "test" {
  name                = "example-lab"
  resource_group_name = "example-resources"
}

output "unique_identifier" {
  value = "${data.azurerm_dev_test_lab.test.unique_identifier}"
}
```

Argument Reference

- `name` - (Required) The name of the Dev Test Lab.
- `resource_group_name` - (Required) The Name of the Resource Group where the Dev Test Lab exists.

Attributes Reference

- `id` - The ID of the Dev Test Lab.
- `artifacts_storage_account_id` - The ID of the Storage Account used for Artifact Storage.
- `default_storage_account_id` - The ID of the Default Storage Account for this Dev Test Lab.
- `default_premium_storage_account_id` - The ID of the Default Premium Storage Account for this Dev Test Lab.
- `key_vault_id` - The ID of the Key used for this Dev Test Lab.
- `location` - The Azure location where the Dev Test Lab exists.
- `premium_data_disk_storage_account_id` - The ID of the Storage Account used for Storage of Premium Data Disk.
- `storage_type` - The type of storage used by the Dev Test Lab.
- `tags` - A mapping of tags to assign to the resource.
- `unique_identifier` - The unique immutable identifier of the Dev Test Lab.

Data Source: azurerm_dev_test_virtual_network

Use this data source to access information about an existing Dev Test Lab Virtual Network.

Example Usage

```
data "azurerm_dev_test_virtual_network" "test" {
  name           = "example-network"
  lab_name       = "examplelab"
  resource_group_name = "example-resource"
}

output "lab_subnet_name" {
  value = "${data.azurerm_dev_test_virtual_network.test.allowed_subnets.0.lab_subnet_name}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Virtual Network.
- `lab_name` - (Required) Specifies the name of the Dev Test Lab.
- `resource_group_name` - (Required) Specifies the name of the resource group that contains the Virtual Network.

Attributes Reference

- `allowed_subnets` - The list of subnets enabled for the virtual network as defined below.
- `subnet_overrides` - The list of permission overrides for the subnets as defined below.
- `unique_identifier` - The unique immutable identifier of the virtual network.

An `allowed_subnets` block supports the following:

- `allow_public_ip` - Indicates if this subnet allows public IP addresses. Possible values are `Allow`, `Default` and `Deny`.
- `lab_subnet_name` - The name of the subnet.
- `resource_id` - The resource identifier for the subnet.

An `subnets_override` block supports the following:

- `lab_subnet_name` - The name of the subnet.
- `resource_id` - The resource identifier for the subnet.

- `use_in_vm_creation_permission` - Indicates if the subnet can be used for VM creation. Possible values are `Allow`, `Default` and `Deny`.
- `use_public_ip_permission` - Indicates if the subnet can be assigned public IP addresses. Possible values are `Allow`, `Default` and `Deny`.
- `virtual_network_pool_name` - The virtual network pool associated with this subnet.

Data Source: azurerm_dns_zone

Use this data source to access information about an existing DNS Zone.

Example Usage

```
data "azurerm_dns_zone" "test" {
  name                = "search-eventhubns"
  resource_group_name = "search-service"
}

output "dns_zone_id" {
  value = "${data.azurerm_dns_zone.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the DNS Zone.
- `resource_group_name` - (Optional) The Name of the Resource Group where the DNS Zone exists. If the Name of the Resource Group is not provided, the first DNS Zone from the list of DNS Zones in your subscription that matches `name` will be returned.

Attributes Reference

- `id` - The ID of the DNS Zone.
- `max_number_of_record_sets` - Maximum number of Records in the zone.
- `number_of_record_sets` - The number of records already in the zone.
- `name_servers` - A list of values that make up the NS record for the zone.
- `registration_virtual_network_ids` - A list of Virtual Network ID's that register hostnames in this DNS zone.
- `resolution_virtual_network_ids` - A list of Virtual Network ID's that resolve records in this DNS zone.
- `tags` - A mapping of tags to assign to the EventHub Namespace.
- `zone_type` - (**Deprecated**) The type of this DNS zone, such as `Public` or `Private`.

NOTE: This field is deprecated since Private DNS is now a separate resource and will be removed in 2.0 of the Azure Provider.

Data Source: azurerm_eventhub_namespace

Use this data source to access information about an existing EventHub Namespace.

Example Usage

```
data "azurerm_eventhub_namespace" "test" {
  name                = "search-eventhubns"
  resource_group_name = "search-service"
}

output "eventhub_namespace_id" {
  value = "${data.azurerm_eventhub_namespace.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the EventHub Namespace.
- `resource_group_name` - (Required) The Name of the Resource Group where the EventHub Namespace exists.

Attributes Reference

- `id` - The ID of the EventHub Namespace.
- `location` - The Azure location where the EventHub Namespace exists
- `sku` - Defines which tier to use.
- `capacity` - The Capacity / Throughput Units for a Standard SKU namespace.
- `auto_inflate_enabled` - Is Auto Inflate enabled for the EventHub Namespace?
- `maximum_throughput_units` - Specifies the maximum number of throughput units when Auto Inflate is Enabled.
- `tags` - A mapping of tags to assign to the EventHub Namespace.

The following attributes are exported only if there is an authorization rule named `RootManageSharedAccessKey` which is created automatically by Azure.

- `default_primary_connection_string` - The primary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_connection_string` - The secondary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_primary_key` - The primary access key for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_key` - The secondary access key for the authorization rule `RootManageSharedAccessKey` .

Data Source: azurerm_express_route_circuit

Use this data source to access information about an existing ExpressRoute circuit.

Example Usage

```
data "azurerm_express_route_circuit" test {
  resource_group_name = "${azurerm_resource_group.test.name}"
  name                = "${azurerm_express_route_circuit.test.name}"
}

output "express_route_circuit_id" {
  value = "${data.azurerm_express_route_circuit.test.id}"
}

output "service_key" {
  value = "${data.azurerm_express_route_circuit.test.service_key}"
}
```

Argument Reference

- `name` - (Required) The name of the ExpressRoute circuit.
- `resource_group_name` - (Required) The Name of the Resource Group where the ExpressRoute circuit exists.

Attributes Reference

- `id` - The ID of the ExpressRoute circuit.
- `location` - The Azure location where the ExpressRoute circuit exists
- `peerings` - A `peerings` block for the ExpressRoute circuit as documented below
- `service_provider_provisioning_state` - The ExpressRoute circuit provisioning state from your chosen service provider. Possible values are "NotProvisioned", "Provisioning", "Provisioned", and "Deprovisioning".
- `service_key` - The string needed by the service provider to provision the ExpressRoute circuit.
- `service_provider_properties` - A `service_provider_properties` block for the ExpressRoute circuit as documented below
- `sku` - A `sku` block for the ExpressRoute circuit as documented below.

`service_provider_properties` supports the following:

- `service_provider_name` - The name of the ExpressRoute Service Provider.
- `peering_location` - The name of the peering location and **not** the Azure resource location.

- `bandwidth_in_mbps` - The bandwidth in Mbps of the ExpressRoute circuit.

`peerings` supports the following:

- `peering_type` - The type of the ExpressRoute Circuit Peering. Acceptable values include `AzurePrivatePeering`, `AzurePublicPeering` and `MicrosoftPeering`. Changing this forces a new resource to be created. ~> **NOTE:** only one Peering of each Type can be created per ExpressRoute circuit.
- `primary_peer_address_prefix` - A /30 subnet for the primary link.
- `secondary_peer_address_prefix` - A /30 subnet for the secondary link.
- `vlan_id` - A valid VLAN ID to establish this peering on.
- `shared_key` - The shared key. Can be a maximum of 25 characters.
- `azure_asn` - The Either a 16-bit or a 32-bit ASN for Azure.
- `peer_asn` - The Either a 16-bit or a 32-bit ASN. Can either be public or private.

`sku` supports the following:

- `tier` - The service tier. Possible values are `Standard` or `Premium`.
- `family` - The billing mode for bandwidth. Possible values are `MeteredData` or `UnlimitedData`.

Data Source: azurerm_firewall

Use this data source to access information about an existing Azure Firewall.

Example Usage

```
data "azurerm_firewall" "test" {
  name                = "firewall1"
  resource_group_name = "firewall-RG"
}

output "firewall_private_ip" {
  value = "${data.azurerm_firewall.test.ip_configuration.0.private_ip_address}"
}
```

Argument Reference

- `name` - (Required) The name of the Azure Firewall.
- `resource_group_name` - (Required) The name of the Resource Group in which the Azure Firewall exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Azure Firewall.
- `ip_configuration` - A `ip_configuration` block as defined below.

A `ip_configuration` block exports the following:

- `subnet_id` - The Resource ID of the subnet where the Azure Firewall is deployed.
- `private_ip_address` - The private IP address of the Azure Firewall.
- `public_ip_address_id` - The Resource ID of the public IP address of the Azure Firewall.

Data Source: azurerm_hdinsight_cluster

Use this data source to access information about an existing HDInsight Cluster.

Example Usage

```
data "azurerm_hdinsight_cluster" "example" {
  name                = "example"
  resource_group_name = "example-resources"
}

output "https_endpoint" {
  value = "${data.azurerm_hdinsight_cluster.example.https_endpoint}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of this HDInsight Cluster.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight Cluster exists.

Attributes Reference

- `location` - The Azure Region in which this HDInsight Cluster exists.
- `cluster_version` - The version of HDInsights which is used on this HDInsight Cluster.
- `component_versions` - A map of versions of software used on this HDInsights Cluster.
- `gateway` - A `gateway` block as defined below.
- `edge_ssh_endpoint` - The SSH Endpoint of the Edge Node for this HDInsight Cluster, if an Edge Node exists.
- `https_endpoint` - The HTTPS Endpoint for this HDInsight Cluster.
- `kind` - The kind of HDInsight Cluster this is, such as a Spark or Storm cluster.
- `tier` - The SKU / Tier of this HDInsight Cluster.
- `ssh_endpoint` - The SSH Endpoint for this HDInsight Cluster.
- `tags` - A map of tags assigned to the HDInsight Cluster.

A `gateway` block exports the following:

- `enabled` - Is the Ambari Portal enabled?
- `username` - The username used for the Ambari Portal.

- password - The password used for the Ambari Portal.

Data Source: azurerm_healthcare_service

Use this data source to access information about an existing Healthcare Service

Example Usage

```
data "azurerm_healthcare_service" "example" {
  name                = "example-healthcare_service"
  resource_group_name = "example-resources"
  location            = "westus2"
}

output "healthcare_service_id" {
  name                = "uniquefhirname"
  resource_group_name = "sample-resource-group"
  kind                = "fhir-R4"
  cosmosdb_throughput = "2000"

  access_policy_object_ids = ["xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxx"]

  tags = {
    "environment" = "testenv"
    "purpose"     = "AcceptanceTests"
  }
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Healthcare Service.
- `resource_group_name` - (Required) The name of the Resource Group in which the Healthcare Service exists.

Attributes Reference

The following attributes are exported:

- `location` - The Azure Region where the Service is located.

Please Note: Not all locations support this resource. Some are West US 2, North Central US, and UK West.

- `kind` - The type of the service.
- `authentication_configuration` - An `authentication_configuration` block as defined below.
- `cosmosdb_offer_throughput` - The provisioned throughput for the backing database.
- `cors_configuration` - A `cors_configuration` block as defined below.

- `tags` - A mapping of tags to assign to the resource.
-

An `authentication_configuration` exports the following:

- `authority` - The Azure Active Directory (tenant) that serves as the authentication authority to access the service.
 - `audience` - The intended audience to receive authentication tokens for the service.
 - `smart_proxy_enabled` - Is the 'SMART on FHIR' option for mobile and web implementations enabled?
-

A `cors_configuration` block exports the following:

- `allowed_origins` - The set of origins to be allowed via CORS.
- `allowed_headers` - The set of headers to be allowed via CORS.
- `allowed_methods` - The methods to be allowed via CORS.
- `max_age_in_seconds` - The max age to be allowed via CORS.
- `allow_credentials` - Are credentials are allowed via CORS?

Data Source: azurerm_image

Use this data source to access information about an existing Image.

Example Usage

```
data "azurerm_image" "search" {
  name                = "search-api"
  resource_group_name = "packerimages"
}

output "image_id" {
  value = "${data.azurerm_image.search.id}"
}
```

Argument Reference

- `name` - (Optional) The name of the Image.
- `name_regex` - (Optional) Regex pattern of the image to match.
- `sort_descending` - (Optional) By default when matching by regex, images are sorted by name in ascending order and the first match is chosen, to sort descending, set this flag.
- `resource_group_name` - (Required) The Name of the Resource Group where this Image exists.

Attributes Reference

- `data_disk` - a collection of `data_disk` blocks as defined below.
- `name` - the name of the Image.
- `location` - the Azure Location where this Image exists.
- `os_disk` - a `os_disk` block as defined below.
- `tags` - a mapping of tags to assigned to the resource.
- `zone_resilient` - is zone resiliency enabled?

`os_disk` supports the following:

- `blob_uri` - the URI in Azure storage of the blob used to create the image.
- `caching` - the caching mode for the OS Disk, such as `ReadWrite`, `ReadOnly`, or `None`.
- `managed_disk_id` - the ID of the Managed Disk used as the OS Disk Image.
- `os_state` - the State of the OS used in the Image, such as `Generalized`.

- `os_type` - the type of Operating System used on the OS Disk, such as `Linux` or `Windows` .
- `size_gb` - the size of the OS Disk in GB.

`data_disk` supports the following:

- `blob_uri` - the URI in Azure storage of the blob used to create the image.
- `caching` - the caching mode for the Data Disk, such as `ReadWrite` , `ReadOnly` , or `None` .
- `lun` - the logical unit number of the data disk.
- `managed_disk_id` - the ID of the Managed Disk used as the Data Disk Image.
- `size_gb` - the size of this Data Disk in GB.

Data Source: azurerm_key_vault_access_policy

Use this data source to access information about the permissions from the Management Key Vault Templates.

Example Usage

```
data "azurerm_key_vault_access_policy" "contributor" {
  name = "Key Management"
}

output "access_policy_key_permissions" {
  value = "${data.azurerm_key_vault_access_policy.key_permissions}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Management Template. Possible values are: `Key Management`, `Secret Management`, `Certificate Management`, `Key & Secret Management`, `Key & Certificate Management`, `Secret & Certificate Management`, `Key`, `Secret`, & `Certificate Management`

Attributes Reference

- `id` - the ID of the Key Vault Access Policy
- `key_permissions` - the key permissions for the access policy
- `secret_permissions` - the secret permissions for the access policy
- `certificate_permissions` - the certificate permissions for the access policy

Data Source: azurerm_key_vault

Use this data source to access information about an existing Key Vault.

Example Usage

```
data "azurerm_key_vault" "test" {
  name                = "mykeyvault"
  resource_group_name = "some-resource-group"
}

output "vault_uri" {
  value = "${data.azurerm_key_vault.test.vault_uri}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault.
- `resource_group_name` - (Required) The name of the Resource Group in which the Key Vault exists.

Attributes Reference

The following attributes are exported:

- `id` - The Vault ID.
- `vault_uri` - The URI of the vault for performing operations on keys and secrets.
- `location` - The Azure Region in which the Key Vault exists.
- `sku` - A `sku` block as described below.
- `tenant_id` - The Azure Active Directory Tenant ID used for authenticating requests to the Key Vault.
- `access_policy` - One or more `access_policy` blocks as defined below.
- `enabled_for_deployment` - Can Azure Virtual Machines retrieve certificates stored as secrets from the Key Vault?
- `enabled_for_disk_encryption` - Can Azure Disk Encryption retrieve secrets from the Key Vault?
- `enabled_for_template_deployment` - Can Azure Resource Manager retrieve secrets from the Key Vault?
- `tags` - A mapping of tags assigned to the Key Vault.

A `sku` block exports the following:

- `name` - The name of the SKU used for this Key Vault.

`access_policy` supports the following:

- `tenant_id` - The Azure Active Directory Tenant ID used to authenticate requests for this Key Vault.
- `object_id` - An Object ID of a User, Service Principal or Security Group.
- `application_id` - The Object ID of a Azure Active Directory Application.
- `certificate_permissions` - A list of certificate permissions applicable to this Access Policy.
- `key_permissions` - A list of key permissions applicable to this Access Policy.
- `secret_permissions` - A list of secret permissions applicable to this Access Policy.
- `storage_permissions` - A list of storage permissions applicable to this Access Policy.

Data Source: azurerm_key_vault_key

Use this data source to access information about an existing Key Vault Key.

Note: All arguments including the secret value will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
data "azurerm_key_vault_key" "test" {
  name          = "secret-sauce"
  key_vault_id = "${data.azurerm_key_vault.existing.id}"
}

output "key_type" {
  value = "${data.azurerm_key_vault_secret.test.key_type}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault Key.
- `key_vault_id` - (Required) Specifies the ID of the Key Vault instance where the Secret resides, available on the `azurerm_key_vault` Data Source / Resource.

NOTE: The vault must be in the same subscription as the provider. If the vault is in another subscription, you must create an aliased provider for that subscription.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Key Vault Key.
- `e` - The RSA public exponent of this Key Vault Key.
- `key_type` - Specifies the Key Type of this Key Vault Key
- `key_size` - Specifies the Size of this Key Vault Key.
- `key_opts` - A list of JSON web key operations assigned to this Key Vault Key
- `n` - The RSA modulus of this Key Vault Key.
- `tags` - A mapping of tags assigned to this Key Vault Key.

- `version` - The current version of the Key Vault Key.

Data Source: azurerm_key_vault_secret

Use this data source to access information about an existing Key Vault Secret.

Note: All arguments including the secret value will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
data "azurerm_key_vault_secret" "test" {
  name          = "secret-sauce"
  key_vault_id = "${data.azurerm_key_vault.existing.id}"
}

output "secret_value" {
  value = "${data.azurerm_key_vault_secret.test.value}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault Secret.
- `key_vault_id` - (Required) Specifies the ID of the Key Vault instance where the Secret resides, available on the `azurerm_key_vault` Data Source / Resource.

NOTE: The vault must be in the same subscription as the provider. If the vault is in another subscription, you must create an aliased provider for that subscription.

Attributes Reference

The following attributes are exported:

- `id` - The Key Vault Secret ID.
- `value` - The value of the Key Vault Secret.
- `version` - The current version of the Key Vault Secret.
- `content_type` - The content type for the Key Vault Secret.
- `tags` - Any tags assigned to this resource.

Data Source: azurerm_kubernetes_cluster

Use this data source to access information about an existing Managed Kubernetes Cluster (AKS).

Note: All arguments including the client secret will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
data "azurerm_kubernetes_cluster" "test" {
  name                = "myakscluster"
  resource_group_name = "my-example-resource-group"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the managed Kubernetes Cluster.
- `resource_group_name` - (Required) The name of the Resource Group in which the managed Kubernetes Cluster exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Kubernetes Managed Cluster.
- `addon_profile` - A `addon_profile` block as documented below.
- `agent_pool_profile` - An `agent_pool_profile` block as documented below.
- `dns_prefix` - The DNS Prefix of the managed Kubernetes cluster.
- `fqdn` - The FQDN of the Azure Kubernetes Managed Cluster.
- `kube_admin_config` - A `kube_admin_config` block as defined below. This is only available when Role Based Access Control with Azure Active Directory is enabled.
- `kube_admin_config_raw` - Raw Kubernetes config for the admin account to be used by kubectl (<https://kubernetes.io/docs/reference/kubectl/overview/>) and other compatible tools. This is only available when Role Based Access Control with Azure Active Directory is enabled.
- `kube_config` - A `kube_config` block as defined below.
- `kube_config_raw` - Base64 encoded Kubernetes configuration.

- `kubernetes_version` - The version of Kubernetes used on the managed Kubernetes Cluster.
 - `location` - The Azure Region in which the managed Kubernetes Cluster exists.
 - `linux_profile` - A `linux_profile` block as documented below.
 - `windows_profile` - A `windows_profile` block as documented below.
 - `network_profile` - A `network_profile` block as documented below.
 - `node_resource_group` - Auto-generated Resource Group containing AKS Cluster resources.
 - `role_based_access_control` - A `role_based_access_control` block as documented below.
 - `service_principal` - A `service_principal` block as documented below.
 - `tags` - A mapping of tags assigned to this resource.
-

A `addon_profile` block exports the following:

- `http_application_routing` - A `http_application_routing` block.
 - `oms_agent` - A `oms_agent` block.
 - `kube_dashboard` - A `kube_dashboard` block.
 - `azure_policy` - A `azure_policy` block.
-

A `agent_pool_profile` block exports the following:

- `type` - The type of the Agent Pool.
 - `count` - The number of Agents (VM's) in the Pool.
 - `max_pods` - The maximum number of pods that can run on each agent.
 - `availability_zones` - The availability zones used for the nodes.
 - `enable_auto_scaling` - If the auto-scaler is enabled.
 - `min_count` - Minimum number of nodes for auto-scaling
 - `max_count` - Maximum number of nodes for auto-scaling
 - `name` - The name assigned to this pool of agents.
 - `os_disk_size_gb` - The size of the Agent VM's Operating System Disk in GB.
 - `os_type` - The Operating System used for the Agents.
 - `vm_size` - The size of each VM in the Agent Pool (e.g. `Standard_F1`).
 - `vnet_subnet_id` - The ID of the Subnet where the Agents in the Pool are provisioned.
 - `node_taints` - The list of Kubernetes taints which are applied to nodes in the agent pool
-

A `azure_active_directory` block exports the following:

- `client_app_id` - The Client ID of an Azure Active Directory Application.
 - `server_app_id` - The Server ID of an Azure Active Directory Application.
 - `tenant_id` - The Tenant ID used for Azure Active Directory Application.
-

A `http_application_routing` block exports the following:

- `enabled` - Is HTTP Application Routing Enabled?
 - `http_application_routing_zone_name` - The Zone Name of the HTTP Application Routing.
-

The `kube_admin_config` and `kube_config` blocks exports the following:

- `client_key` - Base64 encoded private key used by clients to authenticate to the Kubernetes cluster.
- `client_certificate` - Base64 encoded public certificate used by clients to authenticate to the Kubernetes cluster.
- `cluster_ca_certificate` - Base64 encoded public CA certificate used as the root of trust for the Kubernetes cluster.
- `host` - The Kubernetes cluster server host.
- `username` - A username used to authenticate to the Kubernetes cluster.
- `password` - A password or token used to authenticate to the Kubernetes cluster.

NOTE: It's possible to use these credentials with the Kubernetes Provider (</docs/providers/kubernetes/index.html>) like so:

```
provider "kubernetes" {  
  host          = "${data.azurerm_kubernetes_cluster.main.kube_config.0.host}"  
  username      = "${data.azurerm_kubernetes_cluster.main.kube_config.0.username}"  
  password      = "${data.azurerm_kubernetes_cluster.main.kube_config.0.password}"  
  client_certificate = "${base64decode(data.azurerm_kubernetes_cluster.main.kube_config.0.client_certificate)}"  
  client_key     = "${base64decode(data.azurerm_kubernetes_cluster.main.kube_config.0.client_key)}"  
  cluster_ca_certificate = "${base64decode(data.azurerm_kubernetes_cluster.main.kube_config.0.cluster_ca_certificate)}"  
}
```

A `linux_profile` block exports the following:

- `admin_username` - The username associated with the administrator account of the managed Kubernetes Cluster.
 - `ssh_key` - An `ssh_key` block as defined below.
-

A `windows_profile` block exports the following:

- `admin_username` - The username associated with the administrator account of the Windows VMs.
-

A `network_profile` block exports the following:

- `docker_bridge_cidr` - IP address (in CIDR notation) used as the Docker bridge IP address on nodes.
 - `dns_service_ip` - IP address within the Kubernetes service address range used by cluster service discovery (kube-dns).
 - `network_plugin` - Network plugin used such as `azure` or `kubenet`.
 - `network_policy` - Network policy to be used with Azure CNI. Eg: `calico` or `azure`
 - `pod_cidr` - The CIDR used for pod IP addresses.
 - `service_cidr` - Network range used by the Kubernetes service.
-

A `oms_agent` block exports the following:

- `enabled` - Is the OMS Agent Enabled?
 - `log_analytics_workspace_id` - The ID of the Log Analytics Workspace which the OMS Agent should send data to.
-

A `kube_dashboard` block supports the following:

- `enabled` - (Required) Is the Kubernetes Dashboard enabled?
-

A `azure_policy` block supports the following:

- `enabled` - (Required) Is Azure Policy for Kubernetes enabled?
-

A `role_based_access_control` block exports the following:

- `azure_active_directory` - A `azure_active_directory` block as documented above.
 - `enabled` - Is Role Based Access Control enabled?
-

A `service_principal` block supports the following:

- `client_id` - The Client ID of the Service Principal used by this Managed Kubernetes Cluster.
-

A `ssh_key` block exports the following:

- `key_data` - The Public SSH Key used to access the cluster.

Data Source: azurerm_kubernetes_service_versions

Use this data source to retrieve the version of Kubernetes supported by Azure Kubernetes Service.

Example Usage

```
data "azurerm_kubernetes_service_versions" "current" {
  location = "West Europe"
}

output "versions" {
  value = "${data.azurerm_kubernetes_service_versions.current.versions}"
}

output "latest_version" {
  value = "${data.azurerm_kubernetes_service_versions.current.latest_version}"
}
```

Argument Reference

- `location` - (Required) Specifies the location in which to query for versions.
- `version_prefix` - (Optional) A prefix filter for the versions of Kubernetes which should be returned; for example `1.` will return `1.9` to `1.14`, whereas `1.12` will return `1.12.2`.

Attributes Reference

- `versions` - The list of all supported versions.
- `latest_version` - The most recent version available.

Data Source: azurerm_lb_backend_address_pool

Use this data source to access information about an existing Load Balancer's Backend Address Pool.

Example Usage

```
data "azurerm_lb" "test" {
  name                = "example-lb"
  resource_group_name = "example-resources"
}

data "azurerm_lb_backend_address_pool" "test" {
  name                = "first"
  loadbalancer_id    = data.azurerm_lb.test.id
}

output "backend_address_pool_id" {
  value = data.azurerm_lb_backend_address_pool.test.id
}

output "backend_ip_configuration_ids" {
  value = data.azurerm_lb_backend_address_pool.beap.backend_ip_configurations.*.id
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Backend Address Pool.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which the Backend Address Pool exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Backend Address Pool.
- `name` - The name of the Backend Address Pool.
- `backend_ip_configurations` - An array of references to IP addresses defined in network interfaces.

Data Source: azurerm_lb

Use this data source to access information about an existing Load Balancer

Example Usage

```
data "azurerm_lb" "test" {
  name                = "example-lb"
  resource_group_name = "example-resources"
}

output "loadbalancer_id" {
  value = "${data.azurerm_lb.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Load Balancer.
- `resource_group_name` - (Required) The name of the Resource Group in which the Load Balancer exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer.
- `frontend_ip_configuration` - (Optional) A `frontend_ip_configuration` block as documented below.
- `location` - The Azure location where the Load Balancer exists.
- `private_ip_address` - The first private IP address assigned to the load balancer in `frontend_ip_configuration` blocks, if any.
- `private_ip_addresses` - The list of private IP address assigned to the load balancer in `frontend_ip_configuration` blocks, if any.
- `sku` - The SKU of the Load Balancer.
- `tags` - A mapping of tags assigned to the resource.

A `frontend_ip_configuration` block exports the following:

- `name` - The name of the Frontend IP Configuration.
- `subnet_id` - The ID of the Subnet which is associated with the IP Configuration.
- `private_ip_address` - Private IP Address to assign to the Load Balancer.

- `private_ip_address_allocation` - The allocation method for the Private IP Address used by this Load Balancer.
- `public_ip_address_id` - The ID of a Public IP Address which is associated with this Load Balancer.
- `zones` - A list of Availability Zones which the Load Balancer's IP Addresses should be created in.

Data Source: azure_rm_log_analytics_workspace

Use this data source to access information about an existing Log Analytics (formally Operational Insights) Workspace.

Example Usage

```
data "azure_rm_log_analytics_workspace" "test" {
  name                = "acctest-01"
  resource_group_name = "acctest"
}

output "log_analytics_workspace_id" {
  value = "${data.azure_rm_log_analytics_workspace.test.workspace_id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Log Analytics Workspace.
- `resource_group_name` - (Required) The name of the resource group in which the Log Analytics workspace is located in.

Attributes Reference

The following attributes are exported:

- `id` - The Azure Resource ID of the Log Analytics Workspace.
- `primary_shared_key` - The Primary shared key for the Log Analytics Workspace.
- `secondary_shared_key` - The Secondary shared key for the Log Analytics Workspace.
- `workspace_id` - The Workspace (or Customer) ID for the Log Analytics Workspace.
- `portal_url` - The Portal URL for the Log Analytics Workspace.
- `sku` - The Sku of the Log Analytics Workspace.
- `retention_in_days` - The workspace data retention in days.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_logic_app_workflow

Use this data source to access information about an existing Logic App Workflow.

Example Usage

```
data "azurerm_logic_app_workflow" "test" {
  name                = "workflow1"
  resource_group_name = "my-resource-group"
}

output "access_endpoint" {
  value = "${data.azurerm_logic_app_workflow.test.access_endpoint}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Logic App Workflow.
- `resource_group_name` - (Required) The name of the Resource Group in which the Logic App Workflow exists.

Attributes Reference

The following attributes are exported:

- `id` - The Logic App Workflow ID.
- `location` - The Azure location where the Logic App Workflow exists.
- `workflow_schema` - The Schema used for this Logic App Workflow.
- `workflow_version` - The version of the Schema used for this Logic App Workflow. Defaults to `1.0.0.0`.
- `parameters` - A map of Key-Value pairs.
- `tags` - A mapping of tags assigned to the resource.
- `access_endpoint` - The Access Endpoint for the Logic App Workflow

Data Source: azurerm_managed_disk

Use this data source to access information about an existing Managed Disk.

Example Usage

```
data "azurerm_managed_disk" "datasourcemd" {
  name                = "testManagedDisk"
  resource_group_name = "acctestRG"
}

resource "azurerm_virtual_network" "test" {
  name                = "acctvn"
  address_space       = ["10.0.0.0/16"]
  location             = "West US 2"
  resource_group_name = "acctestRG"
}

resource "azurerm_subnet" "test" {
  name                = "acctsub"
  resource_group_name = "acctestRG"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "acctni"
  location             = "West US 2"
  resource_group_name = "acctestRG"

  ip_configuration {
    name                        = "testconfiguration1"
    subnet_id                  = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "acctvm"
  location             = "West US 2"
  resource_group_name = "acctestRG"
  network_interface_ids = ["${azurerm_network_interface.test.id}"]
  vm_size              = "Standard_DS1_v2"

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }

  storage_os_disk {
    name          = "myosdisk1"
    caching       = "ReadWrite"
    create_option = "FromImage"
  }
}
```

```

    managed_disk_type = "Standard_LRS"
  }

  storage_data_disk {
    name           = "datadisk_new"
    managed_disk_type = "Standard_LRS"
    create_option   = "Empty"
    lun             = 0
    disk_size_gb     = "1023"
  }

  storage_data_disk {
    name           = "${data.azurerm_managed_disk.datasourcemd.name}"
    managed_disk_id = "${data.azurerm_managed_disk.datasourcemd.id}"
    create_option   = "Attach"
    lun             = 1
    disk_size_gb     = "${data.azurerm_managed_disk.datasourcemd.disk_size_gb}"
  }

  os_profile {
    computer_name = "hostname"
    admin_username = "testadmin"
    admin_password = "Password1234!"
  }

  os_profile_linux_config {
    disable_password_authentication = false
  }

  tags = {
    environment = "staging"
  }
}

```

Argument Reference

- `name` - (Required) Specifies the name of the Managed Disk.
- `resource_group_name` - (Required) Specifies the name of the resource group.

Attributes Reference

- `storage_account_type` - The storage account type for the managed disk.
- `source_uri` - The source URI for the managed disk
- `source_resource_id` - ID of an existing managed disk that the current resource was created from.
- `os_type` - The operating system for managed disk. Valid values are `Linux` or `Windows`
- `disk_size_gb` - The size of the managed disk in gigabytes.
- `disk_iops_read_write` - The number of IOPS allowed for this disk. One operation can transfer between 4k and 256k bytes.

- `disk_mbps_read_write` - The bandwidth allowed for this disk.
- `tags` - A mapping of tags assigned to the resource.
- `zones` - A collection containing the availability zone the managed disk is allocated in.

Data Source: azurerm_management_group

Use this data source to access information about an existing Management Group.

Example Usage

```
data "azurerm_management_group" "test" {
  group_id = "00000000-0000-0000-0000-000000000000"
}

output "display_name" {
  value = "${data.azurerm_management_group.test.display_name}"
}
```

Argument Reference

The following arguments are supported:

- `group_id` - (Required) Specifies the UUID of this Management Group.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Management Group.
- `display_name` - A friendly name for the Management Group.
- `parent_management_group_id` - The ID of any Parent Management Group.
- `subscription_ids` - A list of Subscription ID's which are assigned to the Management Group.

Data Source: azurerm_maps_account

Use this data source to access information about an existing Azure Maps Account.

Example Usage

```
data "azurerm_maps_account" "test" {
  name                = "production"
  resource_group_name = "maps"
}

output "maps_account_id" {
  value = "${data.azurerm_maps_account.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Maps Account.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which the Maps Account is located.

Attributes Reference

- `id` - The ID of the Maps Account.
- `sku_name` - The sku of the Azure Maps Account.
- `primary_access_key` - The primary key used to authenticate and authorize access to the Maps REST APIs.
- `secondary_access_key` - The primary key used to authenticate and authorize access to the Maps REST APIs. The second key is given to provide seamless key regeneration.
- `x_ms_client_id` - A unique identifier for the Maps Account.

Data Source: azurerm_monitor_action_group

Use this data source to access the properties of an Action Group.

Example Usage

```
data "azurerm_monitor_action_group" "example" {
  resource_group_name = "terraform-example-rg"
  name                = "tfex-actiongroup"
}

output "action_group_id" {
  value = "${data.azurerm_monitor_action_group.example.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Action Group.
- `resource_group_name` - (Required) Specifies the name of the resource group the Action Group is located in.

Attributes Reference

- `id` - The ID of the Action Group.
- `short_name` - The short name of the action group.
- `enabled` - Whether this action group is enabled.
- `email_receiver` - One or more `email_receiver` blocks as defined below.
- `sms_receiver` - One or more `sms_receiver` blocks as defined below.
- `webhook_receiver` - One or more `webhook_receiver` blocks as defined below.

`email_receiver` supports the following:

- `name` - The name of the email receiver.
- `email_address` - The email address of this receiver.

`sms_receiver` supports the following:

- `name` - The name of the SMS receiver.
- `country_code` - The country code of the SMS receiver.
- `phone_number` - The phone number of the SMS receiver.

`webhook_receiver` supports the following:

- `name` - The name of the webhook receiver.
- `service_uri` - The URI where webhooks should be sent.

Data Source: azurerm_monitor_diagnostic_categories

Use this data source to access information about the Monitor Diagnostics Categories supported by an existing Resource.

Example Usage

```
data "azurerm_key_vault" "test" {
  name                       = "${azurerm_key_vault.test.name}"
  resource_group_name       = "${azurerm_key_vault.test.resource_group_name}"
}

data "azurerm_monitor_diagnostic_categories" "test" {
  resource_id = "${data.azurerm_key_vault.test.id}"
}
```

Argument Reference

- `resource_id` - (Required) The ID of an existing Resource which Monitor Diagnostics Categories should be retrieved for.

Attributes Reference

- `id` - The ID of the Resource.
- `logs` - A list of the Log Categories supported for this Resource.
- `metrics` - A list of the Metric Categories supported for this Resource.

Data Source: azurerm_monitor_log_profile

Use this data source to access the properties of a Log Profile.

Example Usage

```
data "azurerm_monitor_log_profile" "test" {
  name = "test-logprofile"
}

output "log_profile_storage_account_id" {
  value = "${data.azurerm_monitor_log_profile.test.storage_account_id}"
}
```

Argument Reference

- `name` - (Required) Specifies the Name of the Log Profile.

Attributes Reference

- `id` - The ID of the Log Profile.
- `storage_account_id` - The resource id of the storage account in which the Activity Log is stored.
- `servicebus_rule_id` - The service bus (or event hub) rule ID of the service bus (or event hub) namespace in which the Activity Log is streamed to.
- `locations` - List of regions for which Activity Log events are stored or streamed.
- `categories` - List of categories of the logs.
- `retention_policy` - a `retention_policy` block as documented below.

The `retention_policy` block supports:

- `enabled` - A boolean value indicating whether the retention policy is enabled.
- `days` - The number of days for the retention policy.

Data Source: azurerm_mssql_elasticpool

Use this data source to access information about an existing SQL elastic pool.

Example Usage

```
data "azurerm_mssql_elasticpool" "example" {
  name                = "mssqlelasticpoolname"
  resource_group_name = "example-resources"
  server_name         = "example-sql-server"
}

output "elasticpool_id" {
  value = "${data.azurerm_mssql_elasticpool.example.id}"
}
```

Argument Reference

- `name` - (Required) The name of the elastic pool.
- `resource_group_name` - (Required) The name of the resource group which contains the elastic pool.
- `server_name` - (Required) The name of the SQL Server which contains the elastic pool.

Attributes Reference

- `location` - Specifies the supported Azure location where the resource exists.
- `max_size_gb` - The max data size of the elastic pool in gigabytes.
- `max_size_bytes` - The max data size of the elastic pool in bytes.
- `per_db_min_capacity` - The minimum capacity all databases are guaranteed.
- `per_db_max_capacity` - The maximum capacity any one database can consume.
- `tags` - A mapping of tags to assign to the resource.
- `zone_redundant` - Whether or not this elastic pool is zone redundant.

Data Source: azurerm_network_ddos_protection_plan

Use this data source to access information about an existing Azure Network DDoS Protection Plan.

Example Usage

```
data "azurerm_network_ddos_protection_plan" "example" {
  name = "${azurerm_network_ddos_protection_plan.example.name}"
  resource_group_name = "${azurerm_network_ddos_protection_plan.example.resource_group_name}"
}

output "ddos_protection_plan_id" {
  value = "${data.azurerm_network_ddos_protection_plan.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Network DDoS Protection Plan.
- `resource_group_name` - (Required) The name of the resource group where the Network DDoS Protection Plan exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the DDoS Protection Plan
- `location` - Specifies the supported Azure location where the resource exists.
- `tags` - A mapping of tags assigned to the resource.
- `virtual_network_ids` - The Resource ID list of the Virtual Networks associated with DDoS Protection Plan.

Data Source: azurerm_network_interface

Use this data source to access information about an existing Network Interface.

Example Usage

```
data "azurerm_network_interface" "test" {
  name                = "acctest-nic"
  resource_group_name = "networking"
}

output "network_interface_id" {
  value = "${data.azurerm_network_interface.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Network Interface.
- `resource_group_name` - (Required) Specifies the name of the resource group the Network Interface is located in.

Attributes Reference

- `id` - The ID of the Network Interface.
- `applied_dns_servers` - List of DNS servers applied to the specified Network Interface.
- `enable_accelerated_networking` - Indicates if accelerated networking is set on the specified Network Interface.
- `enable_ip_forwarding` - Indicate if IP forwarding is set on the specified Network Interface.
- `dns_servers` - The list of DNS servers used by the specified Network Interface.
- `internal_dns_name_label` - The internal dns name label of the specified Network Interface.
- `ip_configuration` - One or more `ip_configuration` blocks as defined below.
- `location` - The location of the specified Network Interface.
- `mac_address` - The MAC address used by the specified Network Interface.
- `network_security_group_id` - The ID of the network security group associated to the specified Network Interface.
- `private_ip_address` - The primary private ip address associated to the specified Network Interface.
- `private_ip_addresses` - The list of private ip addresses associates to the specified Network Interface.
- `tags` - List the tags associated to the specified Network Interface.
- `virtual_machine_id` - The ID of the virtual machine that the specified Network Interface is attached to.

A `ip_configuration` block contains:

- `name` - The name of the IP Configuration.
- `subnet_id` - The ID of the Subnet which the Network Interface is connected to.
- `private_ip_address` - The Private IP Address assigned to this Network Interface.
- `private_ip_address_allocation` - The IP Address allocation type for the Private address, such as `Dynamic` or `Static`.
- `public_ip_address_id` - The ID of the Public IP Address which is connected to this Network Interface.
- `application_gateway_backend_address_pools_ids` - A list of Backend Address Pool ID's within a Application Gateway that this Network Interface is connected to.
- `load_balancer_backend_address_pools_ids` - A list of Backend Address Pool ID's within a Load Balancer that this Network Interface is connected to.
- `load_balancer_inbound_nat_rules_ids` - A list of Inbound NAT Rule ID's within a Load Balancer that this Network Interface is connected to.
- `primary` - is this the Primary IP Configuration for this Network Interface?

Data Source: azurerm_network_security_group

Use this data source to access information about an existing Network Security Group.

Example Usage

```
data "azurerm_network_security_group" "test" {
  name                = "${azurerm_network_security_group.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

output "location" {
  value = "${data.azurerm_network_security_group.test.location}"
}
```

Argument Reference

- `name` - (Required) Specifies the Name of the Network Security Group.
- `resource_group_name` - (Required) Specifies the Name of the Resource Group within which the Network Security Group exists

Attributes Reference

- `id` - The ID of the Network Security Group.
- `location` - The supported Azure location where the resource exists.
- `security_rule` - One or more `security_rule` blocks as defined below.
- `tags` - A mapping of tags assigned to the resource.

The `security_rule` block supports:

- `name` - The name of the security rule.
- `description` - The description for this rule.
- `protocol` - The network protocol this rule applies to.
- `source_port_range` - The Source Port or Range.
- `destination_port_range` - The Destination Port or Range.
- `source_address_prefix` - CIDR or source IP range or * to match any IP.
- `source_address_prefixes` - A list of CIDRs or source IP ranges.
- `destination_address_prefix` - CIDR or destination IP range or * to match any IP.

- `destination_address_prefixes` - A list of CIDRs or destination IP ranges.
- `source_application_security_group_ids` - A List of source Application Security Group ID's
- `destination_application_security_group_ids` - A List of destination Application Security Group ID's
- `access` - Is network traffic is allowed or denied?
- `priority` - The priority of the rule
- `direction` - The direction specifies if rule will be evaluated on incoming or outgoing traffic.

Data Source: azurerm_network_watcher

Use this data source to access information about an existing Network Watcher.

Example Usage

```
data "azurerm_network_watcher" "test" {
  name                       = "${azurerm_network_watcher.test.name}"
  resource_group_name       = "${azurerm_resource_group.test.name}"
}

output "network_watcher_id" {
  value = "${data.azurerm_network_watcher.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the Name of the Network Watcher.
- `resource_group_name` - (Required) Specifies the Name of the Resource Group within which the Network Watcher exists.

Attributes Reference

- `id` - The ID of the Network Watcher.
- `location` - The supported Azure location where the resource exists.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azure_rm_notification_hub

Use this data source to access information about an existing Notification Hub within a Notification Hub Namespace.

Example Usage

```
data "azure_rm_notification_hub" "test" {
  name            = "notification-hub"
  namespace_name  = "namespace-name"
  resource_group_name = "resource-group-name"
}

output "id" {
  value = "${data.azure_rm_notification_hub.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the Name of the Notification Hub.
- `namespace_name` - (Required) Specifies the Name of the Notification Hub Namespace which contains the Notification Hub.
- `resource_group_name` - (Required) Specifies the Name of the Resource Group within which the Notification Hub exists.

Attributes Reference

- `id` - The ID of the Notification Hub.
- `location` - The Azure Region in which this Notification Hub exists.
- `apns_credential` - A `apns_credential` block as defined below.
- `gcm_credential` - A `gcm_credential` block as defined below.

A `apns_credential` block exports:

- `application_mode` - The Application Mode which defines which server the APNS Messages should be sent to. Possible values are `Production` and `Sandbox`.
- `bundle_id` - The Bundle ID of the iOS/macOS application to send push notifications for, such as `com.hashicorp.example`.
- `key_id` - The Apple Push Notifications Service (APNS) Key.
- `team_id` - The ID of the team the Token.

- `token` - The Push Token associated with the Apple Developer Account.
-

A `gcm_credential` block exports:

- `api_key` - The API Key associated with the Google Cloud Messaging service.

Data Source: azurerm_notification_hub_namespace

Use this data source to access information about an existing Notification Hub Namespace.

Example Usage

```
data "azurerm_notification_hub_namespace" "test" {
  name                = "my-namespace"
  resource_group_name = "my-resource-group"
}

output "servicebus_endpoint" {
  value = "${data.azurerm_notification_hub_namespace.test.servicebus_endpoint}"
}
```

Argument Reference

- `name` - (Required) Specifies the Name of the Notification Hub Namespace.
- `resource_group_name` - (Required) Specifies the Name of the Resource Group within which the Notification Hub exists.

Attributes Reference

- `id` - The ID of the Notification Hub Namespace.
- `location` - The Azure Region in which this Notification Hub Namespace exists.
- `namespace_type` - The Type of Namespace, such as `Messaging` or `NotificationHub`.
- `sku` - A `sku` block as defined below.
- `enabled` - Is this Notification Hub Namespace enabled?

A `sku` block exports the following:

- `name` - (Required) The name of the SKU to use for this Notification Hub Namespace. Possible values are `Free`, `Basic` or `Standard`.

Data Source: azurerm_platform_image

Use this data source to access information about a Platform Image.

Example Usage

```
data "azurerm_platform_image" "test" {
  location = "West Europe"
  publisher = "Canonical"
  offer    = "UbuntuServer"
  sku      = "16.04-LTS"
}

output "version" {
  value = "${data.azurerm_platform_image.test.version}"
}
```

Argument Reference

- `location` - (Required) Specifies the Location to pull information about this Platform Image from.
- `publisher` - (Required) Specifies the Publisher associated with the Platform Image.
- `offer` - (Required) Specifies the Offer associated with the Platform Image.
- `sku` - (Required) Specifies the SKU of the Platform Image.

Attributes Reference

- `id` - The ID of the Platform Image.
- `version` - The latest version of the Platform Image.

Data Source: azurerm_policy_definition

Use this data source to access information about a Policy Definition, both custom and built in. Retrieves Policy Definitions from your current subscription by default.

Example Usage

```
data "azurerm_policy_definition" "test" {
  display_name = "Allowed resource types"
}

output "id" {
  value = "${data.azurerm_policy_definition.test.id}"
}
```

Argument Reference

- `display_name` - (Required) Specifies the name of the Policy Definition.
- `management_group_id` - (Optional) Only retrieve Policy Definitions from this Management Group.

Attributes Reference

- `id` - The ID of the Policy Definition.
- `name` - The Name of the Policy Definition.
- `type` - The Type of Policy.
- `description` - The Description of the Policy.
- `policy_type` - The Type of the Policy, such as `Microsoft.Authorization/policyDefinitions`.
- `policy_rule` - The Rule as defined (in JSON) in the Policy.
- `parameters` - Any Parameters defined in the Policy.
- `metadata` - Any Metadata defined in the Policy.

Data Source: azurerm_postgresql_server

Use this data source to access information about an existing PostgreSQL Azure Database Server.

Example Usage

```
data "azurerm_postgresql_server" "test" {
  name                = "postgresql-server-1"
  resource_group_name = "api-rg-pro"
}

output "postgresql_server_id" {
  value = "${data.azurerm_postgresql_server.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the PostgreSQL Server.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the PostgreSQL Server exists.

Attributes Reference

- `location` - The location of the Resource Group in which the PostgreSQL Server exists.
- `fqdn` - The fully qualified domain name of the PostgreSQL Server.
- `version` - The version of the PostgreSQL Server.
- `administrator_login` - The administrator username of the PostgreSQL Server.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_proximity_placement_group

Use this data source to access information about an existing Proximity Placement Group.

Example Usage

```
data "azurerm_proximity_placement_group" "example" {
  name                = "tf-appsecuritygroup"
  resource_group_name = "my-resource-group"
}

output "proximity_placement_group_id" {
  value = "${data.azurerm_proximity_placement_group.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - The name of the Proximity Placement Group.
- `resource_group_name` - The name of the resource group in which the Proximity Placement Group exists.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Proximity Placement Group.

Data Source: azurerm_public_ip

Use this data source to access information about an existing Public IP Address.

Example Usage (reference an existing)

```
data "azurerm_public_ip" "test" {
  name                = "name_of_public_ip"
  resource_group_name = "name_of_resource_group"
}

output "domain_name_label" {
  value = "${data.azurerm_public_ip.test.domain_name_label}"
}

output "public_ip_address" {
  value = "${data.azurerm_public_ip.test.ip_address}"
}
```

Example Usage (Retrieve the Dynamic Public IP of a new VM)

```
resource "azurerm_resource_group" "test" {
  name      = "test-resources"
  location  = "West US 2"
}

resource "azurerm_virtual_network" "test" {
  name            = "test-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "acctsub"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name                = "test-pip"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Dynamic"
  idle_timeout_in_minutes = 30

  tags = {
    environment = "test"
  }
}
```

```

resource "azurerm_network_interface" "test" {
  name                = "test-nic"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id          = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Static"
    private_ip_address   = "10.0.2.5"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "test-vm"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_interface_ids = ["${azurerm_network_interface.test.id}"]

  # ...
}

data "azurerm_public_ip" "test" {
  name                = "${azurerm_public_ip.test.name}"
  resource_group_name = "${azurerm_virtual_machine.test.resource_group_name}"
}

output "public_ip_address" {
  value = "${data.azurerm_public_ip.test.ip_address}"
}

```

Argument Reference

- `name` - (Required) Specifies the name of the public IP address.
- `resource_group_name` - (Required) Specifies the name of the resource group.

Attributes Reference

- `domain_name_label` - The label for the Domain Name.
- `idle_timeout_in_minutes` - Specifies the timeout for the TCP idle connection.
- `fqdn` - Fully qualified domain name of the A DNS record associated with the public IP. This is the concatenation of the `domainNameLabel` and the regionalized DNS zone.
- `ip_address` - The IP address value that was allocated.
- `ip_version` - The IP version being used, for example `IPv4` or `IPv6`.
- `tags` - A mapping of tags to assigned to the resource.

Data Source: azurerm_public_ip_prefix

Use this data source to access information about an existing Public IP Prefix.

Example Usage (reference an existing)

```
data "azurerm_public_ip_prefix" "test" {
  name                = "name_of_public_ip"
  resource_group_name = "name_of_resource_group"
}

output "public_ip_prefix" {
  value = "${data.azurerm_public_ip_prefix.test.ip_prefix}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the public IP prefix.
- `resource_group_name` - (Required) Specifies the name of the resource group.

Attributes Reference

- `name` - The name of the Public IP prefix resource.
- `resource_group_name` - The name of the resource group in which to create the public IP.
- `location` - The supported Azure location where the resource exists.
- `sku` - The SKU of the Public IP Prefix.
- `prefix_length` - The number of bits of the prefix.
- `tags` - A mapping of tags to assigned to the resource.

Data Source: azurerm_public_ips

Use this data source to access information about a set of existing Public IP Addresses.

Example Usage

```
data "azurerm_public_ips" "test" {
  resource_group_name = "pip-test"
  attached             = false
}
```

Argument Reference

- `resource_group_name` - (Required) Specifies the name of the resource group.
- `attached` - (Optional) Filter to include IP Addresses which are attached to a device, such as a VM/LB (`true`) or `unattached` (`false`).
- `name_prefix` - (Optional) A prefix match used for the IP Addresses `name` field, case sensitive.
- `allocation_type` - (Optional) The Allocation Type for the Public IP Address. Possible values include `Static` or `Dynamic` .

Attributes Reference

- `public_ips` - A List of `public_ips` blocks as defined below filtered by the criteria above.

A `public_ips` block contains:

- `id` - The ID of the Public IP Address
- `domain_name_label` - The Domain Name Label of the Public IP Address
- `fqdn` - The FQDN of the Public IP Address
- `name` - The Name of the Public IP Address

Data Source:

azurerm_recovery_services_protection_policy_vm

Use this data source to access information about an existing Recovery Services VM Protection Policy.

Example Usage

```
data "azurerm_recovery_services_protection_policy_vm" "policy" {
  name = "policy"
  recovery_vault_name = "recovery_vault"
  resource_group_name = "resource_group"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Recovery Services VM Protection Policy.
- `recovery_vault_name` - (Required) Specifies the name of the Recovery Services Vault.
- `resource_group_name` - (Required) The name of the resource group in which the Recovery Services VM Protection Policy resides.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Recovery Services VM Protection Policy.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_recovery_services_vault

Use this data source to access information about an existing Recovery Services Vault.

Example Usage

```
data "azurerm_recovery_services_vault" "vault" {
  name = "tfex-recovery_vault"
  resource_group_name = "tfex-resource_group"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Recovery Services Vault.
- `resource_group_name` - (Required) The name of the resource group in which the Recovery Services Vault resides.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Recovery Services Vault.
- `location` - The Azure location where the resource resides.
- `tags` - A mapping of tags assigned to the resource.
- `sku` - The vault's current SKU.

Data Source: azurerm_redis_cache

Use this data source to access information about an existing Redis Cache

Example Usage

```
data "azurerm_redis_cache" "example" {
  name                = "myrediscache"
  resource_group_name = "redis-cache"
}

output "primary_access_key" {
  value = "${data.azurerm_redis_cache.example.primary_access_key}"
}

output "hostname" {
  value = "${data.azurerm_redis_cache.example.hostname}"
}
```

Argument Reference

- `name` - The name of the Redis cache
- `resource_group_name` - The name of the resource group the Redis cache instance is located in.

Attribute Reference

- `id` - The Cache ID.
- `location` - The location of the Redis Cache.
- `capacity` - The size of the Redis Cache deployed.
- `family` - The SKU family/pricing group used. Possible values are `C` (for Basic/Standard SKU family) and `P` (for Premium)
- `sku_name` - The SKU of Redis used. Possible values are `Basic`, `Standard` and `Premium`.
- `enable_non_ssl_port` - Whether the SSL port is enabled.
- `minimum_tls_version` - The minimum TLS version.
- `patch_schedule` - A list of `patch_schedule` blocks as defined below - only available for Premium SKU's.
- `private_static_ip_address` - The Static IP Address assigned to the Redis Cache when hosted inside the Virtual Network.
- `hostname` - The Hostname of the Redis Instance
- `ssl_port` - The SSL Port of the Redis Instance

- `port` - The non-SSL Port of the Redis Instance
- `primary_access_key` - The Primary Access Key for the Redis Instance
- `secondary_access_key` - The Secondary Access Key for the Redis Instance
- `redis_configuration` - A `redis_configuration` block as defined below.

A `patch_schedule` block supports the following (Requires Premium SKU's, attempting to access this value on Basic or Standard SKU's will result in an error):

- `day_of_week` - the Weekday name for the patch item
- `start_hour_utc` - The Start Hour for maintenance in UTC

Note: The Patch Window lasts for 5 hours from the `start_hour_utc`.

A `redis_configuration` block exports the following:

- `enable_authentication` - Specifies if authentication is enabled
- `maxmemory_reserved` - The value in megabytes reserved for non-cache usage e.g. failover
- `maxmemory_delta` - The max-memory delta for this Redis instance.
- `maxmemory_policy` - How Redis will select what to remove when `maxmemory` is reached.
- `maxfragmentationmemory_reserved` - Value in megabytes reserved to accommodate for memory fragmentation.
- `rdb_backup_enabled` - Is Backup Enabled? Only supported on Premium SKU's.
- `rdb_backup_frequency` - The Backup Frequency in Minutes. Only supported on Premium SKU's.
- `rdb_backup_max_snapshot_count` - The maximum number of snapshots that can be created as a backup.
- `rdb_storage_connection_string` - The Connection String to the Storage Account. Only supported for Premium SKU's.

NOTE: There's a bug in the Redis API where the original storage connection string isn't being returned, which is being tracked in this issue (<https://github.com/Azure/azure-rest-api-specs/issues/3037>). In the interim you can use the `ignore_changes` attribute to ignore changes to this field (https://www.terraform.io/docs/configuration/resources.html#ignore_changes) e.g.:

Data Source: azurerm_resource_group

Use this data source to access information about an existing Resource Group.

Example Usage

```
data "azurerm_resource_group" "test" {
  name = "dsrg_test"
}

resource "azurerm_managed_disk" "test" {
  name                = "managed_disk_name"
  location            = "${data.azurerm_resource_group.test.location}"
  resource_group_name = "${data.azurerm_resource_group.test.name}"
  storage_account_type = "Standard_LRS"
  create_option       = "Empty"
  disk_size_gb        = "1"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the resource group.

NOTE: If the specified location doesn't match the actual resource group location, an error message with the actual location value will be shown.

Attributes Reference

- `location` - The location of the resource group.
- `tags` - A mapping of tags assigned to the resource group.

Data Source: azurerm_resources

Use this data source to access information about existing resources.

Example Usage

```
# Get Resources from a Resource Group
data "azurerm_resources" "test" {
  resource_group_name = "example-resources"
}

# Get Resources with specific Tags
data "azurerm_resources" "test" {
  resource_group_name = "example-resources"

  required_tags {
    environment = "production"
    role        = "webserver"
  }
}
```

Argument Reference

NOTE: At least one of `name` , `resource_group_name` or `type` must be specified.

- `name` - (Optional) The name of the Resource.
- `resource_group_name` - (Optional) The name of the Resource group where the Resources are located.
- `type` - (Optional) The Resource Type of the Resources you want to list (e.g. `Microsoft.Network/virtualNetworks`). A full list of available Resource Types can be found here (<https://docs.microsoft.com/en-us/azure/azure-resource-manager/azure-services-resource-providers>).
- `required_tags` - (Optional) A mapping of tags which the resource has to have in order to be included in the result.

Attributes Reference

- `resources` - One or more `resource` blocks as defined below.

The `resource` block contains:

- `name` - The name of this resource.
- `id` - The Resource ID of this resource.
- `type` - The type of this resource.

- `location` - The location of this resource.
- `tags` - The type of resource that this is, such as `Microsoft.Network/virtualNetworks`.

Data Source: azurerm_role_definition

Use this data source to access information about an existing Role Definition.

Example Usage

```
data "azurerm_subscription" "primary" {}

resource "azurerm_role_definition" "custom" {
  role_definition_id = "00000000-0000-0000-0000-000000000000"
  name               = "CustomRoleDef"
  scope              = "${data.azurerm_subscription.primary.id}"

  # ...
}

data "azurerm_role_definition" "custom" {
  role_definition_id = "${azurerm_role_definition.custom.role_definition_id}"
  scope              = "${data.azurerm_subscription.primary.id}" # /subscriptions/00000000-0000-0000-0000-000000000000
}

data "azurerm_role_definition" "custom-byname" {
  name = "${azurerm_role_definition.custom.name}"
  scope = "${data.azurerm_subscription.primary.id}"
}

data "azurerm_builtin_role_definition" "builtin" {
  name = "Contributor"
}

output "custom_role_definition_id" {
  value = "${data.azurerm_role_definition.custom.id}"
}

output "contributor_role_definition_id" {
  value = "${data.azurerm_role_definition.builtin.id}"
}
```

Argument Reference

- `name` - (Optional) Specifies the Name of either a built-in or custom Role Definition.

You can also use this for built-in roles such as `Contributor`, `Owner`, `Reader` and `Virtual Machine Contributor`

- `role_definition_id` - (Optional) Specifies the ID of the Role Definition as a UUID/GUID.
- `scope` - (Optional) Specifies the Scope at which the Custom Role Definition exists.

NOTE: One of `name` or `role_definition_id` must be specified.

Attributes Reference

- `id` - the ID of the built-in Role Definition.
- `description` - the Description of the built-in Role.
- `type` - the Type of the Role.
- `permissions` - a `permissions` block as documented below.
- `assignable_scopes` - One or more assignable scopes for this Role Definition, such as `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333` , `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup` , or `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup/providers/Microsoft.Compute/virtualMachines/myVM` .

A `permissions` block contains:

- `actions` - a list of actions supported by this role
- `not_actions` - a list of actions which are denied by this role

Data Source: azurerm_route_table

Use this data source to access information about an existing Route Table.

Example Usage

```
data "azurerm_route_table" "test" {
  name                = "myroutetable"
  resource_group_name = "some-resource-group"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Route Table.
- `resource_group_name` - (Required) The name of the Resource Group in which the Route Table exists.

Attributes Reference

The following attributes are exported:

- `id` - The Route Table ID.
- `location` - The Azure Region in which the Route Table exists.
- `route` - One or more `route` blocks as documented below.
- `subnets` - The collection of Subnets associated with this route table.
- `tags` - A mapping of tags assigned to the Route Table.

The `route` block exports the following:

- `name` - The name of the Route.
- `address_prefix` - The destination CIDR to which the route applies.
- `next_hop_type` - The type of Azure hop the packet should be sent to.
- `next_hop_in_ip_address` - Contains the IP address packets should be forwarded to.

Data Source: azurerm_scheduler_job_collection

Use this data source to access information about an existing Scheduler Job Collection.

NOTE: Support for Scheduler Job Collections has been deprecated by Microsoft in favour of Logic Apps (more information can be found at this link (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such we plan to remove support for this data source as a part of version 2.0 of the AzureRM Provider.

Example Usage

```
data "azurerm_scheduler_job_collection" "test" {
  name                = "tfex-job-collection"
  resource_group_name = "tfex-job-collection-rg"
}

output "job_collection_state" {
  value = "${data.azurerm_scheduler_job_collection.jobs.state}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Scheduler Job Collection.
- `resource_group_name` - (Required) Specifies the name of the resource group in which the Scheduler Job Collection resides.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Scheduler Job Collection.
- `location` - The Azure location where the resource exists.
- `tags` - A mapping of tags assigned to the resource.
- `sku` - The Job Collection's pricing level's SKU.
- `state` - The Job Collection's state.
- `quota` - The Job collection quotas as documented in the `quota` block below.

The `quota` block supports:

- `max_job_count` - Sets the maximum number of jobs in the collection.

- `max_recurrence_frequency` - The maximum frequency of recurrence.
- `max_retry_interval` - The maximum interval between retries.

Data Source:

azurerm_servicebus_namespace_authorization_rule

Use this data source to access information about an existing ServiceBus Namespace Authorization Rule.

Example Usage

```
data "azurerm_servicebus_namespace_authorization_rule" "example" {
  name                = "examplerule"
  namespace_name      = "examplnamespace"
  resource_group_name = "example-resources"
}

output "rule_id" {
  value = "${data.azurerm_servicebus_namespace_authorization_rule.example.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the ServiceBus Namespace Authorization Rule.
- `namespace_name` - (Required) Specifies the name of the ServiceBus Namespace.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the ServiceBus Namespace exists.

Attributes Reference

- `id` - The id of the the ServiceBus Namespace Authorization Rule.
- `primary_connection_string` - The primary connection string for the authorization rule.
- `primary_key` - The primary access key for the authorization rule.
- `secondary_connection_string` - The secondary connection string for the authorization rule.
- `secondary_key` - The secondary access key for the authorization rule.

Data Source: azurerm_servicebus_namespace

Use this data source to access information about an existing ServiceBus Namespace.

Example Usage

```
data "azurerm_servicebus_namespace" "test" {
  name                = "examplnamespace"
  resource_group_name = "example-resources"
}

output "location" {
  value = "${data.azurerm_servicebus_namespace.test.location}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the ServiceBus Namespace.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the ServiceBus Namespace exists.

Attributes Reference

- `location` - The location of the Resource Group in which the ServiceBus Namespace exists.
- `sku` - The Tier used for the ServiceBus Namespace.
- `capacity` - The capacity of the ServiceBus Namespace.
- `zone_redundant` - Whether or not this ServiceBus Namespace is zone redundant.
- `tags` - A mapping of tags assigned to the resource.

The following attributes are exported only if there is an authorization rule named `RootManageSharedAccessKey` which is created automatically by Azure.

- `default_primary_connection_string` - The primary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_connection_string` - The secondary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_primary_key` - The primary access key for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_key` - The secondary access key for the authorization rule `RootManageSharedAccessKey` .

Data Source: azurerm_shared_image_gallery

Use this data source to access information about an existing Shared Image Gallery.

Example Usage

```
data "azurerm_shared_image_gallery" "test" {
  name                = "my-image-gallery"
  resource_group_name = "example-resources"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Shared Image Gallery.
- `resource_group_name` - (Required) The name of the Resource Group in which the Shared Image Gallery exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Shared Image Gallery.
- `description` - A description for the Shared Image Gallery.
- `unique_name` - The unique name assigned to the Shared Image Gallery.
- `tags` - A mapping of tags which are assigned to the Shared Image Gallery.

Data Source: azure_rm_shared_image

Use this data source to access information about an existing Shared Image within a Shared Image Gallery.

Example Usage

```
data "azure_rm_shared_image" "test" {
  name           = "my-image"
  gallery_name   = "my-image-gallery"
  resource_group = "example-resources"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Shared Image.
- `gallery_name` - (Required) The name of the Shared Image Gallery in which the Shared Image exists.
- `resource_group_name` - (Required) The name of the Resource Group in which the Shared Image Gallery exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Shared Image.
- `description` - The description of this Shared Image.
- `eula` - The End User Licence Agreement for the Shared Image.
- `location` - The supported Azure location where the Shared Image Gallery exists.
- `identifier` - An `identifier` block as defined below.
- `os_type` - The type of Operating System present in this Shared Image.
- `privacy_statement_uri` - The URI containing the Privacy Statement for this Shared Image.
- `release_note_uri` - The URI containing the Release Notes for this Shared Image.
- `tags` - A mapping of tags assigned to the Shared Image.

A `identifier` block exports the following:

- `offer` - The Offer Name for this Shared Image.
- `publisher` - The Publisher Name for this Gallery Image.

- sku - The Name of the SKU for this Gallery Image.

Data Source: azurerm_shared_image_version

Use this data source to access information about an existing Version of a Shared Image within a Shared Image Gallery.

Example Usage

```
data "azurerm_shared_image_version" "test" {
  name           = "1.0.0"
  image_name     = "my-image"
  gallery_name   = "my-image-gallery"
  resource_group_name = "example-resources"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Image Version.
- `image_name` - (Required) The name of the Shared Image in which this Version exists.
- `gallery_name` - (Required) The name of the Shared Image in which the Shared Image exists.
- `resource_group_name` - (Required) The name of the Resource Group in which the Shared Image Gallery exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Shared Image.
- `exclude_from_latest` - Is this Image Version excluded from the `latest` filter?
- `location` - The supported Azure location where the Shared Image Gallery exists.
- `managed_image_id` - The ID of the Managed Image which was the source of this Shared Image Version.
- `target_region` - One or more `target_region` blocks as documented below.
- `tags` - A mapping of tags assigned to the Shared Image.

The `target_region` block exports the following:

- `name` - The Azure Region in which this Image Version exists.
- `regional_replica_count` - The number of replicas of the Image Version to be created per region.

Data Source: azurerm_snapshot

Use this data source to access information about an existing Snapshot.

Example Usage

```
data "azurerm_snapshot" "test" {  
  name = "my-snapshot"  
  resource_group_name = "my-resource-group"  
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Snapshot.
- `resource_group_name` - (Required) Specifies the name of the resource group the Snapshot is located in.

Attributes Reference

- `id` - The ID of the Snapshot.
- `create_option` - How the snapshot was created.
- `source_uri` - The URI to a Managed or Unmanaged Disk.
- `source_resource_id` - The reference to an existing snapshot.
- `storage_account_id` - The ID of an storage account.
- `disk_size_gb` - The size of the Snapshotted Disk in GB.

Data Source: azurerm_sql_database

Use this data source to access information about an existing SQL Azure Database.

Example Usage

```
data "azurerm_sql_database" "example" {
  name                = "example_db"
  server_name         = "example_db_server"
  resource_group_name = "example-resources"
}

output "sql_database_id" {
  value = data.azurerm_sql_database.example.id
}
```

Argument Reference

- `name` - (Required) The name of the SQL Database.
- `server_name` - (Required) The name of the SQL Server.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the Azure SQL Database exists.

Attributes Reference

- `id` - The SQL Database ID.
- `collation` - The name of the collation.
- `creation_date` - The creation date of the SQL Database.
- `default_secondary_location` - The default secondary location of the SQL Database.
- `edition` - The edition of the database.
- `elastic_pool_name` - The name of the elastic database pool the database belongs to.
- `failover_group_id` - The ID of the failover group the database belongs to.
- `location` - The location of the Resource Group in which the SQL Server exists.
- `name` - The name of the database.
- `read_scale` - Indicate if read-only connections will be redirected to a high-available replica.
- `requested_service_objective_id` - The ID pertaining to the performance level of the database.
- `requested_service_objective_name` - The name pertaining to the performance level of the database.

- `resource_group_name` - The name of the resource group in which the database resides. This will always be the same resource group as the Database Server.
- `server_name` - The name of the SQL Server on which to create the database.
- `tags` - A mapping of tags assigned to the resource.

Data Source: azurerm_sql_server

Use this data source to access information about an existing SQL Azure Database Server.

Example Usage

```
data "azurerm_sql_server" "example" {
  name                = "examplesqlservername"
  resource_group_name = "example-resources"
}

output "sql_server_id" {
  value = "${data.azurerm_sql_server.test.id}"
}
```

Argument Reference

- `name` - (Required) The name of the SQL Server.
- `resource_group_name` - (Required) Specifies the name of the Resource Group where the SQL Server exists.

Attributes Reference

- `location` - The location of the Resource Group in which the SQL Server exists.
- `fqdn` - The fully qualified domain name of the SQL Server.
- `version` - The version of the SQL Server.
- `administrator_login` - The administrator username of the SQL Server.
- `tags` - A mapping of tags assigned to the resource.

Data Source:

azurerm_storage_account_blob_container_sas

Use this data source to obtain a Shared Access Signature (SAS Token) for an existing Storage Account Blob Container.

Shared access signatures allow fine-grained, ephemeral access control to various aspects of an Azure Storage Account Blob Container.

Example Usage

```

resource "azurerm_resource_group" "rg" {
  name      = "resourceGroupName"
  location = "westus"
}

resource "azurerm_storage_account" "storage" {
  name                        = "storageaccountname"
  resource_group_name        = "${azurerm_resource_group.rg.name}"
  location                   = "${azurerm_resource_group.rg.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "container" {
  name                = "mycontainer"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  storage_account_name = "${azurerm_storage_account.storage.name}"
  container_access_type = "private"
}

data "azurerm_storage_account_blob_container_sas" "test" {
  connection_string = "${azurerm_storage_account.storage.primary_connection_string}"
  container_name     = "${azurerm_storage_container.container.name}"
  https_only         = true

  ip_address = "168.1.5.65"

  start = "2018-03-21"
  expiry = "2018-03-21"

  permissions {
    read   = true
    add    = true
    create = false
    write  = false
    delete = true
    list   = true
  }

  cache_control       = "max-age=5"
  content_disposition = "inline"
  content_encoding    = "deflate"
  content_language    = "en-US"
  content_type        = "application/json"
}

output "sas_url_query_string" {
  value = "${data.azurerm_storage_account_blob_container_sas.test.sas}"
}

```

Argument Reference

- `connection_string` - (Required) The connection string for the storage account to which this SAS applies. Typically directly from the `primary_connection_string` attribute of a terraform created `azurerm_storage_account` resource.
 - `container_name` - (Required) Name of the container.
 - `https_only` - (Optional) Only permit `https` access. If `false`, both `http` and `https` are permitted. Defaults to `true`.
 - `ip_address` - (Optional) Single ipv4 address or range (connected with a dash) of ipv4 addresses.
 - `start` - (Required) The starting time and date of validity of this SAS. Must be a valid ISO-8601 format time/date string.
 - `expiry` - (Required) The expiration time and date of this SAS. Must be a valid ISO-8601 format time/date string.
 - `permissions` - (Required) A `permissions` block as defined below.
 - `cache_control` - (Optional) The `Cache-Control` response header that is sent when this SAS token is used.
 - `content_disposition` - (Optional) The `Content-Disposition` response header that is sent when this SAS token is used.
 - `content_encoding` - (Optional) The `Content-Encoding` response header that is sent when this SAS token is used.
 - `content_language` - (Optional) The `Content-Language` response header that is sent when this SAS token is used.
 - `content_type` - (Optional) The `Content-Type` response header that is sent when this SAS token is used.
-

A `permissions` block contains:

- `read` - (Required) Should Read permissions be enabled for this SAS?
- `add` - (Required) Should Add permissions be enabled for this SAS?
- `create` - (Required) Should Create permissions be enabled for this SAS?
- `write` - (Required) Should Write permissions be enabled for this SAS?
- `delete` - (Required) Should Delete permissions be enabled for this SAS?
- `list` - (Required) Should List permissions be enabled for this SAS?

Refer to the SAS creation reference from Azure (<https://docs.microsoft.com/en-us/rest/api/storageservices/create-service-sas>) for additional details on the fields above.

Attributes Reference

- `sas` - The computed Blob Container Shared Access Signature (SAS).

Data Source: azurerm_storage_account

Use this data source to access information about an existing Storage Account.

Example Usage

```
data "azurerm_storage_account" "test" {
  name                = "packerimages"
  resource_group_name = "packer-storage"
}

output "storage_account_tier" {
  value = "${data.azurerm_storage_account.test.account_tier}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Storage Account
- `resource_group_name` - (Required) Specifies the name of the resource group the Storage Account is located in.

Attributes Reference

- `id` - The ID of the Storage Account.
- `location` - The Azure location where the Storage Account exists
- `account_kind` - The Kind of account.
- `account_tier` - The Tier of this storage account.
- `account_replication_type` - The type of replication used for this storage account.
- `access_tier` - The access tier for `BlobStorage` accounts.
- `enable_blob_encryption` - Are Encryption Services are enabled for Blob storage? See [here](https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/) (<https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/>) for more information.
- `enable_file_encryption` - Are Encryption Services are enabled for File storage? See [here](https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/) (<https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/>) for more information.
- `enable_https_traffic_only` - Is traffic only allowed via HTTPS? See [here](https://docs.microsoft.com/en-us/azure/storage/storage-require-secure-transfer/) (<https://docs.microsoft.com/en-us/azure/storage/storage-require-secure-transfer/>) for more information.
- `is_hns_enabled` - Is Hierarchical Namespace enabled?
- `account_encryption_source` - The Encryption Source for this Storage Account.
- `custom_domain` - A `custom_domain` block as documented below.

- `tags` - A mapping of tags to assigned to the resource.
- `primary_location` - The primary location of the Storage Account.
- `secondary_location` - The secondary location of the Storage Account.
- `primary_blob_endpoint` - The endpoint URL for blob storage in the primary location.
- `primary_blob_host` - The hostname with port if applicable for blob storage in the primary location.
- `secondary_blob_endpoint` - The endpoint URL for blob storage in the secondary location.
- `secondary_blob_host` - The hostname with port if applicable for blob storage in the secondary location.
- `primary_queue_endpoint` - The endpoint URL for queue storage in the primary location.
- `primary_queue_host` - The hostname with port if applicable for queue storage in the primary location.
- `secondary_queue_endpoint` - The endpoint URL for queue storage in the secondary location.
- `secondary_queue_host` - The hostname with port if applicable for queue storage in the secondary location.
- `primary_table_endpoint` - The endpoint URL for table storage in the primary location.
- `primary_table_host` - The hostname with port if applicable for table storage in the primary location.
- `secondary_table_endpoint` - The endpoint URL for table storage in the secondary location.
- `secondary_table_host` - The hostname with port if applicable for table storage in the secondary location.
- `primary_file_endpoint` - The endpoint URL for file storage in the primary location.
- `primary_file_host` - The hostname with port if applicable for file storage in the primary location.
- `secondary_file_endpoint` - The endpoint URL for file storage in the secondary location.
- `secondary_file_host` - The hostname with port if applicable for file storage in the secondary location.
- `primary_dfs_endpoint` - The endpoint URL for DFS storage in the primary location.
- `primary_dfs_host` - The hostname with port if applicable for DFS storage in the primary location.
- `secondary_dfs_endpoint` - The endpoint URL for DFS storage in the secondary location.
- `secondary_dfs_host` - The hostname with port if applicable for DFS storage in the secondary location.
- `primary_web_endpoint` - The endpoint URL for web storage in the primary location.
- `primary_web_host` - The hostname with port if applicable for web storage in the primary location.
- `secondary_web_endpoint` - The endpoint URL for web storage in the secondary location.
- `secondary_web_host` - The hostname with port if applicable for web storage in the secondary location.
- `primary_access_key` - The primary access key for the Storage Account.
- `secondary_access_key` - The secondary access key for the Storage Account.
- `primary_connection_string` - The connection string associated with the primary location
- `secondary_connection_string` - The connection string associated with the secondary location

- `primary_blob_connection_string` - The connection string associated with the primary blob location
- `secondary_blob_connection_string` - The connection string associated with the secondary blob location

~> **NOTE:** If there's a Write Lock on the Storage Account, or the account doesn't have permission then these fields will have an empty value due to a bug in the Azure API
(<https://github.com/Azure/azure-rest-api-specs/issues/6363>)

- `custom_domain` supports the following:
- `name` - The Custom Domain Name used for the Storage Account.

Data Source: azurerm_storage_account_sas

Use this data source to obtain a Shared Access Signature (SAS Token) for an existing Storage Account.

Shared access signatures allow fine-grained, ephemeral access control to various aspects of an Azure Storage Account.

Note that this is an Account SAS (<https://docs.microsoft.com/en-us/rest/api/storageservices/constructing-an-account-sas>) and *not* a Service SAS (<https://docs.microsoft.com/en-us/rest/api/storageservices/constructing-a-service-sas>).

Example Usage

```

resource "azurerm_resource_group" "testrg" {
  name      = "resourceGroupName"
  location = "westus"
}

resource "azurerm_storage_account" "testsa" {
  name                        = "storageaccountname"
  resource_group_name        = "${azurerm_resource_group.testrg.name}"
  location                   = "westus"
  account_tier               = "Standard"
  account_replication_type   = "GRS"

  tags = {
    environment = "staging"
  }
}

data "azurerm_storage_account_sas" "test" {
  connection_string = "${azurerm_storage_account.testsa.primary_connection_string}"
  https_only        = true

  resource_types {
    service  = true
    container = false
    object   = false
  }

  services {
    blob  = true
    queue = false
    table = false
    file  = false
  }

  start  = "2018-03-21"
  expiry = "2020-03-21"

  permissions {
    read    = true
    write   = true
    delete = false
    list    = false
    add     = true
    create  = true
    update  = false
    process = false
  }
}

output "sas_url_query_string" {
  value = "${data.azurerm_storage_account_sas.test.sas}"
}

```

Argument Reference

- `connection_string` - (Required) The connection string for the storage account to which this SAS applies. Typically directly from the `primary_connection_string` attribute of a terraform created `azurerm_storage_account` resource.
 - `https_only` - (Optional) Only permit `https` access. If `false`, both `http` and `https` are permitted. Defaults to `true`.
 - `resource_types` - (Required) A `resource_types` block as defined below.
 - `services` - (Required) A `services` block as defined below.
 - `start` - (Required) The starting time and date of validity of this SAS. Must be a valid ISO-8601 format time/date string.
 - `expiry` - (Required) The expiration time and date of this SAS. Must be a valid ISO-8601 format time/date string.
 - `permissions` - (Required) A `permissions` block as defined below.
-

`resource_types` is a set of `true / false` flags which define the storage account resource types that are granted access by this SAS. This can be thought of as the scope over which the permissions apply. A `service` will have larger scope (affecting all sub-resources) than `object`.

A `resource_types` block contains:

- `service` - (Required) Should permission be granted to the entire service?
 - `container` - (Required) Should permission be granted to the container?
 - `object` - (Required) Should permission be granted only to a specific object?
-

`services` is a set of `true / false` flags which define the storage account services that are granted access by this SAS.

A `services` block contains:

- `blob` - (Required) Should permission be granted to `blob` services within this storage account?
 - `queue` - (Required) Should permission be granted to `queue` services within this storage account?
 - `table` - (Required) Should permission be granted to `table` services within this storage account?
 - `file` - (Required) Should permission be granted to `file` services within this storage account?
-

A `permissions` block contains:

- `read` - (Required) Should Read permissions be enabled for this SAS?
- `write` - (Required) Should Write permissions be enabled for this SAS?
- `delete` - (Required) Should Delete permissions be enabled for this SAS?
- `list` - (Required) Should List permissions be enabled for this SAS?
- `add` - (Required) Should Add permissions be enabled for this SAS?
- `create` - (Required) Should Create permissions be enabled for this SAS?
- `update` - (Required) Should Update permissions be enabled for this SAS?

- `process` - (Required) Should Process permissions be enabled for this SAS?

Refer to the SAS creation reference from Azure (<https://docs.microsoft.com/en-us/rest/api/storageservices/constructing-an-account-sas>) for additional details on the fields above.

Attributes Reference

- `sas` - The computed Account Shared Access Signature (SAS).

Data Source: azurerm_storage_management_policy

Use this data source to access information about an existing Storage Management Policy.

Example Usage

```
data "azurerm_storage_account" "test" {
  name                = "storageaccountname"
  resource_group_name = "resourcegroupname"
}

data "azurerm_storage_management_policy" "testpolicy" {
  storage_account_id = "${azurerm_storage_account.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `storage_account_id` - (Required) Specifies the id of the storage account to retrieve the management policy for.

Attributes Reference

- `id` - The ID of the Management Policy.
- `rule` - A `rule` block as documented below.

- `rule` supports the following:
 - `name` - (Required) A rule name can contain any combination of alpha numeric characters. Rule name is case-sensitive. It must be unique within a policy.
 - `enabled` - (Required) Boolean to specify whether the rule is enabled.
 - `filters` - A `filter` block as documented below.
 - `actions` - An `actions` block as documented below.

`filters` supports the following:

- `prefix_match` - An array of strings for prefixes to be matched.
- `blob_types` - An array of predefined values. Only `blockBlob` is supported.

`actions` supports the following:

- `base_blob` - A `base_blob` block as documented below.
 - `snapshot` - A `snapshot` block as documented below.
-

`base_blob` supports the following:

- `tier_to_cool_after_days_since_modification_greater_than` - The age in days after last modification to tier blobs to cool storage. Supports blob currently at Hot tier.
 - `tier_to_archive_after_days_since_modification_greater_than` - The age in days after last modification to tier blobs to archive storage. Supports blob currently at Hot or Cool tier.
 - `delete_after_days_since_modification_greater_than` - The age in days after last modification to delete the blob.
-

`snapshot` supports the following:

- `delete_after_days_since_creation_greater_than` - The age in days after create to delete the snapshot.

Data Source: azurerm_stream_analytics_job

Use this data source to access information about an existing Stream Analytics Job.

Example Usage

```
data "azurerm_stream_analytics_job" "test" {
  name                = "example-job"
  resource_group_name = "example-resources"
}

output "job_id" {
  value = "${data.azurerm_stream_analytics_job.test.job_id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Stream Analytics Job.
- `resource_group_name` - (Required) Specifies the name of the resource group the Stream Analytics Job is located in.

Attributes Reference

- `id` - The ID of the Stream Analytics Job.
- `compatibility_level` - The compatibility level for this job.
- `data_locale` - The Data Locale of the Job.
- `events_late_arrival_max_delay_in_seconds` - The maximum tolerable delay in seconds where events arriving late could be included.
- `events_out_of_order_max_delay_in_seconds` - The maximum tolerable delay in seconds where out-of-order events can be adjusted to be back in order.
- `events_out_of_order_policy` - The policy which should be applied to events which arrive out of order in the input event stream.
- `job_id` - The Job ID assigned by the Stream Analytics Job.
- `location` - The Azure location where the Stream Analytics Job exists.
- `output_error_policy` - The policy which should be applied to events which arrive at the output and cannot be written to the external storage due to being malformed (such as missing column values, column values of wrong type or size).
- `streaming_units` - The number of streaming units that the streaming job uses.
- `transformation_query` - The query that will be run in the streaming job, written in Stream Analytics Query Language

(SAQL) (<https://msdn.microsoft.com/library/azure/dn834998>).

Data Source: azurerm_subnet

Use this data source to access information about an existing Subnet within a Virtual Network.

Example Usage

```
data "azurerm_subnet" "test" {
  name                = "backend"
  virtual_network_name = "production"
  resource_group_name = "networking"
}

output "subnet_id" {
  value = "${data.azurerm_subnet.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Subnet.
- `virtual_network_name` - (Required) Specifies the name of the Virtual Network this Subnet is located within.
- `resource_group_name` - (Required) Specifies the name of the resource group the Virtual Network is located in.

Attributes Reference

- `id` - The ID of the Subnet.
- `address_prefix` - The address prefix used for the subnet.
- `network_security_group_id` - The ID of the Network Security Group associated with the subnet.
- `route_table_id` - The ID of the Route Table associated with this subnet.
- `ip_configurations` - The collection of IP Configurations with IPs within this subnet.
- `service_endpoints` - A list of Service Endpoints within this subnet.

Data Source: azurerm_subscription

Use this data source to access information about an existing Subscription.

Example Usage

```
data "azurerm_subscription" "current" {}

output "current_subscription_display_name" {
  value = "${data.azurerm_subscription.current.display_name}"
}
```

Argument Reference

- `subscription_id` - (Optional) Specifies the ID of the subscription. If this argument is omitted, the subscription ID of the current Azure Resource Manager provider is used.

Attributes Reference

- `id` - The ID of the subscription.
- `subscription_id` - The subscription GUID.
- `display_name` - The subscription display name.
- `tenant_id` - The subscription tenant ID.
- `state` - The subscription state. Possible values are Enabled, Warned, PastDue, Disabled, and Deleted.
- `location_placement_id` - The subscription location placement ID.
- `quota_id` - The subscription quota ID.
- `spending_limit` - The subscription spending limit.

Data Source: azurerm_subscriptions

Use this data source to access information about all the Subscriptions currently available.

Example Usage

```
data "azurerm_subscriptions" "available" {}

output "available_subscriptions" {
  value = "${data.azurerm_subscriptions.available.subscriptions}"
}

output "first_available_subscription_display_name" {
  value = "${data.azurerm_subscriptions.available.subscriptions.0.display_name}"
}
```

Argument Reference

- `display_name_prefix` - (Optional) A case-insensitive prefix which can be used to filter on the `display_name` field
- `display_name_contains` - (Optional) A case-insensitive value which must be contained within the `display_name` field, used to filter the results

Attributes Reference

- `subscriptions` - One or more `subscription` blocks as defined below.

The `subscription` block contains:

- `subscription_id` - The subscription GUID.
- `display_name` - The subscription display name.
- `tenant_id` - The subscription tenant ID.
- `state` - The subscription state. Possible values are Enabled, Warned, PastDue, Disabled, and Deleted.
- `location_placement_id` - The subscription location placement ID.
- `quota_id` - The subscription quota ID.
- `spending_limit` - The subscription spending limit.

Data Source:

azurerm_traffic_manager_geographical_location

Use this data source to access the ID of a specified Traffic Manager Geographical Location within the Geographical Hierarchy.

Example Usage (World)

```
data "azurerm_traffic_manager_geographical_location" "test" {
  name = "World"
}

output "location_code" {
  value = "${data.azurerm_traffic_manager_geographical_location.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Location, for example `World`, `Europe` or `Germany` .

Attributes Reference

- `id` - The ID of this Location, also known as the `Code` of this Location.

Data Source: azurerm_user_assigned_identity

Use this data source to access information about an existing User Assigned Identity.

Example Usage (reference an existing)

```
data "azurerm_user_assigned_identity" "example" {
  name                = "name_of_user_assigned_identity"
  resource_group_name = "name_of_resource_group"
}

output "uai_client_id" {
  value = "${data.azurerm_user_assigned_identity.example.client_id}"
}

output "uai_principal_id" {
  value = "${data.azurerm_user_assigned_identity.example.principal_id}"
}
```

Argument Reference

- `name` - (Required) The name of the User Assigned Identity.
- `resource_group_name` - (Required) The name of the Resource Group in which the User Assigned Identity exists.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the User Assigned Identity.
- `location` - The Azure location where the User Assigned Identity exists.
- `principal_id` - The Service Principal ID of the User Assigned Identity.
- `client_id` - The Client ID of the User Assigned Identity.
- `tags` - A mapping of tags assigned to the User Assigned Identity.

Data Source: azurerm_virtual_machine

Use this data source to access information about an existing Virtual Machine.

Example Usage

```
data "azurerm_virtual_machine" "test" {
  name                = "production"
  resource_group_name = "networking"
}

output "virtual_machine_id" {
  value = "${data.azurerm_virtual_machine.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Virtual Machine.
- `resource_group_name` - (Required) Specifies the name of the resource group the Virtual Machine is located in.

Attributes Reference

- `id` - The ID of the Virtual Machine.

Data Source:

azurerm_virtual_network_gateway_connection

Use this data source to access information about an existing Virtual Network Gateway Connection.

Example Usage

```
data "azurerm_virtual_network_gateway_connection" "test" {
  name                = "production"
  resource_group_name = "networking"
}

output "virtual_network_gateway_connection_id" {
  value = "${data.azurerm_virtual_network_gateway_connection.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Virtual Network Gateway Connection.
- `resource_group_name` - (Required) Specifies the name of the resource group the Virtual Network Gateway Connection is located in.

Attributes Reference

- `id` - The ID of the Virtual Network Gateway Connection.
- `location` - The location/region where the connection is located.
- `type` - The type of connection. Valid options are `IPsec` (Site-to-Site), `ExpressRoute` (ExpressRoute), and `Vnet2Vnet` (VNet-to-VNet).
- `virtual_network_gateway_id` - The ID of the Virtual Network Gateway in which the connection is created.
- `authorization_key` - The authorization key associated with the Express Route Circuit. This field is present only if the `type` is an `ExpressRoute` connection.
- `express_route_circuit_id` - The ID of the Express Route Circuit (i.e. when `type` is `ExpressRoute`).
- `peer_virtual_network_gateway_id` - The ID of the peer virtual network gateway when a VNet-to-VNet connection (i.e. when `type` is `Vnet2Vnet`).
- `local_network_gateway_id` - The ID of the local network gateway when a Site-to-Site connection (i.e. when `type` is `IPsec`).
- `routing_weight` - The routing weight.
- `shared_key` - The shared IPsec key.

- `enable_bgp` - If `true`, BGP (Border Gateway Protocol) is enabled for this connection.
- `express_route_gateway_bypass` - If `true`, data packets will bypass ExpressRoute Gateway for data forwarding. This is only valid for ExpressRoute connections.
- `use_policy_based_traffic_selectors` - If `true`, policy-based traffic selectors are enabled for this connection. Enabling policy-based traffic selectors requires an `ipsec_policy` block.
- `ipsec_policy` (Optional) A `ipsec_policy` block which is documented below. Only a single policy can be defined for a connection. For details on custom policies refer to the relevant section in the Azure documentation (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-ipsecikepolicy-rm-powershell>).
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `ipsec_policy` block supports:

- `dh_group` - The DH group used in IKE phase 1 for initial SA. Valid options are `DHGroup1`, `DHGroup14`, `DHGroup2`, `DHGroup2048`, `DHGroup24`, `ECP256`, `ECP384`, or `None`.
- `ike_encryption` - The IKE encryption algorithm. Valid options are `AES128`, `AES192`, `AES256`, `DES`, or `DES3`.
- `ike_integrity` - The IKE integrity algorithm. Valid options are `MD5`, `SHA1`, `SHA256`, or `SHA384`.
- `ipsec_encryption` - The IPSec encryption algorithm. Valid options are `AES128`, `AES192`, `AES256`, `DES`, `DES3`, `GCMAES128`, `GCMAES192`, `GCMAES256`, or `None`.
- `ipsec_integrity` - The IPSec integrity algorithm. Valid options are `GCMAES128`, `GCMAES192`, `GCMAES256`, `MD5`, `SHA1`, or `SHA256`.
- `pfs_group` - The DH group used in IKE phase 2 for new child SA. Valid options are `ECP256`, `ECP384`, `PFS1`, `PFS2`, `PFS2048`, `PFS24`, or `None`.
- `sa_datasize` - The IPSec SA payload size in KB. Must be at least 1024 KB.
- `sa_lifetime` - The IPSec SA lifetime in seconds. Must be at least 300 seconds.

Data Source: azurerm_virtual_network_gateway

Use this data source to access information about an existing Virtual Network Gateway.

Example Usage

```
data "azurerm_virtual_network_gateway" "test" {
  name                = "production"
  resource_group_name = "networking"
}

output "virtual_network_gateway_id" {
  value = "${data.azurerm_virtual_network_gateway.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Virtual Network Gateway.
- `resource_group_name` - (Required) Specifies the name of the resource group the Virtual Network Gateway is located in.

Attributes Reference

- `id` - The ID of the Virtual Network Gateway.
- `location` - The location/region where the Virtual Network Gateway is located.
- `type` - The type of the Virtual Network Gateway.
- `vpn_type` - The routing type of the Virtual Network Gateway.
- `enable_bgp` - Will BGP (Border Gateway Protocol) will be enabled for this Virtual Network Gateway.
- `active_active` - (Optional) Is this an Active-Active Gateway?
- `default_local_network_gateway_id` - The ID of the local network gateway through which outbound Internet traffic from the virtual network in which the gateway is created will be routed (*forced tunneling*). Refer to the Azure documentation on forced tunneling (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-forced-tunneling-rm>).
- `sku` - Configuration of the size and capacity of the Virtual Network Gateway.
- `ip_configuration` - One or two `ip_configuration` blocks documented below.
- `vpn_client_configuration` - A `vpn_client_configuration` block which is documented below.
- `tags` - A mapping of tags assigned to the resource.

The `ip_configuration` block supports:

- `name` - A user-defined name of the IP configuration.
- `private_ip_address_allocation` - Defines how the private IP address of the gateways virtual interface is assigned.
- `subnet_id` - The ID of the gateway subnet of a virtual network in which the virtual network gateway will be created. It is mandatory that the associated subnet is named `GatewaySubnet`. Therefore, each virtual network can contain at most a single Virtual Network Gateway.
- `public_ip_address_id` - The ID of the Public IP Address associated with the Virtual Network Gateway.

The `vpn_client_configuration` block supports:

- `address_space` - The address space out of which ip addresses for vpn clients will be taken. You can provide more than one address space, e.g. in CIDR notation.
- `root_certificate` - One or more `root_certificate` blocks which are defined below. These root certificates are used to sign the client certificate used by the VPN clients to connect to the gateway.
- `revoked_certificate` - One or more `revoked_certificate` blocks which are defined below.
- `radius_server_address` - (Optional) The address of the Radius server. This setting is incompatible with the use of `root_certificate` and `revoked_certificate`.
- `radius_server_secret` - (Optional) The secret used by the Radius server. This setting is incompatible with the use of `root_certificate` and `revoked_certificate`.
- `vpn_client_protocols` - (Optional) List of the protocols supported by the vpn client. The supported values are SSTP, IKEv2 and OpenVPN.

The `bgp_settings` block supports:

- `asn` - The Autonomous System Number (ASN) to use as part of the BGP.
- `peering_address` - The BGP peer IP address of the virtual network gateway. This address is needed to configure the created gateway as a BGP Peer on the on-premises VPN devices.
- `peer_weight` - The weight added to routes which have been learned through BGP peering.

The `root_certificate` block supports:

- `name` - The user-defined name of the root certificate.
- `public_cert_data` - The public certificate of the root certificate authority. The certificate must be provided in Base-64 encoded X.509 format (PEM).

The `root_revoked_certificate` block supports:

- `name` - The user-defined name of the revoked certificate.
- `public_cert_data` - The SHA1 thumbprint of the certificate to be revoked.

Data Source: azurerm_virtual_network

Use this data source to access information about an existing Virtual Network.

Example Usage

```
data "azurerm_virtual_network" "test" {
  name                = "production"
  resource_group_name = "networking"
}

output "virtual_network_id" {
  value = "${data.azurerm_virtual_network.test.id}"
}
```

Argument Reference

- `name` - (Required) Specifies the name of the Virtual Network.
- `resource_group_name` - (Required) Specifies the name of the resource group the Virtual Network is located in.

Attributes Reference

- `id` - The ID of the virtual network.
- `location` - Location of the virtual network.
- `address_space` - The list of address spaces used by the virtual network.
- `dns_servers` - The list of DNS servers used by the virtual network.
- `subnets` - The list of name of the subnets that are attached to this virtual network.
- `vnet_peerings` - A mapping of name - virtual network id of the virtual network peerings.

v2.0 of the AzureRM Provider

Terraform initially shipped support for the AzureRM Provider back in December 2015.

Since then we've added support for 305 Resources, 92 Data Sources and have launched a couple of related Providers in the form of the Azure Active Directory Provider (<https://www.terraform.io/docs/providers/azuread/index.html>) and the Azure Stack Provider (<https://www.terraform.io/docs/providers/azurestack/index.html>).

Version 2.0 of the AzureRM Provider is a major release and as such includes some larger-scale changes which are outlined in this document.

NOTE: This guide is a Work In Progress and additional information may be added to this guide until version 2.0 of the AzureRM Provider is released.

Pinning your Provider Version

We recommend pinning the version of each Provider you use in Terraform - you can do this using the `version` attribute in the `provider` block, either to a specific version of the AzureRM Provider, like so:

```
provider "azurerm" {  
  version = "=1.36.0"  
}
```

.. or to any 1.x release:

```
provider "azurerm" {  
  version = "~> 1.x"  
}
```

More information on how to pin the version of a Terraform Provider being used can be found on the Terraform Website (<https://www.terraform.io/docs/configuration/providers.html#provider-versions>).

Once version 2.0 of the AzureRM Provider is released - you can then upgrade to it by updating the version specified in the Provider block, like so:

```
provider "azurerm" {  
  version = "=2.0.0"  
}
```

What's coming in Version 2.0 of the AzureRM Provider?

At a high level, we're intending for version 2.0 to include the following changes:

- Changes when importing existing resources - when bringing Resources provisioned outside of Terraform under Terraform's control, these resources will now be required to be imported.

- Custom Timeouts will be available on Resources - this will allow you to specify a custom timeout for provisioning the resource in your Terraform Configuration using the `timeouts` block (<https://www.terraform.io/docs/configuration/resources.html#operation-timeouts>).
- New resources for Virtual Machines and Virtual Machine Scale Sets.
- Removing Fields, Data Sources and Resources which have been deprecated.

Each of these topics is covered in more detail below, however please note that this guide is a Work In Progress until version 2.0 of the AzureRM Provider is released and thus things may be added/changed as necessary.

Changes when Importing Existing Resources

Terraform allows for existing resources which have been created outside of Terraform to be Imported into Terraform's State. Once a resource is imported into the state, it's possible for Terraform to track changes and manage this resource. The Azure Provider allows Importing existing resources into the state (using `terraform import`) for (almost) every resource.

Version 2.0 of the Azure Provider aims to solve an issue where it's possible to unintentionally import resources into the state by running `terraform apply`. To explain this further, the majority of Azure's APIs are Upserts - which means that a resource will be updated if it exists, otherwise it'll be created.

Where the unique identifier for (most) Azure resources is the `name` (rather than for example an `aws_instance` where AWS will generate a different unique identifier) - it's possible that users may have unintentionally imported existing resources into Terraform (and made changes to the resource) when running `terraform apply` when using the same unique identifier as an existing resource.

Whilst this may allow resources to work in some cases, it leads to hard-to-diagnose bugs in others (which could have been caught during `terraform plan`).

In order to match the behaviour of other Terraform Providers version 2.0 of the AzureRM Provider will require that existing resources are imported into the state prior to use. This means that Terraform will be checking for the presence of an existing resource prior to creating it - and will return an error similar to below:

```
A resource with the ID /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1 already exists - to be managed via Terraform this resource needs to be imported into the State. Please see the resource documentation for `azurerm_resource_group` for more information.
```

Information on how to import a given Resource can be found in the documentation for that Resource - for example here's how to here's how to import a Resource Group in Terraform ([/docs/providers/azurerm/r/resource_group.html#import](/docs/providers/azurerm/r/resource_group.html#import)).

You can opt into this behaviour from version 1.22 of the AzureRM Provider by setting the Environment Variable `ARM_PROVIDER_STRICT` to `true`.

Custom Timeouts for Resources

Resources can optionally support a `timeouts` block - which allows users to specify a Custom Timeout for resource creation/deletion as part of the Terraform Configuration.

Prior to version 2.0 the Azure Provider has a default value set for resource timeouts for an hour - which cannot be overridden. This works for the most-part but there are certain scenarios where it'd be helpful to override this.

This is useful for resources which can take a long time to delete - for example deleting the `azurerm_resource_group` resource will delete any resources within it, which can take time. Within Terraform your Terraform Configuration this could be represented like so:

```
resource "azurerm_resource_group" "test" {
  name      = "example-resource-group"
  location  = "West Europe"

  timeouts {
    create = "10m"
    delete = "30m"
  }
}
```

We intend to support the `timeouts` block in version 2.0 of the Azure Provider - which will allow timeouts to be specified on resources (as shown above). This feature request is being tracked here (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/171>) and will form part of the 2.0 release of the AzureRM Provider.

New Resources for Virtual Machines and Virtual Machine Scale Sets

We originally shipped support for the `azurerm_virtual_machine` and `azurerm_virtual_machine_scale_set` resources back in March 2016.

Over time new features have been added to these resources by Azure, such as Managed Disks and Managed Service Identity which these resources support. Since these resources first launched Azure's also changed the behaviour of some fields, so that it's now possible to update them where this wasn't previously possible - for example the Custom Data for a Virtual Machine.

We've spent some time thinking about how we can accommodate these changes and about how we can improve the user experience of both resources. In particular we've wanted to be able to give better validation during `terraform plan`, rather than bailing out with an Azure API error during `terraform apply`, however this isn't possible with the current resource structure since they're very generic. The validation requirements also vary substantially based on the fields provided, for example the `name` field for a Virtual Machine can be up to 63 characters for a Linux Virtual Machine but only allows 15 characters for a Windows Virtual Machine.

As such after spending some time reading through bug reports and thinking/prototyping some potential solutions to this - we believe the best path forward here is to split these resources out, so that we would have:

- a Linux Virtual Machine Resource (working name: `azurerm_linux_virtual_machine`)
- a Windows Virtual Machine Resource (working name: `azurerm_windows_virtual_machine`)
- updating the Data Disk Attachment Resource to support Unmanaged Disks
- a Linux Virtual Machine Scale Set Resource (working name: `azurerm_linux_virtual_machine_scale_set`)
- a Windows Virtual Machine Scale Set Resource (working name: `azurerm_windows_virtual_machine_scale_set`)
- a separate resource for Virtual Machine Scale Set Extensions (working name `azurerm_virtual_machine_scale_set_extension`)

Please Note: all of the resources mentioned above currently do not exist but will form part of the 2.0 release.

Whilst we're aware that this isn't ideal since users will eventually have to update their code/import an existing resource - we believe this approach gives us a good footing for the future. In particular this allows us to re-consider the schema design so that we can both support these new use-cases, fix some bugs and improve the user experience with these resources.

The existing `azurerm_virtual_machine` and `azurerm_virtual_machine_scale_set` resources would continue to be available throughout the 2.x releases - but over time we'd end up deprecating these in favour of the new resources.

Removal of Deprecated Fields, Data Sources and Resources

The Azure Provider was first launched in December 2015 - over the years the Provider has accumulated fields which have been deprecated (either by Azure because they're no longer used or because we've created an external resource to use instead).

Since version 2.0 is a major version - we intend to take this opportunity to remove deprecated Fields, Resources and Data Sources.

Below we'll cover each of the Data Sources and Resources which will be affected by the 2.0 upgrade.

NOTE: This guide is a Work In Progress and as such Fields, Resources and Data Sources may be added to this guide until version 2.0 of the AzureRM Provider is released.

Data Source: `azurerm_azuread_application`

The AzureAD Data Sources and Resources have been moved to the new AzureAD Provider (<https://www.terraform.io/docs/providers/azuread/index.html>), as such this Data Source will be removed in v2.0 of the AzureRM Provider.

A guide on how to migrate to using the new Provider can be found here (<https://www.terraform.io/docs/providers/azurerm/guides/migrating-to-azuread.html>).

Data Source: `azurerm_azuread_service_principal`

The AzureAD Data Sources and Resources have been moved to the new AzureAD Provider (<https://www.terraform.io/docs/providers/azuread/index.html>), as such this Data Source will be removed in v2.0 of the AzureRM Provider.

A guide on how to migrate to using the new Provider can be found here (<https://www.terraform.io/docs/providers/azurerm/guides/migrating-to-azuread.html>).

Data Source: `azurerm_api_management`

The deprecated field `sku` will be removed in favour of the `sku_name` property.

Data Source: `azurerm_client_config`

The deprecated fields `service_principal_object_id` and `service_principal_application_id` will be removed. They have been replaced with the `object_id` and `client_id` respectively.

Data Source: `azurerm_dns_zone`

The deprecated field `zone_type` will be removed. Private DNS Zones are now a separate resource in both Azure and Terraform.

Data Source: `azurerm_kubernetes_cluster`

The deprecated field `dns_prefix` within the `agent_pool_profile` block will be removed.

Data Source: `azurerm_network_interface`

The deprecated field `internal_fqdn` will be removed.

Data Source: `azurerm_scheduler_job_collection`

Azure Scheduler is being retired in favour of Logic Apps (more information can be found here (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such this Data Source will be removed.

Data Source: `azurerm_virtual_network`

The deprecated field `address_space` will be removed.

Resource: `azurerm_api_management`

The deprecated field `sku` will be removed in favour of the `sku_name` property.

The deprecated field `disable_backend_ssl30` will be removed in favour of the `enable_backend_ssl30` property.

The deprecated field `disable_backend_tls10` will be removed in favour of the `enable_backend_tls10` property.

The deprecated field `disable_backend_tls11` will be removed in favour of the `enable_backend_tls11` property.

The deprecated field `disable_frontend_ssl30` will be removed in favour of the `enable_frontend_ssl30` property.

The deprecated field `disable_frontend_tls10` will be removed in favour of the `enable_frontend_tls10` property.

The deprecated field `disable_frontend_tls11` will be removed in favour of the `enable_frontend_tls11` property.

The deprecated field `disable_triple_des_ciphers` will be removed in favour of the `enable_triple_des_ciphers` property.

The deprecated field `disable_triple_des_ciphers` will be removed in favour of the `enable_triple_des_ciphers` property.

Resource: `azurerm_app_service_plan`

The fields in the `properties` block (`app_service_environment_id`, `reserved` and `per_site_scaling`) have been moved to the top level - as such the `properties` block will be removed.

Resource: `azurerm_application_gateway`

The deprecated `fqdn_list` field in the `backend_address_pool` block will be removed in favour of the `fqdns` field, which is available from v1.22 of the AzureRM Provider.

The deprecated `ip_address_list` field in the `backend_address_pool` block will be removed in favour of the `ip_addresses` field, which is available from v1.22 of the AzureRM Provider.

Resource: `azurerm_automation_schedule`

The deprecated `account_name` field will be removed. This has been deprecated in favour of the `automation_account_name` field.

Resource: `azurerm_azuread_application`

The AzureAD Data Sources and Resources have been moved to the new AzureAD Provider (<https://www.terraform.io/docs/providers/azuread/index.html>), as such this Resource will be removed in v2.0 of the AzureRM Provider.

A guide on how to migrate to using the new Provider can be found here (<https://www.terraform.io/docs/providers/azurerm/guides/migrating-to-azuread.html>).

Resource: `azurerm_azuread_service_principal`

The AzureAD Data Sources and Resources have been moved to the new AzureAD Provider (<https://www.terraform.io/docs/providers/azuread/index.html>), as such this Resource will be removed.

A guide on how to migrate to using the new Provider can be found here (<https://www.terraform.io/docs/providers/azurerm/guides/migrating-to-azuread.html>).

Resource: `azurerm_azuread_service_principal_password`

The AzureAD Data Sources and Resources have been moved to the new AzureAD Provider (<https://www.terraform.io/docs/providers/azuread/index.html>), as such this Resource will be removed.

A guide on how to migrate to using the new Provider can be found here (<https://www.terraform.io/docs/providers/azurerm/guides/migrating-to-azuread.html>).

Resource: `azurerm_connection_monitor`

The `azurerm_connection_monitor` resource will be deprecated in favour of a new resources `azurerm_network_connection_monitor`.

Resource: `azurerm_container_group`

The deprecated `port` and `protocol` fields in the `container` block will be removed. These fields have been moved into the `ports` block within the `ports` field.

The deprecated `command` field in the `container` block will be removed. This has been replaced by the `commands` field in the `container` block.

Resource: `azurerm_container_registry`

The deprecated `storage_account` block will be removed. This has been replaced by the `storage_account_id` field and is only applicable to `Classic` Container Registries.

Resource: `azurerm_container_service`

Azure Container Service (ACS) is being Deprecated in favour of Azure Kubernetes Service (AKS) (more information can be found here (<https://azure.microsoft.com/en-us/updates/azure-container-service-will-retire-on-january-31-2020/>)), in preparation the resource will be removed.

Resource: `azurerm_cosmosdb_account`

The deprecated `failover_policy` block will be removed. This has been replaced by the `geo_location` block.

Resource: `azurerm_ddos_protection_plan`

The `azurerm_ddos_protection_planr` resource will be deprecated in favour of a new resources `azurerm_network_ddos_protection_plan`.

Resource: `azurerm_dns_mx_record`

The `preference` field in the `record` block will change from a `String` to an `Integer` to better reflect the API.

Resource: `azurerm_dns_ns_record`

The deprecated `record` field will be removed. This has been replaced by the `records` field which accepts multiple values.

Resource: `azurerm_dns_zone`

The deprecated `zone_type` field will be removed. This has been replaced by the `azurerm_private_dns_zone` resource.

Resource: `azurerm_eventhub`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_eventhub_authorization_rule`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_eventhub_consumer_group`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_eventhub_namespace`

The deprecated `kafka_enabled` field will be removed, since this is no longer used.

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_firewall`

The deprecated `internal_public_ip_address_id` field in the `ip_configuration` block will be removed. This field has been replaced by the `public_ip_address_id` field in the `ip_configuration` block.

Resource: `azurerm_kubernetes_cluster`

The deprecated `dns_prefix` field in the `agent_pool_profile` block will be removed. This field has been removed by Azure and is no longer used.

The deprecated `fqdn` field in the `agent_pool_profile` block will be removed. This has been replaced by the top-level field `fqdn`.

The `service_principal` will be changing from a Set to a List, which will allow Terraform to better detect when the values have changed locally and as such can detect when this needs to be recreated.

Resource: `azurerm_lb_backend_address_pool`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_lb_nat_probe`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_lb_nat_rule`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_lb_probe`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_lb_rule`

The deprecated `location` field will be removed, since this is no longer used.

Resource: `azurerm_log_analytics_linked_service`

The `resource_id` field has been moved from the `linked_service_properties` block to the top-level.

The `linked_service_properties` block will be removed, since it's no longer required.

Resource: `azurerm_log_analytics_workspace_linked_service`

This resource has been renamed to `azurerm_log_analytics_linked_service` which is available from v1.22 of the AzureRM Provider - instructions on how to migrate are available in this guide (<https://terraform.io/docs/providers/azurerm/guides/migrating-between-renamed-resources.html>). As such this resource will be removed.

Resource: `azurerm_mssql_elasticpool`

The deprecated `elastic_pool_properties` block will be removed. The fields inside this block have been moved to the top-level.

Resource: `azurerm_network_interface`

The `application_gateway_backend_address_pools_ids` field in the `ip_configuration` block will be removed. This has been replaced by the `azurerm_network_interface_application_gateway_backend_address_pool_association` resource.

The `application_security_group_ids` field in the `ip_configuration` block will be removed. This has been replaced by the `azurerm_network_interface_application_security_group_association` resource.

The `load_balancer_backend_address_pools_ids` field in the `ip_configuration` block will be removed. This has been

replaced by the `azurerm_network_interface_backend_address_pool_association` resource.

The `load_balancer_inbound_nat_rules_ids` field in the `ip_configuration` block will be removed. This has been replaced by the `azurerm_network_interface_nat_rule_association` resource.

Resource: `azurerm_packet_capture`

The `azurerm_packet_capture` resource will be deprecated in favour of a new resources `azurerm_network_packet_capture`.

Resource: `azurerm_public_ip`

The deprecated `public_ip_address_allocation` field will be removed. This field has been replaced by `allocation_method`.

Resource: `azurerm_scheduler_job`

Azure Scheduler is being retired in favour of Logic Apps (more information can be found here (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such this Resource will be removed in v2.0 of the AzureRM Provider.

Resource: `azurerm_scheduler_job_collection`

Azure Scheduler is being retired in favour of Logic Apps (more information can be found here (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such this Resource will be removed in v2.0 of the AzureRM Provider.

Resource: `azurerm_servicebus_queue`

The deprecated `location` field will be removed, since this is no longer used.

The deprecated `enable_batched_operations` field will be removed, since this is no longer used.

The deprecated `support_ordering` field will be removed, since this is no longer used.

Resource: `azurerm_servicebus_subscription`

The deprecated `location` field will be removed, since this is no longer used.

The deprecated `dead_lettering_on_filter_evaluation_exceptions` field will be removed, since this is no longer used.

Resource: `azurerm_servicebus_topic`

The deprecated `location` field will be removed, since this is no longer used.

The deprecated `enable_filtering_messages_before_publishing` field will be removed, since this is no longer used.

Resource: `azurerm_storage_account`

The deprecated `account_type` field will be removed. This has been split into the fields `account_tier` and `account_replication_type`.

Resource: `azurerm_storage_blob`

The deprecated `attempts` field will be removed, since Terraform now handles retries automatically.

The deprecated `resource_group_name` field will be removed, since this is no longer used.

Resource: `azurerm_storage_container`

The deprecated `resource_group_name` field will be removed, since this is no longer used.

The deprecated `properties` block will be removed.

Resource: `azurerm_storage_queue`

The deprecated `resource_group_name` field will be removed, since this is no longer used.

Resource: `azurerm_storage_share`

The deprecated `resource_group_name` field will be removed, since this is no longer used.

Resource: `azurerm_storage_table`

The deprecated `resource_group_name` field will be removed, since this is no longer used.

Resource: `azurerm_subnet`

The deprecated field `network_security_group_id` will be removed. This has been replaced by the `azurerm_subnet_network_security_group_association` resource.

The deprecated field `route_table_id` will be removed. This has been replaced by the `azurerm_subnet_route_table_association` resource.

Resource: `azurerm_virtual_machine`

The `azurerm_virtual_machine` resource will be deprecated in favour of two new resources: `azurerm_linux_virtual_machine` and `azurerm_windows_virtual_machine`.

Splitting the Virtual Machine resource in two allows us to both provide finer-grain validation for this resource and update the schema.

The existing `azurerm_virtual_machine` Resource will continue to be available in its current form, however it will eventually be deprecated and we recommend using the new resources going forward.

Resource: `azurerm_virtual_machine_scale_set`

The `azurerm_virtual_machine_scale_set` resource will be deprecated in favour of two new resources: `azurerm_linux_virtual_machine_scale_set` and `azurerm_windows_virtual_machine_scale_set`.

Splitting the Virtual Machine Scale Set resource in two allows us to both provide finer-grain validation for this resource and update the schema.

The existing `azurerm_virtual_machine_scale_set` Resource will continue to be available in its current form, however it will eventually be deprecated and we recommend using the new resources going forward.

Resource: `azurerm_log_analytics_linked_service`

The deprecated field `linked_service_properties` will be removed. This has been replaced by the `resource_id` resource.

Resource: `azurerm_log_analytics_workspace_linked_service`

The deprecated field `linked_service_properties` will be removed. This has been replaced by the `resource_id` resource.

We've spent the past few months laying the groundwork for these changes - and whilst we appreciate that your Terraform Configurations may require code changes to upgrade to 2.0 - we take Semantic Versioning seriously and so try our best to limit these changes to major versions.

You can follow along with the work in the 2.0 release in this GitHub Milestone (<https://github.com/terraform-providers/terraform-provider-azurerm/milestone/31>) - and in this GitHub Issue (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/2807>).

Azure Provider: Migrating to a renamed resource

In v1.22 of the AzureRM Provider several resources have been deprecated in favour of a renamed version - this guide covers how to migrate from the old resource to the new one and is applicable for the following resources:

Old Name	New Name
azurerm_log_analytics_workspace_linked_service	azurerm_log_analytics_linked_service
azurerm_autoscale_setting	azurerm_monitor_autoscale_setting
azurerm_metric_alertrule	azurerm_monitor_metric_alert
azurerm_connection_monitor	azurerm_network_connection_monitor
azurerm_ddos_protection_plan	azurerm_network_ddos_protection_plan
azurerm_packet_capture	azurerm_network_packet_capture

As the Schema's for each resource are the same at this time - it's possible to migrate between the resources by updating your Terraform Configuration and updating the Statefile.

In this guide we'll assume we're migrating from the `azurerm_autoscale_setting` resource to the new `azurerm_monitor_autoscale_setting` resource, but this should be applicable for any of the resources listed above.

Assuming we have the following Terraform Configuration:

```
resource "azurerm_resource_group" "test" {
  # ...
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_autoscale_setting" "test" {
  name                = "myAutoscaleSetting"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"

  # ...
}
```

We can update the Terraform Configuration to use the new resource by updating the name from `azurerm_autoscale_setting` to `azurerm_monitor_autoscale_setting`:

```

resource "azurerm_resource_group" "test" {
  # ...
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_monitor_autoscale_setting" "test" {
  name                = "myAutoscaleSetting"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"

  # ...
}

```

As the Terraform Configuration has been updated - we now need to update the State. We can view the items Terraform is tracking in it's Statefile using the `terraform state list` command, for example:

```

$ terraform state list
azurerm_autoscale_setting.test
azurerm_resource_group.test
azurerm_virtual_machine.test

```

We can move each of the resources which has been renamed in the state using the `terraform state mv` command - for example:

```

$ terraform state mv azurerm_autoscale_setting.test azurerm_monitor_autoscale_setting.test
Moved azurerm_autoscale_setting.test to azurerm_monitor_autoscale_setting.test

```

Once this has been done, running `terraform plan` should show no changes:

```

$ terraform plan
Refreshing Terraform state in-memory prior to plan...
The refreshed state will be used to calculate this plan, but will not be
persisted to local or remote state storage.

-----

No changes. Infrastructure is up-to-date.

This means that Terraform did not detect any differences between your
configuration and real physical resources that exist. As a result, no
actions need to be performed.

```

At this point you've switched over to using the newly renamed resources and should be able to continue using Terraform as normal.

Azure Active Directory: Migrating to the AzureAD Provider

In v1.21 of the AzureRM Provider the Azure Active Directory Data Sources and Resources have been split out into a new Provider specifically for Azure Active Directory.

This guide covers how to migrate from using the following Data Sources and Resources in the AzureRM Provider to using them in the new AzureAD Provider:

- Data Source: `azurerm_azuread_application`
- Data Source: `azurerm_azuread_service_principal`
- Resource: `azurerm_azuread_application`
- Resource: `azurerm_azuread_service_principal`
- Resource: `azurerm_azuread_service_principal_password`

Updating the Provider block

As the AzureAD and AzureRM Provider support the same authentication methods - it's possible to update the Provider block by setting the new Provider name and version, for example:

```
provider "azurerm" {  
  version = "=1.36.0"  
}
```

can become:

```
provider "azuread" {  
  version = "=0.1.0"  
}
```

Updating the Terraform Configurations

The Azure Active Directory Data Sources and Resources have been split out into the new Provider - which means the name of the Data Sources and Resources has changed slightly.

The main difference in naming is that the `azurerm_` prefix has been removed from the names of the Data Sources and Resources - the following table explains the new name for each of the Azure Active Directory resources:

Type	Old Name	New Name
Data Source	<code>azurerm_azuread_application</code>	<code>azuread_application</code>
Data Source	<code>azurerm_azuread_service_principal</code>	<code>azuread_service_principal</code>
Resource	<code>azurerm_azuread_application</code>	<code>azuread_application</code>
Resource	<code>azurerm_azuread_service_principal</code>	<code>azuread_service_principal</code>

Type	Old Name	New Name
Resource	azurerm_azuread_service_principal_password	azuread_service_principal_password

Once the Provider blocks have been updated, it should be possible to replace the `azurerm_` prefix in your Terraform Configuration from each of the AzureAD resources (and any interpolations) so that the new resources in the AzureAD Provider are used instead.

For example the following Terraform Configuration:

```
resource "azurerm_azuread_application" "test" {
  name = "my-application"
}

resource "azurerm_azuread_service_principal" "test" {
  application_id = "${azurerm_azuread_application.test.application_id}"
}

resource "azurerm_azuread_service_principal_password" "test" {
  service_principal_id = "${azurerm_azuread_service_principal.test.id}"
  value                = "bd018069-622d-4b46-bcb9-2bbee49fe7d9"
  end_date             = "2020-01-01T01:02:03Z"
}
```

we can remove the `azurerm_` prefix from each of the resource names and interpolations to use the `AzureAD` provider instead by making this:

```
resource "azuread_application" "test" {
  name = "my-application"
}

resource "azuread_service_principal" "test" {
  application_id = "${azuread_application.test.application_id}"
}

resource "azuread_service_principal_password" "test" {
  service_principal_id = "${azuread_service_principal.test.id}"
  value                = "bd018069-622d-4b46-bcb9-2bbee49fe7d9"
  end_date             = "2020-01-01T01:02:03Z"
}
```

At this point it should be possible to run `terraform init`, which will download the new AzureAD Provider.

Migrating Resources in the State

Now that we've updated the Provider Block and the Terraform Configuration we need to update the names of the resources in the state.

Firstly, let's list the existing items in the state - we can do this by running `terraform state list`, for example:


```
$ terraform state list
azurerm_azuread_application.test
azurerm_azuread_service_principal.test
azurerm_azuread_service_principal_password.import
azurerm_azuread_service_principal_password.test
```

As the Terraform Configuration has been updated - we can move each of the resources in the state using the `terraform state mv` command, for example:

```
$ terraform state mv azurerm_azuread_application.test azuread_application.test
Moved azurerm_azuread_application.test to azuread_application.test
```

This needs to be repeated for each of the Azure Active Directory resources which exist in the state.

Once this has been done, running `terraform plan` should show no changes:

```
$ terraform plan
Refreshing Terraform state in-memory prior to plan...
The refreshed state will be used to calculate this plan, but will not be
persisted to local or remote state storage.

-----

No changes. Infrastructure is up-to-date.

This means that Terraform did not detect any differences between your
configuration and real physical resources that exist. As a result, no
actions need to be performed.
```

At this point you've switched over to using the new Azure Active Directory provider (<http://terraform.io/docs/providers/azuread/index.html>)! You can stay up to date with Releases (and file Feature Requests/Bugs) on the Github repository (<https://github.com/terraform-providers/terraform-provider-azuread>).

azurerm_analysis_services_server

Manages an Analysis Services Server.

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "analysis-services-server-test"
  location  = "northeurope"
}

resource "azurerm_analysis_services_server" "server" {
  name                  = "analysis-services-server"
  location              = "northeurope"
  resource_group_name   = "${azurerm_resource_group.rg.name}"
  sku                  = "S0"
  admin_users           = ["myuser@domain.tld"]
  enable_power_bi_service = true

  ipv4_firewall_rule {
    name      = "myRule1"
    range_start = "210.117.252.0"
    range_end   = "210.117.252.255"
  }

  tags = {
    abc = 123
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Analysis Services Server. Changing this forces a new resource to be created.
- `location` - (Required) The Azure location where the Analysis Services Server exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Analysis Services Server should be exist. Changing this forces a new resource to be created.
- `sku` - (Required) SKU for the Analysis Services Server. Possible values are: D1 , B1 , B2 , S0 , S1 , S2 , S4 , S8 and S9
- `admin_users` - (Optional) List of email addresses of admin users.
- `querypool_connection_mode` - (Optional) Controls how the read-write server is used in the query pool. If this values is set to `All` then read-write servers are also used for queries. Otherwise with `ReadOnly` these servers do not participate in query operations.

- `backup_blob_container_uri` - (Optional) URI and SAS token for a blob container to store backups.
 - `enable_power_bi_service` - (Optional) Indicates if the Power BI service is allowed to access or not.
 - `ipv4_firewall_rule` - (Optional) One or more `ipv4_firewall_rule` block(s) as defined below.
-

A `ipv4_firewall_rule` block supports the following:

- `name` - (Required) Specifies the name of the firewall rule.
- `range_start` - (Required) Start of the firewall rule range as IPv4 address.
- `range_end` - (Required) End of the firewall rule range as IPv4 address.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the Analysis Services Server.
- `server_full_name` - The full name of the Analysis Services Server.

Import

Analysis Services Server can be imported using the `resource id`, e.g.

```
terraform import azurerm_analysis_services_server.server /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resourcegroup1/providers/Microsoft.AnalysisServices/servers/server1
```

azurerm_api_management_api

Manages an API within an API Management Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name       = "My Company"
  publisher_email      = "company@terraform.io"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_api" "test" {
  name                = "example-api"
  resource_group_name = "${azurerm_resource_group.test.name}"
  api_management_name = "${azurerm_api_management.test.name}"
  revision            = "1"
  display_name        = "Example API"
  path                = "example"
  protocols            = ["https"]

  import {
    content_format = "swagger-link-json"
    content_value  = "http://conferenceapi.azurewebsites.net/?format=json"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management API. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The Name of the API Management Service where this API should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The Name of the Resource Group where the API Management API exists. Changing this forces a new resource to be created.
- `revision` - (Required) The Revision which used for this API.
- `display_name` - (Required) The display name of the API.

- `path` - (Required) The Path for this API Management API, which is a relative URL which uniquely identifies this API and all of its resource paths within the API Management Service.
 - `protocols` - (Required) A list of protocols the operations in this API can be invoked. Possible values are `http` and `https`.
-

- `description` - (Optional) A description of the API Management API, which may include HTML formatting tags.
 - `import` - (Optional) An `import` block as documented below.
 - `service_url` - (Optional) Absolute URL of the backend service implementing this API.
 - `soap_pass_through` - (Optional) Should this API expose a SOAP frontend, rather than a HTTP frontend? Defaults to `false`.
 - `subscription_key_parameter_names` - (Optional) A `subscription_key_parameter_names` block as documented below.
-

A `import` block supports the following:

- `content_format` - (Required) The format of the content from which the API Definition should be imported. Possible values are: `swagger-json`, `swagger-link-json`, `wadl-link-json`, `wadl-xml`, `wsdl` and `wsdl-link`.
 - `content_value` - (Required) The Content from which the API Definition should be imported. When a `content_format` of `*-link-*` is specified this must be a URL, otherwise this must be defined inline.
 - `wsdl_selector` - (Optional) A `wsdl_selector` block as defined below, which allows you to limit the import of a WSDL to only a subset of the document. This can only be specified when `content_format` is `wsdl` or `wsdl-link`.
-

A `subscription_key_parameter_names` block supports the following:

- `header` - (Required) The name of the HTTP Header which should be used for the Subscription Key.
 - `query` - (Required) The name of the QueryString parameter which should be used for the Subscription Key.
-

A `wsdl_selector` block supports the following:

- `service_name` - (Required) The name of service to import from WSDL.
 - `endpoint_name` - (Required) The name of endpoint (port) to import from WSDL.
-

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API.
- `is_current` - Is this the current API Revision?
- `is_online` - Is this API Revision online/accessible via the Gateway?

- `version` - The Version number of this API, if this API is versioned.
- `version_set_id` - The ID of the Version Set which this API is associated with.

Import

API Management API's can be imported using the `resource_id`, e.g.

```
terraform import azurerm_api_management_api.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/apis/api1
```

azurerm_api_management_api_operation

Manages an API Operation within an API Management Service.

Example Usage

```
data "azurerm_api_management_api" "example" {
  name                = "search-api"
  api_management_name = "search-api-management"
  resource_group_name = "search-service"
  revision            = "2"
}

resource "azurerm_api_management_api_operation" "example" {
  operation_id      = "user-delete"
  api_name          = "${data.azurerm_api_management_api.example.name}"
  api_management_name = "${data.azurerm_api_management_api.example.api_management_name}"
  resource_group_name = "${data.azurerm_api_management_api.example.resource_group_name}"
  display_name      = "Delete User Operation"
  method            = "DELETE"
  url_template      = "/users/{id}/delete"
  description        = "This can only be done by the logged in user."

  response {
    status_code = 200
  }
}
```

Argument Reference

The following arguments are supported:

- `operation_id` - (Required) A unique identifier for this API Operation. Changing this forces a new resource to be created.
- `api_name` - (Required) The name of the API within the API Management Service where this API Operation should be created. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The Name of the API Management Service where the API exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `display_name` - (Required) The Display Name for this API Management Operation.
- `method` - (Required) The HTTP Method used for this API Management Operation, like `GET` , `DELETE` , `PUT` or `POST` - but not limited to these values.

- `url_template` - (Required) The relative URL Template identifying the target resource for this operation, which may include parameters.
-

- `description` - (Optional) A description for this API Operation, which may include HTML formatting tags.
 - `request` - (Optional) A `request` block as defined below.
 - `response` - (Optional) One or more `response` blocks as defined below.
 - `template_parameter` - (Optional) One or more `template_parameter` blocks as defined below.
-

A `form_parameter` block supports the following:

- `name` - (Required) The Name of this Form Parameter.
 - `required` - (Required) Is this Form Parameter Required?
 - `type` - (Required) The Type of this Form Parameter, such as a `string`.
 - `description` - (Optional) A description of this Form Parameter.
 - `default_value` - (Optional) The default value for this Form Parameter.
 - `values` - (Optional) One or more acceptable values for this Form Parameter.
-

A `header` block supports the following:

- `name` - (Required) The Name of this Header.
 - `required` - (Required) Is this Header Required?
 - `type` - (Required) The Type of this Header, such as a `string`.
 - `description` - (Optional) A description of this Header.
 - `default_value` - (Optional) The default value for this Header.
 - `values` - (Optional) One or more acceptable values for this Header.
-

A `query_parameter` block supports the following:

- `name` - (Required) The Name of this Query Parameter.
 - `required` - (Required) Is this Query Parameter Required?
 - `type` - (Required) The Type of this Query Parameter, such as a `string`.
 - `description` - (Optional) A description of this Query Parameter.
 - `default_value` - (Optional) The default value for this Query Parameter.
 - `values` - (Optional) One or more acceptable values for this Query Parameter.
-

A `request` block supports the following:

- `description` - (Required) A description of the HTTP Request, which may include HTML tags.
 - `header` - (Optional) One or more `header` blocks as defined above.
 - `query_parameter` - (Optional) One or more `query_parameter` blocks as defined above.
 - `representation` - (Optional) One or more `representation` blocks as defined below.
-

A `representation` block supports the following:

- `content_type` - (Required) The Content Type of this representation, such as `application/json`.
- `form_parameter` - (Optional) One or more `form_parameter` block as defined above.

NOTE: This is Required when `content_type` is set to `application/x-www-form-urlencoded` or `multipart/form-data`.

- `sample` - (Optional) An example of this representation.
- `schema_id` - (Optional) The ID of an API Management Schema which represents this Response.

NOTE: This can only be specified when `content_type` is not set to `application/x-www-form-urlencoded` or `multipart/form-data`.

- `type_name` - (Optional) The Type Name defined by the Schema.

NOTE: This can only be specified when `content_type` is not set to `application/x-www-form-urlencoded` or `multipart/form-data`.

A `response` block supports the following:

- `status_code` - (Required) The HTTP Status Code.
 - `description` - (Required) A description of the HTTP Response, which may include HTML tags.
 - `header` - (Optional) One or more `header` blocks as defined above.
 - `representation` - (Optional) One or more `representation` blocks as defined below.
-

A `template_parameter` block supports the following:

- `name` - (Required) The Name of this Template Parameter.
- `required` - (Required) Is this Template Parameter Required?
- `type` - (Required) The Type of this Template Parameter, such as a `string`.
- `description` - (Optional) A description of this Template Parameter.
- `default_value` - (Optional) The default value for this Template Parameter.
- `values` - (Optional) One or more acceptable values for this Template Parameter.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API Operation.

Import

API Management API Operation's can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_api_operation.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/apis/api1/operations/operation1
```

azurerm_api_management_api_operation_policy

Manages an API Management API Operation Policy

Example Usage

```
resource "azurerm_api_management_api_operation" "example" {
  #...
}

resource "azurerm_api_management_api_operation_policy" "example" {
  api_name           = "${azurerm_api_management_api_operation.example.api_name}"
  api_management_name = "${azurerm_api_management_api_operation.example.api_management_name}"
  resource_group_name = "${azurerm_api_management_api_operation.example.resource_group_name}"
  operation_id       = "${azurerm_api_management_api_operation.example.operation_id}"

  xml_content = <<XML
<policies>
  <inbound>
    <find-and-replace from="xyz" to="abc" />
  </inbound>
</policies>
XML
}
```

Argument Reference

The following arguments are supported:

- `api_name` - (Required) The ID of the API Management API Operation within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `xml_content` - (Optional) The XML Content for this Policy.
- `xml_link` - (Optional) A link to a Policy XML Document, which must be publicly available.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API Operation Policy.

Import

API Management API Operation Policy can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_api_operation_policy.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/instance1/apis/api1/operations/operation1/policies/policy
```

azurerm_api_management_api_policy

Manages an API Management API Policy

Example Usage

```
data "azurerm_api_management_api" "example" {
  api_name            = "my-api"
  api_management_name = "example-apim"
  resource_group_name = "search-service"
}

resource "azurerm_api_management_api_policy" "example" {
  api_name            = "${data.azurerm_api_management_api.example.name}"
  api_management_name = "${data.azurerm_api_management_api.example.api_management_name}"
  resource_group_name = "${data.azurerm_api_management_api.example.resource_group_name}"

  xml_content = <<XML
<policies>
  <inbound>
    <find-and-replace from="xyz" to="abc" />
  </inbound>
</policies>
XML
}
```

Argument Reference

The following arguments are supported:

- `api_name` - (Required) The ID of the API Management API within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `xml_content` - (Optional) The XML Content for this Policy as a string. An XML file can be used here with Terraform's file function (<https://www.terraform.io/docs/configuration/functions/file.html>) that is similar to Microsoft's `PolicyFilePath` option.
- `xml_link` - (Optional) A link to a Policy XML Document, which must be publicly available.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API Policy.

Import

API Management API Policy can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_api_policy.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/apis/exampleId/policies/policy
```

azurerm_api_management_api_schema

Manages an API Schema within an API Management Service.

Example Usage

```
data "azurerm_api_management_api" "example" {
  name                = "search-api"
  api_management_name = "search-api-management"
  resource_group_name = "search-service"
  revision            = "2"
}

resource "azurerm_api_management_api_schema" "example" {
  api_name                = "${data.azurerm_api_management_api.example.name}"
  api_management_name     = "${data.azurerm_api_management_api.example.api_management_name}"
  resource_group_name     = "${data.azurerm_api_management_api.example.resource_group_name}"
  schema_id              = "example-sche,a"
  content_type            = "application/vnd.ms-azure-apim.xsd+xml"
  value                   = "${file("api_management_api_schema.xml")}"
}
```

Argument Reference

The following arguments are supported:

- `schema_id` - (Required) A unique identifier for this API Schema. Changing this forces a new resource to be created.
- `api_name` - (Required) The name of the API within the API Management Service where this API Schema should be created. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The Name of the API Management Service where the API exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The Name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `content_type` - (Required) The content type of the API Schema.
- `value` - (Required) The JSON escaped string defining the document representing the Schema.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API Schema.

Import

API Management API Schema's can be imported using the `resource_id`, e.g.

```
terraform import azurerm_api_management_api_schema.test /subscriptions/00000000-0000-0000-0000-000000000000/00/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/schemas/schema1
```


azurerm_api_management_version_set

Manages an API Version Set within a API Management Service.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example-resources"
  location = "West US"
}

resource "azurerm_api_management" "example" {
  name            = "example-apim"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  publisher_name  = "pub1"
  publisher_email  = "pub1@email.com"

  sku {
    name     = "Developer"
    capacity = 1
  }
}

resource "azurerm_api_management_version_set" "example" {
  name                = "example-apimapi-1.0.0"
  resource_group_name = "${azurerm_resource_group.example.name}"
  api_management_name = "${azurerm_api_management.example.name}"
  display_name        = "ExampleAPIVersionSet"
  versioning_scheme   = "Segment"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Version Set. Changing this forces a new resource to be created.
 - `api_management_name` - (Required) The name of the API Management Service ([/docs/providers/azurerm/r/api_management.html](/docs/providers/azurerm/r/api_management.html)) in which the API Version Set should exist. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the Resource Group in which the parent API Management Service exists. Changing this forces a new resource to be created.
 - `display_name` - (Required) The display name of this API Version Set.
 - `versioning_scheme` - (Required) Specifies where in an Inbound HTTP Request that the API Version should be read from. Possible values are `Header`, `Query` and `Segment`.
-

- `description` - (Optional) The description of API Version Set.
- `version_header_name` - (Optional) The name of the Header which should be read from Inbound Requests which defines the API Version.

NOTE: This must be specified when `versioning_scheme` is set to `Header`.

- `version_query_name` - (Optional) The name of the Query String which should be read from Inbound Requests which defines the API Version.

NOTE: This must be specified when `versioning_scheme` is set to `Query`.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Version Set.

Import

API Version Set can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_version_set.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-resources/providers/Microsoft.ApiManagement/service/example-apim/api-version-sets/example-apimp
```

azurerm_api_management_authorization_server

Manages an Authorization Server within an API Management Service.

Example Usage

```
data "azurerm_api_management_api" "example" {
  name = "search-api"
  api_management_name = "search-api-management"
  resource_group_name = "search-service"
  revision = "2"
}

resource "azurerm_api_management_authorization_server" "example" {
  name = "test-server"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
  display_name = "Test Server"
  authorization_endpoint = "https://example.mydomain.com/client/authorize"
  client_id = "42424242-4242-4242-4242-424242424242"
  client_registration_endpoint = "https://example.mydomain.com/client/register"

  grant_types = [
    "authorizationCode",
  ]
}
```

Argument Reference

The following arguments are supported:

- `api_management_name` - (Required) The name of the API Management Service in which this Authorization Server should be created. Changing this forces a new resource to be created.
- `authorization_methods` - (Required) The HTTP Verbs supported by the Authorization Endpoint. Possible values are DELETE, GET, HEAD, OPTIONS, PATCH, POST, PUT and TRACE.

NOTE: GET must always be present.

- `authorization_endpoint` - (Required) The OAUTH Authorization Endpoint.
- `client_id` - (Required) The Client/App ID registered with this Authorization Server.
- `client_registration_endpoint` - (Required) The URI of page where Client/App Registration is performed for this Authorization Server.
- `display_name` - (Required) The user-friendly name of this Authorization Server.

- `grant_types` - (Required) Form of Authorization Grants required when requesting an Access Token. Possible values are `authorizationCode`, `clientCredentials`, `implicit` and `resourceOwnerPassword`.
- `name` - (Required) The name of this Authorization Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.

- `bearer_token_sending_methods` - (Optional) The mechanism by which Access Tokens are passed to the API. Possible values are `authorizationHeader` and `query`.
- `client_authentication_method` - (Optional) The Authentication Methods supported by the Token endpoint of this Authorization Server.. Possible values are `Basic` and `Body`.
- `client_secret` - (Optional) The Client/App Secret registered with this Authorization Server.
- `default_scope` - (Optional) The Default Scope used when requesting an Access Token, specified as a string containing space-delimited values.
- `description` - (Optional) A description of the Authorization Server, which may contain HTML formatting tags.
- `resource_owner_password` - (Optional) The password associated with the Resource Owner.

NOTE: This can only be specified when `grant_type` includes `resourceOwnerPassword`.

- `resource_owner_username` - (Optional) The username associated with the Resource Owner.

NOTE: This can only be specified when `grant_type` includes `resourceOwnerPassword`.

- `support_state` - (Optional) Does this Authorization Server support State? If this is set to `true` the client may use the `state` parameter to raise protocol security.
- `token_body_parameters` - (Optional) A `token_body_parameters` block as defined below.
- `token_endpoint` - (Optional) The OAUTH Token Endpoint.

A `token_body_parameter` block supports the following:

- `name` - (Required) The Name of the Parameter.
- `value` - (Required) The Value of the Parameter.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Authorization Server.

Import

API Management Authorization Servers can be imported using the `resource_id`, e.g.

```
terraform import azurerm_api_management_authorization_server.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/authorizationServers/server1
```

azurerm_api_management_backend

Manages a backend within an API Management Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name      = "My Company"
  publisher_email     = "company@terraform.io"

  sku {
    name      = "Developer"
    capacity = 1
  }
}

resource "azurerm_api_management_backend" "test" {
  name                = "example-backend"
  resource_group_name = "${azurerm_resource_group.test.name}"
  api_management_name = "${azurerm_api_management.test.name}"
  protocol            = "http"
  url                 = "https://backend"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management backend. Changing this forces a new resource to be created.
 - `api_management_name` - (Required) The Name of the API Management Service where this backend should be created. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The Name of the Resource Group where the API Management Service exists. Changing this forces a new resource to be created.
 - `protocol` - (Required) The protocol used by the backend host. Possible values are `http` or `soap`.
 - `url` - (Required) The URL of the backend host.
-
- `credentials` - (Optional) A `credentials` block as documented below.

- `description` - (Optional) The description of the backend.
 - `proxy` - (Optional) A `proxy` block as documented below.
 - `resource_id` - (Optional) The management URI of the backend host in an external system. This URI can be the ARM Resource ID of Logic Apps, Function Apps or API Apps, or the management endpoint of a Service Fabric cluster.
 - `service_fabric_cluster` - (Optional) A `service_fabric_cluster` block as documented below.
 - `title` - (Optional) The title of the backend.
 - `tls` - (Optional) A `tls` block as documented below.
-

A `credentials` block supports the following:

- `authorization` - (Optional) An `authorization` block as defined below.
 - `certificate` - (Optional) A list of client certificate thumbprints to present to the backend host. The certificates must exist within the API Management Service.
 - `header` - (Optional) A mapping of header parameters to pass to the backend host. The keys are the header names and the values are a comma separated string of header values. This is converted to a list before being passed to the API.
 - `query` - (Optional) A mapping of query parameters to pass to the backend host. The keys are the query names and the values are a comma separated string of query values. This is converted to a list before being passed to the API.
-

An `authorization` block supports the following:

- `parameter` - (Optional) The authentication Parameter value.
 - `scheme` - (Optional) The authentication Scheme name.
-

A `proxy` block supports the following:

- `password` - (Optional) The password to connect to the proxy server.
 - `url` - (Optional) The URL of the proxy server.
 - `username` - (Optional) The username to connect to the proxy server.
-

A `service_fabric_cluster` block supports the following:

- `client_certificate_thumbprint` - (Required) The client certificate thumbprint for the management endpoint.
 - `management_endpoints` - (Required) A list of cluster management endpoints.
 - `max_partition_resolution_retries` - (Required) The maximum number of retries when attempting resolve the partition.
 - `server_certificate_thumbprints` - (Optional) A list of thumbprints of the server certificates of the Service Fabric cluster.
 - `server_x509_name` - (Optional) One or more `server_x509_name` blocks as documented below.
-

A `server_x509_name` block supports the following:

- `issuer_certificate_thumbprint` - (Required) The thumbprint for the issuer of the certificate.
 - `name` - (Required) The common name of the certificate.
-

A `tls` block supports the following:

- `validate_certificate_chain` - (Optional) Flag indicating whether SSL certificate chain validation should be done when using self-signed certificates for the backend host.
 - `validate_certificate_name` - (Optional) Flag indicating whether SSL certificate name validation should be done when using self-signed certificates for the backend host.
-

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management API.

Import

API Management backends can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_backend.test /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/backends/backend1
```


azurerm_api_management_certificate

Manages an Certificate within an API Management Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name      = "My Company"
  publisher_email     = "company@terraform.io"

  sku {
    name      = "Developer"
    capacity = 1
  }
}

resource "azurerm_api_management_certificate" "test" {
  name                = "example-cert"
  api_management_name = "${azurerm_api_management.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  data                = "${filebase64("example.pfx")}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management Certificate. Changing this forces a new resource to be created.
 - `api_management_name` - (Required) The Name of the API Management Service where this Service should be created. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The Name of the Resource Group where the API Management Service exists. Changing this forces a new resource to be created.
 - `data` - (Required) The base-64 encoded certificate data, which must be a PFX file. Changing this forces a new resource to be created.
 - `password` - (Optional) The password used for this certificate. Changing this forces a new resource to be created.
-

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Certificate.
- `expiration` - The Expiration Date of this Certificate, formatted as an RFC3339 string.
- `subject` - The Subject of this Certificate.
- `thumbprint` - The Thumbprint of this Certificate.

Import

API Management Certificates can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_certificate.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/certificates/certificate1
```

azurerm_api_management_group

Manages an API Management Group.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_api_management" "example" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  publisher_name       = "pub1"
  publisher_email      = "pub1@email.com"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_group" "example" {
  name                = "example-apimg"
  resource_group_name = "${azurerm_resource_group.example.name}"
  api_management_name = "${azurerm_api_management.example.name}"
  display_name        = "Example Group"
  description          = "This is an example API management group."
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management Group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Group should exist. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service ([/docs/providers/azurerm/r/api_management.html](/docs/providers/azurerm/r/api_management.html)) in which the API Management Group should exist. Changing this forces a new resource to be created.
- `display_name` - (Required) The display name of this API Management Group.
- `description` - (Optional) The description of this API Management Group.
- `external_id` - (Optional) The identifier of the external Group. For example, an Azure Active Directory group `aad://<tenant>.onmicrosoft.com/groups/<group object id>`.
- `type` - (Optional) The type of this API Management Group. Possible values are `custom` and `external`. Default is `custom`.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Group.

Import

API Management Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_group.example /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/example-resources/providers/Microsoft.ApiManagement/service/example-apim/groups/example-a  
pimg
```

azurerm_api_management_group_user

Manages an API Management User Assignment to a Group.

Example Usage

```
data "azurerm_api_management_user" "example" {
  user_id            = "my-user"
  api_management_name = "example-apim"
  resource_group_name = "search-service"
}

resource "azurerm_api_management_group_user" "example" {
  user_id            = "${data.azurerm_api_management_user.example.id}"
  group_name         = "example-group"
  resource_group_name = "${data.azurerm_api_management_user.example.resource_group_name}"
  api_management_name = "${data.azurerm_api_management_user.example.api_management_name}"
}
```

Argument Reference

The following arguments are supported:

- `user_id` - (Required) The ID of the API Management User which should be assigned to this API Management Group. Changing this forces a new resource to be created.
- `group_name` - (Required) The Name of the API Management Group within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Group User.

Import

API Management Group Users can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_group_user.example /subscriptions/00000000-0000-0000-0000-00000000  
00000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/groups/groupId/users/user1  
23
```

azurerm_api_management

Manages an API Management Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name       = "My Company"
  publisher_email      = "company@terraform.io"

  sku_name = "Developer_1"

  policy {
    xml_content = <<XML
    <policies>
      <inbound />
      <backend />
      <outbound />
      <on-error />
    </policies>
XML
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `location` - (Required) The Azure location where the API Management Service exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service should be exist. Changing this forces a new resource to be created.
- `publisher_name` - (Required) The name of publisher/company.
- `publisher_email` - (Required) The email of publisher/company.
- `sku` - (Deprecated) A `sku` block as documented below

- `sku_name` - (Required) `sku_name` is a string consisting of two parts separated by an underscore(_). The first part is the name, valid values include: `Developer`, `Basic`, `Standard` and `Premium`. The second part is the `capacity` (e.g. the number of deployed units of the `sku`), which must be a positive integer (e.g. `Developer_1`).
-

- `additional_location` - (Optional) One or more `additional_location` blocks as defined below.
 - `certificate` - (Optional) One or more (up to 10) `certificate` blocks as defined below.
 - `identity` - (Optional) An `identity` block is documented below.
 - `hostname_configuration` - (Optional) A `hostname_configuration` block as defined below.
 - `notification_sender_email` - (Optional) Email address from which the notification will be sent.
 - `policy` - (Optional) A `policy` block as defined below.
 - `security` - (Optional) A `security` block as defined below.
 - `sign_in` - (Optional) A `sign_in` block as defined below.
 - `sign_up` - (Optional) A `sign_up` block as defined below.
 - `tags` - (Optional) A mapping of tags assigned to the resource.
-

A `additional_location` block supports the following:

- `location` - (Required) The name of the Azure Region in which the API Management Service should be expanded to.
-

A `certificate` block supports the following:

- `encoded_certificate` - (Required) The Base64 Encoded PFX Certificate.
 - `certificate_password` - (Required) The password for the certificate.
 - `store_name` - (Required) The name of the Certificate Store where this certificate should be stored. Possible values are `CertificateAuthority` and `Root`.
-

A `hostname_configuration` block supports the following:

- `management` - (Optional) One or more `management` blocks as documented below.
 - `portal` - (Optional) One or more `portal` blocks as documented below.
 - `proxy` - (Optional) One or more `proxy` blocks as documented below.
 - `scm` - (Optional) One or more `scm` blocks as documented below.
-

A `identity` block supports the following:

- `type` - (Required) Specifies the type of Managed Service Identity that should be configured on this API Management Service. At this time the only supported value is `SystemAssigned`.
-

A `management`, `portal` and `scm` block supports the following:

- `host_name` - (Required) The Hostname to use for the Management API.
- `key_vault_id` - (Optional) The ID of the Key Vault Secret containing the SSL Certificate, which must be should be of the type `application/x-pkcs12`.

NOTE: Setting this field requires the `identity` block to be specified, since this identity is used for to retrieve the Key Vault Certificate. Auto-updating the Certificate from the Key Vault requires the Secret version isn't specified.

- `certificate` - (Optional) The Base64 Encoded Certificate.
- `certificate_password` - (Optional) The password associated with the certificate provided above.

NOTE: Either `key_vault_id` or `certificate` and `certificate_password` must be specified.

- `negotiate_client_certificate` - (Optional) Should Client Certificate Negotiation be enabled for this Hostname? Defaults to `false`.

A `policy` block supports the following:

- `xml_content` - (Optional) The XML Content for this Policy.
- `xml_link` - (Optional) A link to an API Management Policy XML Document, which must be publicly available.

A `proxy` block supports the following:

- `default_ssl_binding` - (Optional) Is the certificate associated with this Hostname the Default SSL Certificate? This is used when an SNI header isn't specified by a client. Defaults to `false`.
- `host_name` - (Required) The Hostname to use for the Management API.
- `key_vault_id` - (Optional) The ID of the Key Vault Secret containing the SSL Certificate, which must be should be of the type `application/x-pkcs12`.

NOTE: Setting this field requires the `identity` block to be specified, since this identity is used for to retrieve the Key Vault Certificate. Auto-updating the Certificate from the Key Vault requires the Secret version isn't specified.

- `certificate` - (Optional) The Base64 Encoded Certificate.
- `certificate_password` - (Optional) The password associated with the certificate provided above.

NOTE: Either `key_vault_id` or `certificate` and `certificate_password` must be specified.

- `negotiate_client_certificate` - (Optional) Should Client Certificate Negotiation be enabled for this Hostname? Defaults to `false`.

A `security` block supports the following:

- `enable_backend_ssl30` - (Optional) Should SSL 3.0 be enabled on the backend of the gateway? Defaults to `false`.

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Backend.Protocols.Ssl30` field

- `enable_backend_tls10` - (Optional) Should TLS 1.0 be enabled on the backend of the gateway? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Backend.Protocols.Tls10` field

- `enable_backend_tls11` - (Optional) Should TLS 1.1 be enabled on the backend of the gateway? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Backend.Protocols.Tls11` field

- `enable_frontend_ssl30` - (Optional) Should SSL 3.0 be enabled on the frontend of the gateway? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Protocols.Ssl30` field

- `enable_frontend_tls10` - (Optional) Should TLS 1.0 be enabled on the frontend of the gateway? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Protocols.Tls10` field

- `enable_frontend_tls11` - (Optional) Should TLS 1.1 be enabled on the frontend of the gateway? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Protocols.Tls11` field

- `enable_triple_des_ciphers` - (Optional) Should the `TLS_RSA_WITH_3DES_EDE_CBC_SHA` cipher be enabled for all TLS versions (1.0, 1.1 and 1.2)? Defaults to `false` .

info: This maps to the `Microsoft.WindowsAzure.ApiManagement.Gateway.Security.Ciphers.TripleDes168` field

- `disable_backend_ssl30` - (Optional) Should SSL 3.0 be disabled on the backend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false` .

Note: This property has been deprecated in favour of the `enable_backend_ssl30` property and will be removed in version 2.0 of the provider.

- `disable_backend_tls10` - (Optional) Should TLS 1.0 be disabled on the backend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false` .

Note: This property has been deprecated in favour of the `enable_backend_tls10` property and will be removed in version 2.0 of the provider.

- `disable_backend_tls11` - (Optional) Should TLS 1.1 be disabled on the backend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false` .

Note: This property has been deprecated in favour of the `enable_backend_tls11` property and will be removed in version 2.0 of the provider.

- `disable_frontend_ssl30` - (Optional) Should SSL 3.0 be disabled on the frontend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false`.

Note: This property has been deprecated in favour of the `enable_frontend_ssl30` property and will be removed in version 2.0 of the provider.

- `disable_frontend_tls10` - (Optional) Should TLS 1.0 be disabled on the frontend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false`.

Note: This property has been deprecated in favour of the `enable_frontend_tls10` property and will be removed in version 2.0 of the provider.

- `disable_frontend_tls11` - (Optional) Should TLS 1.1 be disabled on the frontend of the gateway? This property was mistakenly inverted and `true` actually enables it. Defaults to `false`.

Note: This property has been deprecated in favour of the `enable_frontend_tls11` property and will be removed in version 2.0 of the provider.

- `disable_triple_des_ciphers` - (Optional) Should the `TLS_RSA_WITH_3DES_EDE_CBC_SHA` cipher be disabled for all TLS versions (1.0, 1.1 and 1.2)? This property was mistakenly inverted and `true` actually enables it. Defaults to `false`.

Note: This property has been deprecated in favour of the `enable_triple_des_ciphers` property and will be removed in version 2.0 of the provider.

A `sku` block supports the following: (Deprecated)

- `name` - (Required) Specifies the Pricing Tier for the API Management Service. Possible values include: Developer, Basic, Standard and Premium.
- `capacity` - (Required) Specifies the Pricing Capacity for the API Management Service.

Note: This property has been deprecated in favour of the `sku_name` property and will be removed in version 2.0 of the provider.

A `sign_in` block supports the following:

- `enabled` - (Required) Should anonymous users be redirected to the sign in page?

A `sign_up` block supports the following:

- `enabled` - (Required) Can users sign up on the development portal?

- `terms_of_service` - (Optional) A `terms_of_service` block as defined below.
-

A `terms_of_service` block supports the following:

- `consent_required` - (Required) Should the user be asked for consent during sign up?
- `enabled` - (Required) Should Terms of Service be displayed during sign up?.
- `text` - (Required) The Terms of Service which users are required to agree to in order to sign up.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Service.
 - `additional_location` - One or more `additional_location` blocks as documented below.
 - `gateway_url` - The URL of the Gateway for the API Management Service.
 - `gateway_regional_url` - The Region URL for the Gateway of the API Management Service.
 - `identity` - An `identity` block as defined below.
 - `management_api_url` - The URL for the Management API associated with this API Management service.
 - `portal_url` - The URL for the Publisher Portal associated with this API Management service.
 - `public_ip_addresses` - The Public IP addresses of the API Management Service.
 - `scm_url` - The URL for the SCM (Source Code Management) Endpoint associated with this API Management service.
-

An `additional_location` block exports the following:

- `gateway_regional_url` - The URL of the Regional Gateway for the API Management Service in the specified region.
 - `public_ip_addresses` - Public Static Load Balanced IP addresses of the API Management service in the additional location. Available only for Basic, Standard and Premium SKU.
-

An `identity` block exports the following:

- `principal_id` - The Principal ID associated with this Managed Service Identity.
- `tenant_id` - The Tenant ID associated with this Managed Service Identity.

Import

API Management Services can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management.test /subscriptions/00000000-0000-0000-0000-000000000000/resource  
Groups/mygroup1/providers/Microsoft.ApiManagement/service/instance1
```

azurerm_api_management_logger

Manages a Logger within an API Management Service.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_application_insights" "example" {
  name                = "example-appinsights"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  application_type    = "Other"
}

resource "azurerm_api_management" "example" {
  name                = "example-apim"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  publisher_name      = "My Company"
  publisher_email     = "company@terraform.io"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_logger" "example" {
  name                = "example-logger"
  api_management_name = "${azurerm_api_management.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  application_insights {
    instrumentation_key = "${azurerm_application_insights.example.instrumentation_key}"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of this Logger, which must be unique within the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.

- `application_insights` - (Optional) An `application_insights` block as documented below.
- `buffered` - (Optional) Specifies whether records should be buffered in the Logger prior to publishing. Defaults to `true`.
- `description` - (Optional) A description of this Logger.
- `eventhub` - (Optional) An `eventhub` block as documented below.

An `application_insights` block supports the following:

- `instrumentation_key` - (Required) The instrumentation key used to push data to Application Insights.

An `eventhub` block supports the following:

- `name` - (Required) The name of an EventHub.
- `connection_string` - (Required) The connection string of an EventHub Namespace.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Logger.

Import

API Management Loggers can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_api_management_logger.example /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/example-rg/Microsoft.ApiManagement/service/example-apim/loggers/example-logger
```

azurerm_api_management_openid_connect_provider

Manages an OpenID Connect Provider within a API Management Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name      = "My Company"
  publisher_email     = "company@terraform.io"

  sku {
    name      = "Developer"
    capacity = 1
  }
}

resource "azurerm_api_management_openid_connect_provider" "test" {
  name                = "example-provider"
  api_management_name = "${azurerm_api_management.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  client_id           = "00001111-2222-3333-4444-555566667777"
  display_name        = "Example Provider"
  metadata_endpoint   = "https://example.com/example"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) the Name of the OpenID Connect Provider which should be created within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service in which this OpenID Connect Provider should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the API Management Service exists. Changing this forces a new resource to be created.
- `client_id` - (Required) The Client ID used for the Client Application.
- `client_secret` - (Required) The Client Secret used for the Client Application.
- `display_name` - (Required) A user-friendly name for this OpenID Connect Provider.

- `metadata_endpoint` - (Required) The URI of the Metadata endpoint.
-

- `description` - (Optional) A description of this OpenID Connect Provider.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management OpenID Connect Provider.

Import

API Management OpenID Connect Providers can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_openid_connect_provider.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/openidConnectProviders/provider1
```

azurerm_api_management_product_api

Manages an API Management API Assignment to a Product.

Example Usage

```
data "azurerm_api_management" "example" {
  name                = "example-api"
  resource_group_name = "example-resources"
}

data "azurerm_api_management_api" "example" {
  name                = "search-api"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
  revision            = "2"
}

data "azurerm_api_management_product" "test" {
  product_id          = "my-product"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
}

resource "azurerm_api_management_product_api" "example" {
  api_name                = "${data.azurerm_api_management_api.example.name}"
  product_id              = "${data.azurerm_api_management_product.example.product_id}"
  api_management_name     = "${data.azurerm_api_management.example.name}"
  resource_group_name     = "${data.azurerm_api_management.example.resource_group_name}"
}
```

Argument Reference

The following arguments are supported:

- `api_name` - (Required) The Name of the API Management API within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `product_id` - (Required) The ID of the API Management Product within the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Product API.

Import

API Management Product API's can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_product_api.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/products/exampleId/apis/apiId
```

azurerm_api_management_product_group

Manages an API Management Product Assignment to a Group.

Example Usage

```
data "azurerm_api_management" "example" {
  name                = "example-api"
  resource_group_name = "example-resources"
}

data "azurerm_api_management_product" "example" {
  product_id          = "my-product"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
}

data "azurerm_api_management_group" "example" {
  name                = "my-group"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
}

resource "azurerm_api_management_product_group" "example" {
  product_id          = "${data.azurerm_api_management_product.example.product_id}"
  group_name          = "${data.azurerm_api_management_group.example.name}"
  api_management_name = "${data.azurerm_api_management.example.name}"
  resource_group_name = "${data.azurerm_api_management.example.resource_group_name}"
}
```

Argument Reference

The following arguments are supported:

- `product_id` - (Required) The ID of the API Management Product within the API Management Service. Changing this forces a new resource to be created.
- `group_name` - (Required) The Name of the API Management Group within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Product Group.

Import

API Management Product Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_product_group.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/products/exampleId/groups/groupId
```

azurerm_api_management_product

Manages an API Management Product.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name      = "My Company"
  publisher_email     = "company@terraform.io"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_product" "test" {
  product_id          = "test-product"
  api_management_name = "${azurerm_api_management.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  display_name        = "Test Product"
  subscription_required = true
  approval_required   = true
  published            = true
}
```

Argument Reference

The following arguments are supported:

- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `approval_required` - (Optional) Do subscribers need to be approved prior to being able to use the Product?

NOTE: `approval_required` can only be set when `subscription_required` is set to `true`.

- `display_name` - (Required) The Display Name for this API Management Product.
- `product_id` - (Required) The Identifier for this Product, which must be unique within the API Management Service. Changing this forces a new resource to be created.
- `published` - (Required) Is this Product Published?

- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service should be exist. Changing this forces a new resource to be created.
- `subscription_required` - (Required) Is a Subscription required to access API's included in this Product?

- `description` - (Optional) A description of this Product, which may include HTML formatting tags.
- `subscriptions_limit` - (Optional) The number of subscriptions a user can have to this Product at the same time.

NOTE: `subscriptions_limit` can only be set when `subscription_required` is set to `true`.

- `terms` - (Optional) The Terms and Conditions for this Product, which must be accepted by Developers before they can begin the Subscription process.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Product.

Import

API Management Products can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_product.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/products/myproduct
```

azurerm_api_management_product_policy

Manages an API Management Product Policy

Example Usage

```
data "azurerm_api_management_product" "example" {
  product_id            = "my-product"
  api_management_name   = "example-apim"
  resource_group_name   = "search-service"
}

resource "azurerm_api_management_product_policy" "example" {
  product_id            = "${data.azurerm_api_management_product.example.product_id}"
  api_management_name   = "${data.azurerm_api_management_product.example.api_management_name}"
  resource_group_name   = "${data.azurerm_api_management_product.example.resource_group_name}"

  xml_content = <<XML
<policies>
  <inbound>
    <find-and-replace from="xyz" to="abc" />
  </inbound>
</policies>
XML
}
```

Argument Reference

The following arguments are supported:

- `product_id` - (Required) The ID of the API Management Product within the API Management Service. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
- `xml_content` - (Optional) The XML Content for this Policy.
- `xml_link` - (Optional) A link to a Policy XML Document, which must be publicly available.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Product Policy.

Import

API Management Product Policy can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_product_policy.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ApiManagement/service/service1/products/exampleId/policies/policy
```

azurerm_api_management_property

Manages an API Management Property.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_api_management" "example" {
  name                = "example-apim"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  publisher_name      = "pub1"
  publisher_email     = "pub1@email.com"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_property" "example" {
  name                = "example-apimg"
  resource_group_name = "${azurerm_resource_group.example.name}"
  api_management_name = "${azurerm_api_management.example.name}"
  display_name        = "ExampleProperty"
  value               = "Example Value"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the API Management Property. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the API Management Property should exist. Changing this forces a new resource to be created.
- `api_management_name` - (Required) The name of the API Management Service ([/docs/providers/azurerm/r/api_management.html](/docs/providers/azurerm/r/api_management.html)) in which the API Management Property should exist. Changing this forces a new resource to be created.
- `display_name` - (Required) The display name of this API Management Property.
- `value` - (Required) The value of this API Management Property.
- `secret` - (Optional) Specifies whether the API Management Property is secret. Valid values are `true` or `false`. The default value is `false`.

NOTE: setting the field `secret` to `true` doesn't make this field sensitive in Terraform, instead it marks the value as secret and encrypts the value in Azure.

- `tags` - (Optional) A list of tags to be applied to the API Management Property.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Property.

Import

API Management Properties can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_property.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-resources/providers/Microsoft.ApiManagement/service/example-apim/properties/example-apimp
```

azurerm_api_management_subscription

Manages a Subscription within a API Management Service.

Example Usage

```
data "azurerm_api_management" "test" {
  name                = "example-apim"
  resource_group_name = "example-resources"
}

data "azurerm_api_management_product" "test" {
  product_id          = "00000000-0000-0000-0000-000000000000"
  api_management_name = "${data.azurerm_api_management.test.name}"
  resource_group_name = "${data.azurerm_api_management.test.resource_group_name}"
}

data "azurerm_api_management_user" "test" {
  user_id             = "11111111-1111-1111-1111-111111111111"
  api_management_name = "${data.azurerm_api_management.test.name}"
  resource_group_name = "${data.azurerm_api_management.test.resource_group_name}"
}

resource "azurerm_api_management_subscription" "test" {
  api_management_name = "${data.azurerm_api_management.test.name}"
  resource_group_name = "${data.azurerm_api_management.test.resource_group_name}"
  user_id             = "${data.azurerm_api_management_user.test.id}"
  product_id          = "${data.azurerm_api_management_product.test.id}"
  display_name        = "Parser API"
}
```

Argument Reference

The following arguments are supported:

- `api_management_name` - (Required) The name of the API Management Service where this Subscription should be created. Changing this forces a new resource to be created.
 - `display_name` - (Required) The display name of this Subscription.
 - `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
 - `product_id` - (Required) The ID of the Product which should be assigned to this Subscription. Changing this forces a new resource to be created.
 - `user_id` - (Required) The ID of the User which should be assigned to this Subscription. Changing this forces a new resource to be created.
-
- `state` - (Optional) The state of this Subscription. Possible values are `active`, `cancelled`, `expired`, `rejected`,

submitted and suspended. Defaults to submitted.

- `subscription_id` - (Optional) An Identifier which should be used as the ID of this Subscription. If not specified a new Subscription ID will be generated. Changing this forces a new resource to be created.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management Subscription.

Import

API Management Subscriptions can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_subscription.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-resources/providers/Microsoft.ApiManagement/service/example-apim/properties/example-apimp
```

azurerm_api_management_user

Manages an API Management User.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_api_management" "test" {
  name                = "example-apim"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  publisher_name       = "My Company"
  publisher_email      = "company@terraform.io"

  sku_name = "Developer_1"
}

resource "azurerm_api_management_user" "test" {
  user_id            = "5931a75ae4bbd512288c680b"
  api_management_name = "${azurerm_api_management.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  first_name         = "Example"
  last_name          = "User"
  email              = "tom+tfdev@hashicorp.com"
  state              = "active"
}
```

Argument Reference

The following arguments are supported:

- `api_management_name` - (Required) The name of the API Management Service in which the User should be created. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the Resource Group in which the API Management Service exists. Changing this forces a new resource to be created.
 - `email` - (Required) The email address associated with this user.
 - `first_name` - (Required) The first name for this user.
 - `last_name` - (Required) The last name for this user.
 - `user_id` - (Required) The Identifier for this User, which must be unique within the API Management Service. Changing this forces a new resource to be created.
-

- `confirmation` - (Optional) The kind of confirmation email which will be sent to this user. Possible values are `invite` and `signup`. Changing this forces a new resource to be created.
- `note` - (Optional) A note about this user.
- `password` - (Optional) The password associated with this user.
- `state` - (Optional) The state of this user. Possible values are `active`, `blocked` and `pending`.

NOTE: the State can be changed from Pending Active/Blocked but not from Active/Blocked Pending.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the API Management User.

Import

API Management Users can be imported using the `resource id`, e.g.

```
terraform import azurerm_api_management_user.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ApiManagement/service/instance1/users/abc123
```

azurerm_application_gateway

Manages an Application Gateway.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_virtual_network" "test" {
  name                = "example-network"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  address_space       = ["10.254.0.0/16"]
}

resource "azurerm_subnet" "frontend" {
  name                = "frontend"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix      = "10.254.0.0/24"
}

resource "azurerm_subnet" "backend" {
  name                = "backend"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix      = "10.254.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name                = "example-pip"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  allocation_method   = "Dynamic"
}

# since these variables are re-used - a locals block makes this more maintainable
locals {
  backend_address_pool_name = "${azurerm_virtual_network.test.name}-beap"
  frontend_port_name       = "${azurerm_virtual_network.test.name}-feport"
  frontend_ip_configuration_name = "${azurerm_virtual_network.test.name}-feip"
  http_setting_name        = "${azurerm_virtual_network.test.name}-be-htst"
  listener_name            = "${azurerm_virtual_network.test.name}-httplstn"
  request_routing_rule_name = "${azurerm_virtual_network.test.name}-rqrt"
  redirect_configuration_name = "${azurerm_virtual_network.test.name}-rdrcfg"
}

resource "azurerm_application_gateway" "network" {
  name                = "example-appgateway"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
```



```

sku {
  name      = "Standard_Small"
  tier       = "Standard"
  capacity = 2
}

gateway_ip_configuration {
  name      = "my-gateway-ip-configuration"
  subnet_id = "${azurerm_subnet.frontend.id}"
}

frontend_port {
  name = "${local.frontend_port_name}"
  port = 80
}

frontend_ip_configuration {
  name      = "${local.frontend_ip_configuration_name}"
  public_ip_address_id = "${azurerm_public_ip.test.id}"
}

backend_address_pool {
  name = "${local.backend_address_pool_name}"
}

backend_http_settings {
  name              = "${local.http_setting_name}"
  cookie_based_affinity = "Disabled"
  path              = "/path1/"
  port              = 80
  protocol           = "Http"
  request_timeout    = 1
}

http_listener {
  name                  = "${local.listener_name}"
  frontend_ip_configuration_name = "${local.frontend_ip_configuration_name}"
  frontend_port_name     = "${local.frontend_port_name}"
  protocol                = "Http"
}

request_routing_rule {
  name              = "${local.request_routing_rule_name}"
  rule_type         = "Basic"
  http_listener_name = "${local.listener_name}"
  backend_address_pool_name = "${local.backend_address_pool_name}"
  backend_http_settings_name = "${local.http_setting_name}"
}
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Application Gateway. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which the Application Gateway should exist. Changing this forces a new resource to be created.
- `location` - (Required) The Azure region where the Application Gateway should exist. Changing this forces a new resource to be created.
- `backend_address_pool` - (Required) One or more `backend_address_pool` blocks as defined below.
- `backend_http_settings` - (Required) One or more `backend_http_settings` blocks as defined below.
- `frontend_ip_configuration` - (Required) One or more `frontend_ip_configuration` blocks as defined below.
- `frontend_port` - (Required) One or more `frontend_port` blocks as defined below.
- `gateway_ip_configuration` - (Required) One or more `gateway_ip_configuration` blocks as defined below.
- `http_listener` - (Required) One or more `http_listener` blocks as defined below.
- `identity` - (Optional) A `identity` block.
- `request_routing_rule` - (Required) One or more `request_routing_rule` blocks as defined below.
- `sku` - (Required) A `sku` block as defined below.
- `zones` - (Optional) A collection of availability zones to spread the Application Gateway over.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>). They are also only supported for v2 SKUs (<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-autoscaling-zone-redundant>)

- `authentication_certificate` - (Optional) One or more `authentication_certificate` blocks as defined below.
- `trusted_root_certificate` - (Optional) One or more `trusted_root_certificate` blocks as defined below.
- `disabled_ssl_protocols` - (Optional / **Deprecated**) A list of SSL Protocols which should be disabled on this Application Gateway. Possible values are `TLSv1_0`, `TLSv1_1` and `TLSv1_2`. ~> **NOTE:** `disabled_ssl_protocols` has been deprecated in favour of `disabled_protocols` in the `ssl_policy` block.
- `ssl_policy` (Optional) a `ssl_policy` block as defined below.
- `enable_http2` - (Optional) Is HTTP2 enabled on the application gateway resource? Defaults to `false`.
- `probe` - (Optional) One or more `probe` blocks as defined below.
- `ssl_certificate` - (Optional) One or more `ssl_certificate` blocks as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `url_path_map` - (Optional) One or more `url_path_map` blocks as defined below.
- `waf_configuration` - (Optional) A `waf_configuration` block as defined below.
- `custom_error_configuration` - (Optional) One or more `custom_error_configuration` blocks as defined below.
- `redirect_configuration` - (Optional) A `redirect_configuration` block as defined below.
- `autoscale_configuration` - (Optional) A `autoscale_configuration` block as defined below.

- `rewrite_rule_set` - (Optional) One or more `rewrite_rule_set` blocks as defined below. Only valid for v2 SKUs.
-

A `authentication_certificate` block supports the following:

- `name` - (Required) The Name of the Authentication Certificate to use.
 - `data` - (Required) The contents of the Authentication Certificate which should be used.
-

A `trusted_root_certificate` block supports the following:

- `name` - (Required) The Name of the Authentication Certificate to use.
 - `data` - (Required) The contents of the Authentication Certificate which should be used.
-

A `authentication_certificate` block, within the `backend_http_settings` block supports the following:

- `name` - (Required) The name of the Authentication Certificate.
-

A `backend_address_pool` block supports the following:

- `name` - (Required) The name of the Backend Address Pool.
 - `fqdns` - (Optional) A list of FQDN's which should be part of the Backend Address Pool.
 - `fqdn_list` - (Optional **Deprecated**) A list of FQDN's which should be part of the Backend Address Pool. This field has been deprecated in favour of `fqdns` and will be removed in v2.0 of the AzureRM Provider.
 - `ip_addresses` - (Optional) A list of IP Addresses which should be part of the Backend Address Pool.
 - `ip_address_list` - (Optional **Deprecated**) A list of IP Addresses which should be part of the Backend Address Pool. This field has been deprecated in favour of `ip_addresses` and will be removed in v2.0 of the AzureRM Provider.
-

A `backend_http_settings` block supports the following:

- `cookie_based_affinity` - (Required) Is Cookie-Based Affinity enabled? Possible values are `Enabled` and `Disabled`.
 - `affinity_cookie_name` - (Optional) The name of the affinity cookie.
 - `name` - (Required) The name of the Backend HTTP Settings Collection.
 - `path` - (Optional) The Path which should be used as a prefix for all HTTP requests.
 - `port` - (Required) The port which should be used for this Backend HTTP Settings Collection.
 - `probe_name` - (Required) The name of an associated HTTP Probe.
 - `protocol` - (Required) The Protocol which should be used. Possible values are `Http` and `Https`.
 - `request_timeout` - (Required) The request timeout in seconds, which must be between 1 and 86400 seconds.
 - `host_name` - (Optional) Host header to be sent to the backend servers. Cannot be set if `pick_host_name_from_backend_address` is set to `true`.
-

- `pick_host_name_from_backend_address` - (Optional) Whether host header should be picked from the host name of the backend server. Defaults to `false`.
 - `authentication_certificate` - (Optional) One or more `authentication_certificate` blocks.
 - `connection_draining` - (Optional) A `connection_draining` block as defined below.
-

A `connection_draining` block supports the following:

- `enabled` - (Required) If connection draining is enabled or not.
 - `drain_timeout_sec` - (Required) The number of seconds connection draining is active. Acceptable values are from 1 second to 3600 seconds.
-

A `frontend_ip_configuration` block supports the following:

- `name` - (Required) The name of the Frontend IP Configuration.
- `subnet_id` - (Required) The ID of the Subnet which the Application Gateway should be connected to.
- `private_ip_address` - (Optional) The Private IP Address to use for the Application Gateway.
- `public_ip_address_id` - (Optional) The ID of a Public IP Address which the Application Gateway should use.

NOTE: The Allocation Method for this Public IP Address should be set to `Dynamic`.

- `private_ip_address_allocation` - (Optional) The Allocation Method for the Private IP Address. Possible values are `Dynamic` and `Static`.
-

A `frontend_port` block supports the following:

- `name` - (Required) The name of the Frontend Port.
 - `port` - (Required) The port used for this Frontend Port.
-

A `gateway_ip_configuration` block supports the following:

- `name` - (Required) The Name of this Gateway IP Configuration.
 - `subnet_id` - (Required) The ID of a Subnet.
-

A `http_listener` block supports the following:

- `name` - (Required) The Name of the HTTP Listener.
- `frontend_ip_configuration_name` - (Required) The Name of the Frontend IP Configuration used for this HTTP Listener.
- `frontend_port_name` - (Required) The Name of the Frontend Port use for this HTTP Listener.
- `host_name` - (Optional) The Hostname which should be used for this HTTP Listener.
- `protocol` - (Required) The Protocol to use for this HTTP Listener. Possible values are `Http` and `Https`.

- `require_sni` - (Optional) Should Server Name Indication be Required? Defaults to `false`.
 - `ssl_certificate_name` - (Optional) The name of the associated SSL Certificate which should be used for this HTTP Listener.
 - `custom_error_configuration` - (Optional) One or more `custom_error_configuration` blocks as defined below.
-

A `identity` block supports the following:

- `type` - (Optional) The Managed Service Identity Type of this Application Gateway. The only possible value is `UserAssigned`. Defaults to `UserAssigned`.
 - `identity_ids` - (Required) Specifies a list with a single user managed identity id to be assigned to the Application Gateway.
-

A `match` block supports the following:

- `body` - (Optional) A snippet from the Response Body which must be present in the Response. Defaults to `*`.
 - `status_code` - (Optional) A list of allowed status codes for this Health Probe.
-

A `path_rule` block supports the following:

- `name` - (Required) The Name of the Path Rule.
 - `paths` - (Required) A list of Paths used in this Path Rule.
 - `backend_address_pool_name` - (Optional) The Name of the Backend Address Pool to use for this Path Rule. Cannot be set if `redirect_configuration_name` is set.
 - `backend_http_settings_name` - (Optional) The Name of the Backend HTTP Settings Collection to use for this Path Rule. Cannot be set if `redirect_configuration_name` is set.
 - `redirect_configuration_name` - (Optional) The Name of a Redirect Configuration to use for this Path Rule. Cannot be set if `backend_address_pool_name` or `backend_http_settings_name` is set.
 - `rewrite_rule_set_name` - (Optional) The Name of the Rewrite Rule Set which should be used for this URL Path Map. Only valid for v2 SKUs.
-

A `probe` block support the following:

- `host` - (Optional) The Hostname used for this Probe. If the Application Gateway is configured for a single site, by default the Host name should be specified as '127.0.0.1', unless otherwise configured in custom probe. Cannot be set if `pick_host_name_from_backend_http_settings` is set to `true`.
- `interval` - (Required) The Interval between two consecutive probes in seconds. Possible values range from 1 second to a maximum of 86,400 seconds.
- `name` - (Required) The Name of the Probe.
- `protocol` - (Required) The Protocol used for this Probe. Possible values are `Http` and `Https`.
- `path` - (Required) The Path used for this Probe.

- `timeout` - (Required) The Timeout used for this Probe, which indicates when a probe becomes unhealthy. Possible values range from 1 second to a maximum of 86,400 seconds.
 - `unhealthy_threshold` - (Required) The Unhealthy Threshold for this Probe, which indicates the amount of retries which should be attempted before a node is deemed unhealthy. Possible values are from 1 - 20 seconds.
 - `pick_host_name_from_backend_http_settings` - (Optional) Whether the host header should be picked from the backend http settings. Defaults to `false`.
 - `match` - (Optional) A `match` block as defined above.
 - `minimum_servers` - (Optional) The minimum number of servers that are always marked as healthy. Defaults to `0`.
-

A `request_routing_rule` block supports the following:

- `name` - (Required) The Name of this Request Routing Rule.
 - `rule_type` - (Required) The Type of Routing that should be used for this Rule. Possible values are `Basic` and `PathBasedRouting`.
 - `http_listener_name` - (Required) The Name of the HTTP Listener which should be used for this Routing Rule.
 - `backend_address_pool_name` - (Optional) The Name of the Backend Address Pool which should be used for this Routing Rule. Cannot be set if `redirect_configuration_name` is set.
 - `backend_http_settings_name` - (Optional) The Name of the Backend HTTP Settings Collection which should be used for this Routing Rule. Cannot be set if `redirect_configuration_name` is set.
 - `redirect_configuration_name` - (Optional) The Name of the Redirect Configuration which should be used for this Routing Rule. Cannot be set if either `backend_address_pool_name` or `backend_http_settings_name` is set.
 - `rewrite_rule_set_name` - (Optional) The Name of the Rewrite Rule Set which should be used for this Routing Rule. Only valid for v2 SKUs.
 - `url_path_map_name` - (Optional) The Name of the URL Path Map which should be associated with this Routing Rule.
-

A `sku` block supports the following:

- `name` - (Required) The Name of the SKU to use for this Application Gateway. Possible values are `Standard_Small`, `Standard_Medium`, `Standard_Large`, `Standard_v2`, `WAF_Medium`, `WAF_Large`, and `WAF_v2`.
 - `tier` - (Required) The Tier of the SKU to use for this Application Gateway. Possible values are `Standard`, `Standard_v2`, `WAF` and `WAF_v2`.
 - `capacity` - (Required) The Capacity of the SKU to use for this Application Gateway - which must be between 1 and 10, optional if `autoscale_configuration` is set
-

A `ssl_certificate` block supports the following:

- `name` - (Required) The Name of the SSL certificate that is unique within this Application Gateway
 - `data` - (Required) PFX certificate.
 - `password` - (Required) Password for the pfx file specified in data.
-

A `url_path_map` block supports the following:

- `name` - (Required) The Name of the URL Path Map.
- `default_backend_address_pool_name` - (Optional) The Name of the Default Backend Address Pool which should be used for this URL Path Map. Cannot be set if `default_redirect_configuration_name` is set.
- `default_backend_http_settings_name` - (Optional) The Name of the Default Backend HTTP Settings Collection which should be used for this URL Path Map. Cannot be set if `default_redirect_configuration_name` is set.
- `default_redirect_configuration_name` - (Optional) The Name of the Default Redirect Configuration which should be used for this URL Path Map. Cannot be set if either `default_backend_address_pool_name` or `default_backend_http_settings_name` is set.
- `default_rewrite_rule_set_name` - (Optional) The Name of the Default Rewrite Rule Set which should be used for this URL Path Map. Only valid for v2 SKUs.
- `path_rule` - (Required) One or more `path_rule` blocks as defined above.

A `ssl_policy` block supports the following:

- `disabled_protocols` - (Optional) A list of SSL Protocols which should be disabled on this Application Gateway. Possible values are `TLSv1_0`, `TLSv1_1` and `TLSv1_2`.

NOTE: `disabled_protocols` cannot be set when `policy_name` or `policy_type` are set.

- `policy_type` - (Optional) The Type of the Policy. Possible values are `Predefined` and `Custom`.

NOTE: `policy_type` is Required when `policy_name` is set - cannot be set if `disabled_protocols` is set.

When using a `policy_type` of `Predefined` the following fields are supported:

- `policy_name` - (Optional) The Name of the Policy e.g `AppGwSslPolicy20170401S`. Required if `policy_type` is set to `Predefined`. Possible values can change over time and are published here <https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-ssl-policy-overview> (<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-ssl-policy-overview>). Not compatible with `disabled_protocols`.

When using a `policy_type` of `Custom` the following fields are supported:

- `cipher_suites` - (Required) A List of accepted cipher suites. Possible values are:
`TLS_DHE_DSS_WITH_AES_128_CBC_SHA`, `TLS_DHE_DSS_WITH_AES_128_CBC_SHA256`,
`TLS_DHE_DSS_WITH_AES_256_CBC_SHA`, `TLS_DHE_DSS_WITH_AES_256_CBC_SHA256`,
`TLS_DHE_RSA_WITH_AES_128_CBC_SHA`, `TLS_DHE_RSA_WITH_AES_128_GCM_SHA256`,
`TLS_DHE_RSA_WITH_AES_256_CBC_SHA`, `TLS_DHE_RSA_WITH_AES_256_GCM_SHA384`,
`TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA`, `TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256`,
`TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA`,
`TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`,
`TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA`, `TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256`,
`TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA`, `TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384`,

TLS_RSA_WITH_3DES_EDE_CBC_SHA , TLS_RSA_WITH_AES_128_CBC_SHA , TLS_RSA_WITH_AES_128_CBC_SHA256 , TLS_RSA_WITH_AES_128_GCM_SHA256 , TLS_RSA_WITH_AES_256_CBC_SHA , TLS_RSA_WITH_AES_256_CBC_SHA256 and TLS_RSA_WITH_AES_256_GCM_SHA384 .

- `min_protocol_version` - (Required) The minimal TLS version. Possible values are `TLSv1_0` , `TLSv1_1` and `TLSv1_2` .

A `waf_configuration` block supports the following:

- `enabled` - (Required) Is the Web Application Firewall be enabled?
- `firewall_mode` - (Required) The Web Application Firewall Mode. Possible values are `Detection` and `Prevention` .
- `rule_set_type` - (Required) The Type of the Rule Set used for this Web Application Firewall.
- `rule_set_version` - (Required) The Version of the Rule Set used for this Web Application Firewall. Possible values are `2.2.9` , `3.0` , and `3.1` .
- `disabled_rule_group` - (Optional) one or more `disabled_rule_group` blocks as defined below.
- `file_upload_limit_mb` - (Optional) The File Upload Limit in MB. Accepted values are in the range `1 MB` to `500 MB`. Defaults to `100 MB`.
- `request_body_check` - (Optional) Is Request Body Inspection enabled? Defaults to `true` .
- `max_request_body_size_kb` - (Optional) The Maximum Request Body Size in KB. Accepted values are in the range `1 KB` to `128 KB`. Defaults to `128 KB`.
- `exclusion` - (Optional) one or more `exclusion` blocks as defined below.

A `disabled_rule_group` block supports the following:

- `rule_group_name` - (Required) The rule group where specific rules should be disabled. Accepted values are: `crs_20_protocol_violations` , `crs_21_protocol_anomalies` , `crs_23_request_limits` , `crs_30_http_policy` , `crs_35_bad_robots` , `crs_40_generic_attacks` , `crs_41_sql_injection_attacks` , `crs_41_xss_attacks` , `crs_42_tight_security` , `crs_45_trojans` , `General` , `REQUEST-911-METHOD-ENFORCEMENT` , `REQUEST-913-SCANNER-DETECTION` , `REQUEST-920-PROTOCOL-ENFORCEMENT` , `REQUEST-921-PROTOCOL-ATTACK` , `REQUEST-930-APPLICATION-ATTACK-LFI` , `REQUEST-931-APPLICATION-ATTACK-RFI` , `REQUEST-932-APPLICATION-ATTACK-RCE` , `REQUEST-933-APPLICATION-ATTACK-PHP` , `REQUEST-941-APPLICATION-ATTACK-XSS` , `REQUEST-942-APPLICATION-ATTACK-SQLI` , `REQUEST-943-APPLICATION-ATTACK-SESSION-FIXATION`
- `rules` - (Optional) A list of rules which should be disabled in that group. Disables all rules in the specified group if `rules` is not specified.

A `exclusion` block supports the following:

- `match_variable` - (Required) Match variable of the exclusion rule to exclude header, cookie or GET arguments. Possible values are `RequestHeaderNames` , `RequestArgNames` and `RequestCookieNames`
- `selector_match_operator` - (Optional) Operator which will be used to search in the variable content. Possible values are `Equals` , `StartsWith` , `EndsWith` , `Contains` . If empty will exclude all traffic on this `match_variable`

- `selector` - (Optional) String value which will be used for the filter operation. If empty will exclude all traffic on this `match_variable`
-

A `custom_error_configuration` block supports the following:

- `status_code` - (Required) Status code of the application gateway customer error. Possible values are `HttpStatus403` and `HttpStatus502`
 - `custom_error_page_url` - (Required) Error page URL of the application gateway customer error.
-

A `redirect_configuration` block supports the following:

- `name` - (Required) Unique name of the redirect configuration block
 - `redirect_type` - (Required) The type of redirect. Possible values are `Permanent`, `Temporary`, `Found` and `SeeOther`
 - `target_listener_name` - (Optional) The name of the listener to redirect to. Cannot be set if `target_url` is set.
 - `target_url` - (Optional) The Url to redirect the request to. Cannot be set if `target_listener_name` is set.
 - `include_path` - (Optional) Whether or not to include the path in the redirected Url. Defaults to `false`
 - `include_query_string` - (Optional) Whether or not to include the query string in the redirected Url. Default to `false`
-

A `autoscale_configuration` block supports the following:

- `min_capacity` - (Required) Minimum capacity for autoscaling.
 - `max_capacity` - (Optional) Maximum capacity for autoscaling.
-

A `rewrite_rule_set` block supports the following:

- `name` - (Required) Unique name of the rewrite rule set block
 - `rewrite_rule` - (Required) One or more `rewrite_rule` blocks as defined above.
-

A `rewrite_rule` block supports the following:

- `name` - (Required) Unique name of the rewrite rule block
 - `rule_sequence` - (Required) Rule sequence of the rewrite rule that determines the order of execution in a set.
 - `condition` - (Optional) One or more `condition` blocks as defined above.
 - `request_header_configuration` - (Optional) One or more `request_header_configuration` blocks as defined above.
 - `response_header_configuration` - (Optional) One or more `response_header_configuration` blocks as defined above.
-

A `condition` block supports the following:

- `variable` - (Required) The variable (<https://docs.microsoft.com/en-us/azure/application-gateway/rewrite-http-headers#server-variables>) of the condition.
 - `pattern` - (Required) The pattern, either fixed string or regular expression, that evaluates the truthfulness of the condition.
 - `ignore_case` - (Optional) Perform a case in-sensitive comparison. Defaults to `false`
 - `negate` - (Optional) Negate the result of the condition evaluation. Defaults to `false`
-

A `request_header_configuration` block supports the following:

- `header_name` - (Required) Header name of the header configuration.
 - `header_value` - (Required) Header value of the header configuration. To delete a request header set this property to an empty string.
-

A `response_header_configuration` block supports the following:

- `header_name` - (Required) Header name of the header configuration.
- `header_value` - (Required) Header value of the header configuration. To delete a response header set this property to an empty string.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Gateway.
- `authentication_certificate` - A list of `authentication_certificate` blocks as defined below.
- `backend_address_pool` - A list of `backend_address_pool` blocks as defined below.
- `backend_http_settings` - A list of `backend_http_settings` blocks as defined below.
- `frontend_ip_configuration` - A list of `frontend_ip_configuration` blocks as defined below.
- `frontend_port` - A list of `frontend_port` blocks as defined below.
- `gateway_ip_configuration` - A list of `gateway_ip_configuration` blocks as defined below.
- `enable_http2` - (Optional) Is HTTP2 enabled on the application gateway resource? Defaults to `false`.
- `http_listener` - A list of `http_listener` blocks as defined below.
- `probe` - A `probe` block as defined below.
- `request_routing_rule` - A list of `request_routing_rule` blocks as defined below.
- `ssl_certificate` - A list of `ssl_certificate` blocks as defined below.
- `url_path_map` - A list of `url_path_map` blocks as defined below.

- `custom_error_configuration` - A list of `custom_error_configuration` blocks as defined below.
 - `redirect_configuration` - A list of `redirect_configuration` blocks as defined below.
-

A `authentication_certificate` block exports the following:

- `id` - The ID of the Authentication Certificate.
-

A `authentication_certificate` block, within the `backend_http_settings` block exports the following:

- `id` - The ID of the Authentication Certificate.
-

A `backend_address_pool` block exports the following:

- `id` - The ID of the Backend Address Pool.
-

A `backend_http_settings` block exports the following:

- `id` - The ID of the Backend HTTP Settings Configuration.
 - `probe_id` - The ID of the associated Probe.
-

A `frontend_ip_configuration` block exports the following:

- `id` - The ID of the Frontend IP Configuration.
-

A `frontend_port` block exports the following:

- `id` - The ID of the Frontend Port.
-

A `gateway_ip_configuration` block exports the following:

- `id` - The ID of the Gateway IP Configuration.
-

A `http_listener` block exports the following:

- `id` - The ID of the HTTP Listener.
 - `frontend_ip_configuration_id` - The ID of the associated Frontend Configuration.
 - `frontend_port_id` - The ID of the associated Frontend Port.
 - `ssl_certificate_id` - The ID of the associated SSL Certificate.
-

A `path_rule` block exports the following:

- `id` - The ID of the Path Rule.
- `backend_address_pool_id` - The ID of the Backend Address Pool used in this Path Rule.
- `backend_http_settings_id` - The ID of the Backend HTTP Settings Collection used in this Path Rule.

- `redirect_configuration_id` - The ID of the Redirect Configuration used in this Path Rule.
 - `rewrite_rule_set_id` - The ID of the Rewrite Rule Set used in this Path Rule.
-

A `probe` block exports the following:

- `id` - The ID of the Probe.
-

A `request_routing_rule` block exports the following:

- `id` - The ID of the Request Routing Rule.
 - `http_listener_id` - The ID of the associated HTTP Listener.
 - `backend_address_pool_id` - The ID of the associated Backend Address Pool.
 - `backend_http_settings_id` - The ID of the associated Backend HTTP Settings Configuration.
 - `redirect_configuration_id` - The ID of the associated Redirect Configuration.
 - `rewrite_rule_set_id` - The ID of the associated Rewrite Rule Set.
 - `url_path_map_id` - The ID of the associated URL Path Map.
-

A `ssl_certificate` block exports the following:

- `id` - The ID of the SSL Certificate.
 - `public_cert_data` - The Public Certificate Data associated with the SSL Certificate.
-

A `url_path_map` block exports the following:

- `id` - The ID of the URL Path Map.
 - `default_backend_address_pool_id` - The ID of the Default Backend Address Pool.
 - `default_backend_http_settings_id` - The ID of the Default Backend HTTP Settings Collection.
 - `default_redirect_configuration_id` - The ID of the Default Redirect Configuration.
 - `path_rule` - A list of `path_rule` blocks as defined above.
-

A `custom_error_configuration` block exports the following:

- `id` - The ID of the Custom Error Configuration.
-

A `redirect_configuration` block exports the following:

- `id` - The ID of the Redirect Configuration.
-

A `rewrite_rule_set` block exports the following:

- `id` - The ID of the Rewrite Rule Set

Import

Application Gateway's can be imported using the `resource_id`, e.g.

```
terraform import azurerm_application_gateway.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/applicationGateways/myGateway1
```

azurerm_application_insights_analytics_item

Manages an Application Insights Analytics Item component.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location  = "West Europe"
}

resource "azurerm_application_insights" "test" {
  name                = "tf-test-appinsights"
  location            = "West Europe"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_type     = "web"
}

resource "azurerm_application_insights_analytics_item" "test" {
  name                = "testquery"
  application_insights_id = "${azurerm_application_insights.test.id}"
  content             = "requests //simple example query"
  scope              = "shared"
  type               = "query"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Insights Analytics Item. Changing this forces a new resource to be created.
- `application_insights_id` - (Required) The ID of the Application Insights component on which the Analytics Item exists. Changing this forces a new resource to be created.
- `type` - (Required) The type of Analytics Item to create. Can be one of `query`, `function`, `folder`, `recent`. Changing this forces a new resource to be created.
- `scope` - (Required) The scope for the Analytics Item. Can be `shared` or `user`. Changing this forces a new resource to be created. Must be `shared` for functions.
- `content` - (Required) The content for the Analytics Item, for example the query text if `type` is `query`.
- `function_alias` - (Optional) The alias to use for the function. Required when `type` is `function`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Insights Analytics Item.
- `time_created` - A string containing the time the Analytics Item was created.
- `time_modified` - A string containing the time the Analytics Item was last modified.
- `version` - A string indicating the version of the query format

Import

Application Insights Analytics Items can be imported using the `resource id`, e.g.

```
terraform import azurerm_application_insights_analytics_item.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/components/analyticsItems/11111111-1111-1111-1111-111111111111
```

Please Note: This is a Terraform Unique ID matching the format: `{appInsightsID}/analyticsItems/{itemId}` for items with `scope` set to `shared`, or `{appInsightsID}/myanalyticsItems/{itemId}` for items with `scope` set to `user`

To find the Analytics Item ID you can query the REST API using the `az rest` CLI command (<https://docs.microsoft.com/en-us/cli/azure/reference-index?view=azure-cli-latest#az-rest>), e.g.

```
az rest --method GET --uri "https://management.azure.com/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/components/appinsightstest/analyticsItems?api-version=2015-05-01"
```

azurerm_application_insights_api_key

Manages an Application Insights API key.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location = "West Europe"
}

resource "azurerm_application_insights" "test" {
  name                = "tf-test-appinsights"
  location            = "West Europe"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_type    = "web"
}

resource "azurerm_application_insights_api_key" "read_telemetry" {
  name                = "tf-test-appinsights-read-telemetry-api-key"
  application_insights_id = "${azurerm_application_insights.test.id}"
  read_permissions    = ["aggregate", "api", "draft", "extendqueries", "search"]
}

resource "azurerm_application_insights_api_key" "write_annotations" {
  name                = "tf-test-appinsights-write-annotations-api-key"
  application_insights_id = "${azurerm_application_insights.test.id}"
  write_permissions   = ["annotations"]
}

resource "azurerm_application_insights_api_key" "authenticate_sdk_control_channel" {
  name                = "tf-test-appinsights-authenticate-sdk-control-channel-api-key"
  application_insights_id = "${azurerm_application_insights.test.id}"
  read_permissions    = ["agentconfig"]
}

resource "azurerm_application_insights_api_key" "full_permissions" {
  name                = "tf-test-appinsights-full-permissions-api-key"
  application_insights_id = "${azurerm_application_insights.test.id}"
  read_permissions    = ["agentconfig", "aggregate", "api", "draft", "extendqueries", "search"]
  write_permissions   = ["annotations"]
}

output "read_telemetry_api_key" {
  value = "${azurerm_application_insights_api_key.read_telemetry.api_key}"
}

output "write_annotations_api_key" {
  value = "${azurerm_application_insights_api_key.write_annotations.api_key}"
}

output "authenticate_sdk_control_channel" {
  value = "${azurerm_application_insights_api_key.authenticate_sdk_control_channel.api_key}"
}

output "full_permissions_api_key" {
  value = "${azurerm_application_insights_api_key.full_permissions.api_key}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Insights API key. Changing this forces a new resource to be created.
- `application_insights_id` - (Required) The ID of the Application Insights component on which the API key operates. Changing this forces a new resource to be created.
- `read_permissions` - (Optional) Specifies the list of read permissions granted to the API key. Valid values are `agentconfig`, `aggregate`, `api`, `draft`, `extendqueries`, `search`. Please note these values are case sensitive. Changing this forces a new resource to be created.
- `write_permissions` - (Optional) Specifies the list of write permissions granted to the API key. Valid values are `annotations`. Please note these values are case sensitive. Changing this forces a new resource to be created.

Note: At least one read or write permission must be defined.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Insights API key.
- `api_key` - The API Key secret (Sensitive).

Import

Application Insights API keys can be imported using the `resource id`, e.g.

```
terraform import azurerm_application_insights_api_key.my_key/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/components/instance1/apikeys/00000000-0000-0000-0000-000000000000
```

Note: The secret `api_key` cannot be retrieved during an import. You will need to edit the state by hand to set the secret value if you happen to have it backed up somewhere.

azurerm_application_insights

Manages an Application Insights component.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location  = "West Europe"
}

resource "azurerm_application_insights" "test" {
  name                = "tf-test-appinsights"
  location            = "West Europe"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_type    = "web"
}

output "instrumentation_key" {
  value = "${azurerm_application_insights.test.instrumentation_key}"
}

output "app_id" {
  value = "${azurerm_application_insights.test.app_id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Insights component. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Application Insights component.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `application_type` - (Required) Specifies the type of Application Insights to create. Valid values are `ios` for *iOS*, `java` for *Java web*, `MobileCenter` for *App Center*, `Node.JS` for *Node.js*, `other` for *General*, `phone` for *Windows Phone*, `store` for *Windows Store* and `web` for *ASP.NET*. Please note these values are case sensitive; unmatched values are treated as *ASP.NET* by Azure. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Insights component.
- `app_id` - The App ID associated with this Application Insights component.
- `instrumentation_key` - The Instrumentation Key for this Application Insights component.

Import

Application Insights instances can be imported using the `resource id`, e.g.

```
terraform import azurerm_application_insights.instance1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/components/instance1
```

azurerm_application_insights_web_test

Manages an Application Insights WebTest.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location  = "West Europe"
}

resource "azurerm_application_insights" "test" {
  name                = "tf-test-appinsights"
  location             = "West Europe"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_type     = "web"
}

resource "azurerm_application_insights_web_test" "test" {
  name                = "tf-test-appinsights-webtest"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_insights_id = "${azurerm_application_insights.test.id}"
  kind                = "ping"
  frequency            = 300
  timeout              = 60
  enabled              = true
  geo_locations        = ["us-tx-sn1-azr", "us-il-ch1-azr"]

  configuration = <<XML
<WebTest Name="WebTest1" Id="ABD48585-0831-40CB-9069-682EA6BB3583" Enabled="True" CssProjectStructure=""
CssIteration="" Timeout="0" WorkItemIds="" xmlns="http://microsoft.com/schemas/VisualStudio/TeamTest/2010
" Description="" CredentialUserName="" CredentialPassword="" PreAuthenticate="True" Proxy="default" StopOn
nError="False" RecordedResultFile="" ResultsLocale="">
  <Items>
    <Request Method="GET" Guid="a5f10126-e4cd-570d-961c-cea43999a200" Version="1.1" Url="http://microsoft
.com" ThinkTime="0" Timeout="300" ParseDependentRequests="True" FollowRedirects="True" RecordResult="True
" Cache="False" ResponseTimeGoal="0" Encoding="utf-8" ExpectedHttpStatusCode="200" ExpectedResponseUrl=""
ReportingName="" IgnoreHttpStatusCode="False" />
  </Items>
</WebTest>
XML
}

output "webtest_id" {
  value = "${azurerm_application_insights_web_test.test.id}"
}

output "webtest_provisioning_state" {
  value = "${azurerm_application_insights_web_test.test.provisioning_state}"
}

output "webtests_synthetic_id" {
  value = "${azurerm_application_insights_web_test.test.synthetic_monitor_id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Insights WebTest. Changing this forces a new resource to be created.
- `application_insights_id` - (Required) The ID of the Application Insights component on which the WebTest operates. Changing this forces a new resource to be created.
- `location` - (Required) The location of the resource group.
- `kind` - (Required) The kind of web test that this web test watches. Choices are `ping` and `multistep`.
- `geo_locations` - (Required) A list of where to physically run the tests from to give global coverage for accessibility of your application.
- `configuration` - (Required) An XML configuration specification for a WebTest.
- `frequency` - (Optional) Interval in seconds between test runs for this WebTest. Default is `300`.
- `timeout` - (Optional) Seconds until this WebTest will timeout and fail. Default is `30`.
- `enabled` - (Optional) Is the test actively being monitored.
- `retry_enabled` - (Optional) Allow for retries should this WebTest fail.
- `description` - (Optional) Purpose/user defined descriptive test for this WebTest.
- `tags` - (Optional) Resource tags.

Import

Application Insights Web Tests can be imported using the `resource_id`, e.g.

```
terraform import azurerm_application_insights_web_test.my_test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/webtests/my_test
```

azurerm_application_insights_web_test

Manages an Application Insights WebTest.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location  = "West Europe"
}

resource "azurerm_application_insights" "test" {
  name                = "tf-test-appinsights"
  location             = "West Europe"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_type     = "web"
}

resource "azurerm_application_insights_web_test" "test" {
  name                = "tf-test-appinsights-webtest"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  application_insights_id = "${azurerm_application_insights.test.id}"
  kind                = "ping"
  frequency            = 300
  timeout              = 60
  enabled              = true
  geo_locations        = ["us-tx-sn1-azr", "us-il-ch1-azr"]

  configuration = <<XML
<WebTest Name="WebTest1" Id="ABD48585-0831-40CB-9069-682EA6BB3583" Enabled="True" CssProjectStructure=""
CssIteration="" Timeout="0" WorkItemIds="" xmlns="http://microsoft.com/schemas/VisualStudio/TeamTest/2010
" Description="" CredentialUserName="" CredentialPassword="" PreAuthenticate="True" Proxy="default" StopOn
nError="False" RecordedResultFile="" ResultsLocale="">
  <Items>
    <Request Method="GET" Guid="a5f10126-e4cd-570d-961c-cea43999a200" Version="1.1" Url="http://microsoft
.com" ThinkTime="0" Timeout="300" ParseDependentRequests="True" FollowRedirects="True" RecordResult="True
" Cache="False" ResponseTimeGoal="0" Encoding="utf-8" ExpectedHttpStatusCode="200" ExpectedResponseUrl=""
ReportingName="" IgnoreHttpStatusCode="False" />
  </Items>
</WebTest>
XML
}

output "webtest_id" {
  value = "${azurerm_application_insights_web_test.test.id}"
}

output "webtest_provisioning_state" {
  value = "${azurerm_application_insights_web_test.test.provisioning_state}"
}

output "webtests_synthetic_id" {
  value = "${azurerm_application_insights_web_test.test.synthetic_monitor_id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Insights WebTest. Changing this forces a new resource to be created.
- `application_insights_id` - (Required) The ID of the Application Insights component on which the WebTest operates. Changing this forces a new resource to be created.
- `location` - (Required) The location of the resource group.
- `kind` - (Required) The kind of web test that this web test watches. Choices are `ping` and `multistep`.
- `geo_locations` - (Required) A list of where to physically run the tests from to give global coverage for accessibility of your application.
- `configuration` - (Required) An XML configuration specification for a WebTest.
- `frequency` - (Optional) Interval in seconds between test runs for this WebTest. Default is `300`.
- `timeout` - (Optional) Seconds until this WebTest will timeout and fail. Default is `30`.
- `enabled` - (Optional) Is the test actively being monitored.
- `retry_enabled` - (Optional) Allow for retries should this WebTest fail.
- `description` - (Optional) Purpose/user defined descriptive test for this WebTest.
- `tags` - (Optional) Resource tags.

Import

Application Insights Web Tests can be imported using the `resource_id`, e.g.

```
terraform import azurerm_application_insights_web_test.my_test/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/webtests/my_test
```

azurerm_application_security_group

Manages an Application Security Group.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "tf-test"
  location  = "West Europe"
}

resource "azurerm_application_security_group" "test" {
  name                = "tf-appsecuritygroup"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    Hello = "World"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Security Group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Application Security Group.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Application Security Group.

Import

Application Security Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_application_security_group.securitygroup1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/applicationSecurityGroups/securitygroup1
```

azurerm_app_service_active_slot

Promotes an App Service Slot to Production within an App Service.

Note: When using Slots - the `app_settings`, `connection_string` and `site_config` blocks on the `azurerm_app_service` resource will be overwritten when promoting a Slot using the `azurerm_app_service_active_slot` resource.

Example Usage

```
resource "random_id" "server" {
  # ...
}

resource "azurerm_resource_group" "test" {
  # ...
}

resource "azurerm_app_service_plan" "test" {
  # ...
}

resource "azurerm_app_service" "test" {
  # ...
}

resource "azurerm_app_service_slot" "test" {
  # ...
}

resource "azurerm_app_service_active_slot" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_name     = "${azurerm_app_service.test.name}"
  app_service_slot_name = "${azurerm_app_service_slot.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which the App Service exists. Changing this forces a new resource to be created.
- `app_service_name` - (Required) The name of the App Service within which the Slot exists. Changing this forces a new resource to be created.
- `app_service_slot_name` - (Required) The name of the App Service Slot which should be promoted to the Production Slot within the App Service.

azurerm_app_service_certificate

Manages an App Service certificate.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_app_service_certificate" "example" {
  name                = "example-cert"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"
  pfx_blob            = "${filebase64("certificate.pfx")}"
  password            = "terraform"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the certificate. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the certificate. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `pfx_blob` - (Optional) The base64-encoded contents of the certificate. Changing this forces a new resource to be created.

NOTE: Either `pfx_blob` or `key_vault_secret_id` must be set - but not both.

- `password` - (Optional) The password to access the certificate's private key. Changing this forces a new resource to be created.
- `key_vault_secret_id` - (Optional) The ID of the Key Vault secret. Changing this forces a new resource to be created.

NOTE: If using `key_vault_secret_id`, the magic Resource Principal with id of `abfa0a7c-a6b6-4736-8310-5855508787cd` must have 'Secret get' and 'Certificate get' permissions on the Key Vault containing the certificate. (Source: App Service Blog (<https://azure.github.io/AppService/2016/05/24/Deploying-Azure-Web-App-Certificate-through-Key-Vault.html>))

Attributes Reference

The following attributes are exported:

- `id` - The App Service certificate ID.
- `friendly_name` - The friendly name of the certificate.
- `subject_name` - The subject name of the certificate.
- `host_names` - List of host names the certificate applies to.
- `issuer` - The name of the certificate issuer.
- `issue_date` - The issue date for the certificate.
- `expiration_date` - The expiration date for the certificate.
- `thumbprint` - The thumbprint for the certificate.

Import

App Service certificates can be imported using the `resource id`, e.g.

```
terraform import azurerm_app_certificate.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/certificates/certificate1
```

azurerm_app_service_certificate_order

Manages an App Service Certificate Order.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_app_service_certificate_order" "example" {
  name                        = "example-cert-order"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "global"
  distinguished_name         = "CN=example.com"
  product_type               = "standard"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the certificate. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the certificate. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `auto_renew` - (Optional) true if the certificate should be automatically renewed when it expires; otherwise, false.
- `csr` - (Optional) Last CSR that was created for this order.
- `distinguished_name` - (Optional) The Distinguished Name for the App Service Certificate Order.

NOTE: Either `csr` or `distinguished_name` must be set - but not both.

- `key_size` - (Optional) Certificate key size.
- `product_type` - (Optional) Certificate product type, such as `Standard` or `Wildcard`.
- `validity_in_years` - (Optional) Duration in years (must be between 1 and 3).

Attributes Reference

The following attributes are exported:

- `id` - The App Service Certificate Order ID.
 - `certificates` - State of the Key Vault secret. A `certificates` block as defined below.
 - `domain_verification_token` - Domain verification token.
 - `status` - Current order status.
 - `expiration_time` - Certificate expiration time.
 - `is_private_key_external` - Whether the private key is external or not.
 - `app_service_certificate_not_renewable_reasons` - Reasons why App Service Certificate is not renewable at the current moment.
 - `signed_certificate_thumbprint` - Certificate thumbprint for signed certificate.
 - `root_thumbprint` - Certificate thumbprint for root certificate.
 - `intermediate_thumbprint` - Certificate thumbprint intermediate certificate.
 - `tags` - A mapping of tags to assign to the resource.
-

`certificates` supports the following:

- `certificate_name` - The name of the App Service Certificate.
- `key_vault_id` - Key Vault resource Id.
- `key_vault_secret_name` - Key Vault secret name.
- `provisioning_state` - Status of the Key Vault secret.

Import

App Service certificate order can be imported using the `resource_id`, e.g.

```
terraform import azurerm_app_certificate_order.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.CertificateRegistration/certificateOrders/certificateorder1
```

azurerm_app_service_custom_hostname_binding

Manages a Hostname Binding within an App Service.

Example Usage

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name      = "some-resource-group"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "some-app-service-plan"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}

resource "azurerm_app_service" "test" {
  name                = "${random_id.server.hex}"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"
}

resource "azurerm_app_service_custom_hostname_binding" "test" {
  hostname            = "www.mywebsite.com"
  app_service_name    = "${azurerm_app_service.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `hostname` - (Required) Specifies the Custom Hostname to use for the App Service, example `www.example.com`. Changing this forces a new resource to be created.

NOTE: A CNAME needs to be configured from this Hostname to the Azure Website - otherwise Azure will reject the Hostname Binding.

- `app_service_name` - (Required) The name of the App Service in which to add the Custom Hostname Binding. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the App Service exists. Changing this forces a new resource to be created.
- `ssl_state` - (Optional) The SSL type. Possible values are `IpBasedEnabled` and `SniEnabled` . Changing this forces a new resource to be created.
- `thumbprint` - (Optional) The SSL certificate thumbprint. Changing this forces a new resource to be created.

NOTE: `thumbprint` must be specified when `ssl_state` is set.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the App Service Custom Hostname Binding
- `virtual_ip` - The virtual IP address assigned to the hostname if IP based SSL is enabled.

Import

App Service Custom Hostname Bindings can be imported using the `resource id` , e.g.

```
terraform import azurerm_app_service_custom_hostname_binding.mywebsite /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/sites/instance1/hostnameBindings/mywebsite.com
```

azurerm_app_service

Manages an App Service (within an App Service Plan).

Note: When using Slots - the `app_settings`, `connection_string` and `site_config` blocks on the `azurerm_app_service` resource will be overwritten when promoting a Slot using the `azurerm_app_service_active_slot` resource.

Example Usage

This example provisions a Windows App Service. Other examples of the `azurerm_app_service` resource can be found in the `./examples/app-service` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/app-service>)

```
resource "azurerm_resource_group" "test" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "example-appserviceplan"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}

resource "azurerm_app_service" "test" {
  name                = "example-app-service"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"

  site_config {
    dotnet_framework_version = "v4.0"
    scm_type                 = "LocalGit"
  }

  app_settings = {
    "SOME_KEY" = "some-value"
  }

  connection_string {
    name = "Database"
    type = "SQLServer"
    value = "Server=some-server.mydomain.com;Integrated Security=SSPI"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the App Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the App Service.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `app_service_plan_id` - (Required) The ID of the App Service Plan within which to create this App Service.
- `app_settings` - (Optional) A key-value pair of App Settings.
- `auth_settings` - (Optional) A `auth_settings` block as defined below.
- `storage_account` - (Optional) One or more `storage_account` blocks as defined below.
- `connection_string` - (Optional) One or more `connection_string` blocks as defined below.
- `client_affinity_enabled` - (Optional) Should the App Service send session affinity cookies, which route client requests in the same session to the same instance?
- `client_cert_enabled` - (Optional) Does the App Service require client certificates for incoming requests? Defaults to `false`.
- `enabled` - (Optional) Is the App Service Enabled?
- `https_only` - (Optional) Can the App Service only be accessed via HTTPS? Defaults to `false`.
- `logs` - (Optional) A `logs` block as defined below.
- `site_config` - (Optional) A `site_config` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `identity` - (Optional) A Managed Service Identity block as defined below.

A `storage_account` block supports the following:

- `name` - (Required) The name of the storage account identifier.
- `type` - (Required) The type of storage. Possible values are `AzureBlob` and `AzureFiles`.
- `account_name` - (Required) The name of the storage account.
- `share_name` - (Required) The name of the file share (container name, for Blob storage).
- `access_key` - (Required) The access key for the storage account.
- `mount_path` - (Optional) The path to mount the storage within the site's runtime environment.

A `connection_string` block supports the following:

- `name` - (Required) The name of the Connection String.

- `type` - (Required) The type of the Connection String. Possible values are `APIHub`, `Custom`, `DocDb`, `EventHub`, `MySQL`, `NotificationHub`, `PostgreSQL`, `RedisCache`, `ServiceBus`, `SQLAzure` and `SQLServer`.
 - `value` - (Required) The value for the Connection String.
-

A `identity` block supports the following:

- `type` - (Required) Specifies the identity type of the App Service. Possible values are `SystemAssigned` (where Azure will generate a Service Principal for you), `UserAssigned` where you can specify the Service Principal IDs in the `identity_ids` field, and `SystemAssigned, UserAssigned` which assigns both a system managed identity as well as the specified user assigned identities.

NOTE: When `type` is set to `SystemAssigned`, The assigned `principal_id` and `tenant_id` can be retrieved after the App Service has been created. More details are available below.

- `identity_ids` - (Optional) Specifies a list of user managed identity ids to be assigned. Required if `type` is `UserAssigned`.
-

A `logs` block supports the following:

- `application_logs` - (Optional) An `application_logs` block as defined below.
 - `http_logs` - (Optional) An `http_logs` block as defined below.
-

An `application_logs` block supports the following:

- `azure_blob_storage` - (Optional) An `azure_blob_storage` block as defined below.
-

An `http_logs` block supports *one* of the following:

- `file_system` - (Optional) A `file_system` block as defined below.
 - `azure_blob_storage` - (Optional) An `azure_blob_storage` block as defined below.
-

An `azure_blob_storage` block supports the following:

- `level` - (Required) The level at which to log. Possible values include `Error`, `Warning`, `Information`, `Verbose` and `Off`. **NOTE:** this field is not available for `http_logs`
 - `sas_url` - (Required) The URL to the storage container, with a Service SAS token appended. **NOTE:** there is currently no means of generating Service SAS tokens with the `azurerm` provider.
 - `retention_in_days` - (Required) The number of days to retain logs for.
-

A `file_system` block supports the following:

- `retention_in_days` - (Required) The number of days to retain logs for.
 - `retention_in_mb` - (Required) The maximum size in megabytes that http log files can use before being removed.
-

A `site_config` block supports the following:

- `always_on` - (Optional) Should the app be loaded at all times? Defaults to `false`.
- `app_command_line` - (Optional) App command line to launch, e.g. `/sbin/myserv -b 0.0.0.0`.
- `cors` - (Optional) A `cors` block as defined below.
- `default_documents` - (Optional) The ordering of default documents to load, if an address isn't specified.
- `dotnet_framework_version` - (Optional) The version of the .net framework's CLR used in this App Service. Possible values are `v2.0` (which will use the latest version of the .net framework for the .net CLR v2 - currently `.net 3.5`) and `v4.0` (which corresponds to the latest version of the .net CLR v4 - which at the time of writing is `.net 4.7.1`). For more information on which .net CLR version to use based on the .net framework you're targeting - please see this table (https://en.wikipedia.org/wiki/.NET_Framework_version_history#Overview). Defaults to `v4.0`.
- `ftps_state` - (Optional) State of FTP / FTPS service for this App Service. Possible values include: `AllAllowed`, `FtpsOnly` and `Disabled`.
- `http2_enabled` - (Optional) Is HTTP2 Enabled on this App Service? Defaults to `false`.
- `ip_restriction` - (Optional) A List of objects (</docs/configuration/attr-as-blocks.html>) representing ip restrictions as defined below.
- `java_version` - (Optional) The version of Java to use. If specified `java_container` and `java_container_version` must also be specified. Possible values are `1.7`, `1.8` and `11`.
- `java_container` - (Optional) The Java Container to use. If specified `java_version` and `java_container_version` must also be specified. Possible values are `JETTY` and `TOMCAT`.
- `java_container_version` - (Optional) The version of the Java Container to use. If specified `java_version` and `java_container` must also be specified.
- `local_mysql_enabled` - (Optional) Is "MySQL In App" Enabled? This runs a local MySQL instance with your app and shares resources from the App Service plan.

NOTE: MySQL In App is not intended for production environments and will not scale beyond a single instance. Instead you may wish to use Azure Database for MySQL ([/docs/providers/azurerm/r/mysql_database.html](/docs/providers/azurerm/r/mysql_database.html)).

- `linux_fx_version` - (Optional) Linux App Framework and version for the App Service. Possible options are a Docker container (`DOCKER|<user/image:tag>`), a base-64 encoded Docker Compose file (`COMPOSE|${filebase64("compose.yml")}`) or a base-64 encoded Kubernetes Manifest (`KUBE|${filebase64("kubernetes.yml")}`).
- `windows_fx_version` - (Optional) The Windows Docker container image (`DOCKER|<user/image:tag>`)

Additional examples of how to run Containers via the `azurerm_app_service` resource can be found in the `./examples/app-service` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/app-service>).

- `managed_pipeline_mode` - (Optional) The Managed Pipeline Mode. Possible values are `Integrated` and `Classic`. Defaults to `Integrated`.
- `min_tls_version` - (Optional) The minimum supported TLS version for the app service. Possible values are `1.0`, `1.1`, and `1.2`. Defaults to `1.2` for new app services.

- `php_version` - (Optional) The version of PHP to use in this App Service. Possible values are `5.5`, `5.6`, `7.0`, `7.1` and `7.2`.
- `python_version` - (Optional) The version of Python to use in this App Service. Possible values are `2.7` and `3.4`.
- `remote_debugging_enabled` - (Optional) Is Remote Debugging Enabled? Defaults to `false`.
- `remote_debugging_version` - (Optional) Which version of Visual Studio should the Remote Debugger be compatible with? Possible values are `VS2012`, `VS2013`, `VS2015` and `VS2017`.
- `scm_type` - (Optional) The type of Source Control enabled for this App Service. Defaults to `None`. Possible values are: `BitbucketGit`, `BitbucketHg`, `CodePlexGit`, `CodePlexHg`, `Dropbox`, `ExternalGit`, `ExternalHg`, `GitHub`, `LocalGit`, `None`, `OneDrive`, `Tfs`, `VSO` and `VSTSRM`.
- `use_32_bit_worker_process` - (Optional) Should the App Service run in 32 bit mode, rather than 64 bit mode?

NOTE: when using an App Service Plan in the `Free` or `Shared` Tiers `use_32_bit_worker_process` must be set to `true`.

- `virtual_network_name` - (Optional) The name of the Virtual Network which this App Service should be attached to.
- `websockets_enabled` - (Optional) Should WebSockets be enabled?

A `cors` block supports the following:

- `allowed_origins` - (Optional) A list of origins which should be able to make cross-origin calls. `*` can be used to allow all calls.
- `support_credentials` - (Optional) Are credentials supported?

A `auth_settings` block supports the following:

- `enabled` - (Required) Is Authentication enabled?
- `active_directory` - (Optional) A `active_directory` block as defined below.
- `additional_login_params` - (Optional) Login parameters to send to the OpenID Connect authorization endpoint when a user logs in. Each parameter must be in the form "key=value".
- `allowed_external_redirect_urls` - (Optional) External URLs that can be redirected to as part of logging in or logging out of the app.
- `default_provider` - (Optional) The default provider to use when multiple providers have been set up. Possible values are `AzureActiveDirectory`, `Facebook`, `Google`, `MicrosoftAccount` and `Twitter`.

NOTE: When using multiple providers, the default provider must be set for settings like `unauthenticated_client_action` to work.

- `facebook` - (Optional) A `facebook` block as defined below.
- `google` - (Optional) A `google` block as defined below.

- `issuer` - (Optional) Issuer URI. When using Azure Active Directory, this value is the URI of the directory tenant, e.g. <https://sts.windows.net/{tenant-guid}/> (<https://sts.windows.net/%7Btenant-guid%7D/>).
 - `microsoft` - (Optional) A `microsoft` block as defined below.
 - `runtime_version` - (Optional) The runtime version of the Authentication/Authorization module.
 - `token_refresh_extension_hours` - (Optional) The number of hours after session token expiration that a session token can be used to call the token refresh API. Defaults to 72.
 - `token_store_enabled` - (Optional) If enabled the module will durably store platform-specific security tokens that are obtained during login flows. Defaults to false.
 - `twitter` - (Optional) A `twitter` block as defined below.
 - `unauthenticated_client_action` - (Optional) The action to take when an unauthenticated client attempts to access the app. Possible values are `AllowAnonymous` and `RedirectToLoginPage`.
-

A `active_directory` block supports the following:

- `client_id` - (Required) The Client ID of this relying party application. Enables OpenIDConnection authentication with Azure Active Directory.
 - `client_secret` - (Optional) The Client Secret of this relying party application. If no secret is provided, implicit flow will be used.
 - `allowed_audiences` (Optional) Allowed audience values to consider when validating JWTs issued by Azure Active Directory.
-

A `facebook` block supports the following:

- `app_id` - (Required) The App ID of the Facebook app used for login
 - `app_secret` - (Required) The App Secret of the Facebook app used for Facebook Login.
 - `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Facebook Login authentication. <https://developers.facebook.com/docs/facebook-login> (<https://developers.facebook.com/docs/facebook-login>)
-

A `google` block supports the following:

- `client_id` - (Required) The OpenID Connect Client ID for the Google web application.
 - `client_secret` - (Required) The client secret associated with the Google web application.
 - `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Google Sign-In authentication. <https://developers.google.com/identity/sign-in/web/> (<https://developers.google.com/identity/sign-in/web/>)
-

A `ip_restriction` block supports the following:

- `ip_address` - (Optional) The IP Address used for this IP Restriction.
- `subnet_mask` - (Optional) The Subnet mask used for this IP Restriction. Defaults to `255.255.255.255`.
- `virtual_network_subnet_id` - (Optional) The Virtual Network Subnet ID used for this IP Restriction.

NOTE: One of either `ip_address` or `virtual_network_subnet_id` must be specified

A `microsoft` block supports the following:

- `client_id` - (Required) The OAuth 2.0 client ID that was created for the app used for authentication.
- `client_secret` - (Required) The OAuth 2.0 client secret that was created for the app used for authentication.
- `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Microsoft Account authentication.
<https://msdn.microsoft.com/en-us/library/dn631845.aspx> (<https://msdn.microsoft.com/en-us/library/dn631845.aspx>)

A `backup` block supports the following:

- `name` (Required) Specifies the name for this Backup.
- `enabled` - (Required) Is this Backup enabled?
- `storage_account_url` (Optional) The SAS URL to a Storage Container where Backups should be saved.
- `schedule` - (Optional) A `schedule` block as defined below.

A `schedule` block supports the following:

- `frequency_interval` - (Required) Sets how often the backup should be executed.
- `frequency_unit` - (Optional) Sets the unit of time for how often the backup should be executed. Possible values are `Day` or `Hour`.
- `keep_at_least_one_backup` - (Optional) Should at least one backup always be kept in the Storage Account by the Retention Policy, regardless of how old it is?
- `retention_period_in_days` - (Optional) Specifies the number of days after which Backups should be deleted.
- `start_time` - (Optional) Sets when the schedule should start working.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the App Service.
- `default_site_hostname` - The Default Hostname associated with the App Service - such as `mysite.azurewebsites.net`
- `outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12`
- `possible_outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12,52.143.43.17` - not all of which are necessarily in use. Superset of `outbound_ip_addresses`.
- `source_control` - A `source_control` block as defined below, which contains the Source Control information when `scm_type` is set to `LocalGit`.

- `site_credential` - A `site_credential` block as defined below, which contains the site-level credentials used to publish to this App Service.
- `identity` - An `identity` block as defined below, which contains the Managed Service Identity information for this App Service.

A `identity` block exports the following:

- `principal_id` - The Principal ID for the Service Principal associated with the Managed Service Identity of this App Service.
- `tenant_id` - The Tenant ID for the Service Principal associated with the Managed Service Identity of this App Service.

You can access the Principal ID via `${azurerm_app_service.test.identity.0.principal_id}` and the Tenant ID via `${azurerm_app_service.test.identity.0.tenant_id}`

A `site_credential` block exports the following:

- `username` - The username which can be used to publish to this App Service
- `password` - The password associated with the username, which can be used to publish to this App Service.

NOTE: both `username` and `password` for the `site_credential` block are only exported when `scm_type` is set to `LocalGit`

A `source_control` block exports the following:

- `repo_url` - URL of the Git repository for this App Service.
- `branch` - Branch name of the Git repository for this App Service.

Import

App Services can be imported using the `resource_id`, e.g.

```
terraform import azurerm_app_service.instance1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/sites/instance1
```

azurerm_app_service_plan

Manages an App Service Plan component.

Example Usage (Dedicated)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "api-appserviceplan-pro"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}
```

Example Usage (Shared / Consumption Plan)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "api-appserviceplan-pro"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  kind                = "FunctionApp"

  sku {
    tier = "Dynamic"
    size = "Y1"
  }
}
```

Example Usage (Linux)

```

resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "api-appserviceplan-pro"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  kind                = "Linux"
  reserved            = true

  sku {
    tier = "Standard"
    size = "S1"
  }
}

```

Example Usage (Windows Container)

```

resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "api-appserviceplan-pro"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  kind                = "xenon"
  is_xenon            = true

  sku {
    tier = "PremiumContainer"
    size = "PC2"
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the App Service Plan component. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the App Service Plan component.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.

- `kind` - (Optional) The kind of the App Service Plan to create. Possible values are `Windows` (also available as `App`), `Linux`, `elastic` (for Premium Consumption) and `FunctionApp` (for a Consumption Plan). Defaults to `Windows`. Changing this forces a new resource to be created.

NOTE: When creating a `Linux` App Service Plan, the `reserved` field must be set to `true`.

- `maximum_elastic_worker_count` - The maximum number of total workers allowed for this ElasticScaleEnabled App Service Plan.
- `sku` - (Required) A `sku` block as documented below.
- `app_service_environment_id` - (Optional) The ID of the App Service Environment where the App Service Plan should be located. Changing forces a new resource to be created.

NOTE: Attaching to an App Service Environment requires the App Service Plan use a `Premium` SKU (when using an ASEv1) and the `Isolated` SKU (for an ASEv2).

- `reserved` - (Optional) Is this App Service Plan `Reserved`. Defaults to `false`.
- `per_site_scaling` - (Optional) Can Apps assigned to this App Service Plan be scaled independently? If set to `false` apps assigned to this plan will scale to all instances of the plan. Defaults to `false`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`sku` supports the following:

- `tier` - (Required) Specifies the plan's pricing tier.
- `size` - (Required) Specifies the plan's instance size.
- `capacity` - (Optional) Specifies the number of workers associated with this App Service Plan.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the App Service Plan component.
- `maximum_number_of_workers` - The maximum number of workers supported with the App Service Plan's `sku`.

Import

App Service Plan instances can be imported using the `resource id`, e.g.

```
terraform import azurerm_app_service_plan.instance1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/serverfarms/instance1
```

azurerm_app_service_slot

Manages an App Service Slot (within an App Service).

Note: When using Slots - the `app_settings`, `connection_string` and `site_config` blocks on the `azurerm_app_service` resource will be overwritten when promoting a Slot using the `azurerm_app_service_active_slot` resource.

Example Usage (.net 4.x)

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name     = "some-resource-group"
  location = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "some-app-service-plan"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}

resource "azurerm_app_service" "test" {
  name                = "${random_id.server.hex}"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"

  site_config {
    dotnet_framework_version = "v4.0"
  }

  app_settings = {
    "SOME_KEY" = "some-value"
  }

  connection_string {
    name = "Database"
    type = "SQLServer"
    value = "Server=some-server.mydomain.com;Integrated Security=SSPI"
  }
}
```

```
}

resource "azurerm_app_service_slot" "test" {
  name                = "${random_id.server.hex}"
  app_service_name    = "${azurerm_app_service.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"

  site_config {
    dotnet_framework_version = "v4.0"
  }

  app_settings = {
    "SOME_KEY" = "some-value"
  }

  connection_string {
    name = "Database"
    type = "SQLServer"
    value = "Server=some-server.mydomain.com;Integrated Security=SSPI"
  }
}
```

Example Usage (Java 1.8)

```

resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name      = "some-resource-group"
  location  = "West Europe"
}

resource "azurerm_app_service_plan" "test" {
  name                = "some-app-service-plan"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}

resource "azurerm_app_service" "test" {
  name                = "${random_id.server.hex}"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"

  site_config {
    java_version      = "1.8"
    java_container    = "JETTY"
    java_container_version = "9.3"
  }
}

resource "azurerm_app_service_slot" "test" {
  name                = "${random_id.server.hex}"
  app_service_name    = "${azurerm_app_service.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"

  site_config {
    java_version      = "1.8"
    java_container    = "JETTY"
    java_container_version = "9.3"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the App Service Slot component. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the App Service Slot component.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `app_service_plan_id` - (Required) The ID of the App Service Plan within which to create this App Service Slot. Changing this forces a new resource to be created.
- `app_service_name` - (Required) The name of the App Service within which to create the App Service Slot. Changing this forces a new resource to be created.
- `app_settings` - (Optional) A key-value pair of App Settings.
- `auth_settings` - (Optional) A `auth_settings` block as defined below.
- `connection_string` - (Optional) An `connection_string` block as defined below.
- `client_affinity_enabled` - (Optional) Should the App Service Slot send session affinity cookies, which route client requests in the same session to the same instance?
- `enabled` - (Optional) Is the App Service Slot Enabled?
- `https_only` - (Optional) Can the App Service Slot only be accessed via HTTPS? Defaults to `false`.
- `site_config` - (Optional) A `site_config` object as defined below.
- `identity` - (Optional) A Managed Service Identity block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`connection_string` supports the following:

- `name` - (Required) The name of the Connection String.
- `type` - (Required) The type of the Connection String. Possible values are `APIHub`, `Custom`, `DocDb`, `EventHub`, `MySQL`, `NotificationHub`, `PostgreSQL`, `RedisCache`, `ServiceBus`, `SQLAzure` and `SQLServer`.
- `value` - (Required) The value for the Connection String.

`site_config` supports the following:

- `app_command_line` - (Optional) App command line to launch, e.g. `/sbin/myserver -b 0.0.0.0`.
- `always_on` - (Optional) Should the app be loaded at all times? Defaults to `false`.
- `cors` - (Optional) A `cors` block as defined below.
- `default_documents` - (Optional) The ordering of default documents to load, if an address isn't specified.
- `dotnet_framework_version` - (Optional) The version of the .net framework's CLR used in this App Service Slot. Possible values are `v2.0` (which will use the latest version of the .net framework for the .net CLR v2 - currently .net 3.5) and `v4.0` (which corresponds to the latest version of the .net CLR v4 - which at the time of writing is .net

4.7.1). For more information on which .net CLR version to use based on the .net framework you're targeting - please see this table (https://en.wikipedia.org/wiki/.NET_Framework_version_history#Overview). Defaults to v4.0 .

- `http2_enabled` - (Optional) Is HTTP2 Enabled on this App Service? Defaults to `false` .
- `ip_restriction` - (Optional) A List of objects (</docs/configuration/attr-as-blocks.html>) representing ip restrictions as defined below.
- `java_container` - (Optional) The Java Container to use. If specified `java_version` and `java_container_version` must also be specified. Possible values are `JETTY` and `TOMCAT` .
- `java_container_version` - (Optional) The version of the Java Container to use. If specified `java_version` and `java_container` must also be specified.
- `java_version` - (Optional) The version of Java to use. If specified `java_container` and `java_container_version` must also be specified. Possible values are `1.7` , `1.8` and `11` .
- `local_mysql_enabled` - (Optional) Is "MySQL In App" Enabled? This runs a local MySQL instance with your app and shares resources from the App Service plan.

NOTE: MySQL In App is not intended for production environments and will not scale beyond a single instance. Instead you may wish to use Azure Database for MySQL ([/docs/providers/azurerms/r/mysql_database.html](/docs/providers/azurerms/r/mysql_database.html)).

- `managed_pipeline_mode` - (Optional) The Managed Pipeline Mode. Possible values are `Integrated` and `Classic` . Defaults to `Integrated` .
- `min_tls_version` - (Optional) The minimum supported TLS version for the app service. Possible values are `1.0` , `1.1` , and `1.2` . Defaults to `1.2` for new app services.
- `php_version` - (Optional) The version of PHP to use in this App Service Slot. Possible values are `5.5` , `5.6` , `7.0` , `7.1` and `7.2` .
- `python_version` - (Optional) The version of Python to use in this App Service Slot. Possible values are `2.7` and `3.4` .
- `remote_debugging_enabled` - (Optional) Is Remote Debugging Enabled? Defaults to `false` .
- `remote_debugging_version` - (Optional) Which version of Visual Studio should the Remote Debugger be compatible with? Possible values are `VS2012` , `VS2013` , `VS2015` and `VS2017` .
- `scm_type` - (Optional) The type of Source Control enabled for this App Service Slot. Defaults to `None` . Possible values are: `BitbucketGit` , `BitbucketHg` , `CodePlexGit` , `CodePlexHg` , `Dropbox` , `ExternalGit` , `ExternalHg` , `GitHub` , `LocalGit` , `None` , `OneDrive` , `Tfs` , `VS0` and `VSTRM`
- `use_32_bit_worker_process` - (Optional) Should the App Service Slot run in 32 bit mode, rather than 64 bit mode?

Note: Deployment Slots are not supported in the `Free` , `Shared` , or `Basic` App Service Plans.

- `virtual_network_name` - (Optional) The name of the Virtual Network which this App Service Slot should be attached to.
- `websockets_enabled` - (Optional) Should WebSockets be enabled?

A `cors` block supports the following:

- `allowed_origins` - (Optional) A list of origins which should be able to make cross-origin calls. `*` can be used to allow all calls.
 - `support_credentials` - (Optional) Are credentials supported?
-

A `auth_settings` block supports the following:

- `enabled` - (Required) Is Authentication enabled?
- `active_directory` - (Optional) A `active_directory` block as defined below.
- `additional_login_params` - (Optional) Login parameters to send to the OpenID Connect authorization endpoint when a user logs in. Each parameter must be in the form "key=value".
- `allowed_external_redirect_urls` - (Optional) External URLs that can be redirected to as part of logging in or logging out of the app.
- `default_provider` - (Optional) The default provider to use when multiple providers have been set up. Possible values are `AzureActiveDirectory`, `Facebook`, `Google`, `MicrosoftAccount` and `Twitter`.

NOTE: When using multiple providers, the default provider must be set for settings like `unauthenticated_client_action` to work.

- `facebook` - (Optional) A `facebook` block as defined below.
 - `google` - (Optional) A `google` block as defined below.
 - `issuer` - (Optional) Issuer URI. When using Azure Active Directory, this value is the URI of the directory tenant, e.g. `https://sts.windows.net/{tenant-guid}/` (`https://sts.windows.net/%7Btenant-guid%7D/`).
 - `microsoft` - (Optional) A `microsoft` block as defined below.
 - `runtime_version` - (Optional) The runtime version of the Authentication/Authorization module.
 - `token_refresh_extension_hours` - (Optional) The number of hours after session token expiration that a session token can be used to call the token refresh API. Defaults to 72.
 - `token_store_enabled` - (Optional) If enabled the module will durably store platform-specific security tokens that are obtained during login flows. Defaults to false.
 - `twitter` - (Optional) A `twitter` block as defined below.
 - `unauthenticated_client_action` - (Optional) The action to take when an unauthenticated client attempts to access the app. Possible values are `AllowAnonymous` and `RedirectToLoginPage`.
-

A `active_directory` block supports the following:

- `client_id` - (Required) The Client ID of this relying party application. Enables OpenIDConnection authentication with Azure Active Directory.
- `client_secret` - (Optional) The Client Secret of this relying party application. If no secret is provided, implicit flow will be used.

- `allowed_audiences` (Optional) Allowed audience values to consider when validating JWTs issued by Azure Active Directory.
-

A `facebook` block supports the following:

- `app_id` - (Required) The App ID of the Facebook app used for login
 - `app_secret` - (Required) The App Secret of the Facebook app used for Facebook Login.
 - `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Facebook Login authentication. <https://developers.facebook.com/docs/facebook-login> (<https://developers.facebook.com/docs/facebook-login>)
-

A `google` block supports the following:

- `client_id` - (Required) The OpenID Connect Client ID for the Google web application.
 - `client_secret` - (Required) The client secret associated with the Google web application.
 - `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Google Sign-In authentication. <https://developers.google.com/identity/sign-in/web/> (<https://developers.google.com/identity/sign-in/web/>)
-

A `ip_restriction` block supports the following:

- `ip_address` - (Optional) The IP Address used for this IP Restriction.
- `subnet_mask` - (Optional) The Subnet mask used for this IP Restriction. Defaults to `255.255.255.255`.
- `virtual_network_subnet_id` - (Optional) The Virtual Network Subnet ID used for this IP Restriction.

NOTE: One of either `ip_address` or `virtual_network_subnet_id` must be specified

A `microsoft` block supports the following:

- `client_id` - (Required) The OAuth 2.0 client ID that was created for the app used for authentication.
 - `client_secret` - (Required) The OAuth 2.0 client secret that was created for the app used for authentication.
 - `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Microsoft Account authentication. <https://msdn.microsoft.com/en-us/library/dn631845.aspx> (<https://msdn.microsoft.com/en-us/library/dn631845.aspx>)
-

A `identity` block supports the following:

- `type` - (Required) Specifies the identity type of the App Service. Possible values are `SystemAssigned` (where Azure will generate a Service Principal for you), `UserAssigned` where you can specify the Service Principal IDs in the `identity_ids` field, and `SystemAssigned, UserAssigned` which assigns both a system managed identity as well as the specified user assigned identities.

NOTE: When `type` is set to `SystemAssigned`, The assigned `principal_id` and `tenant_id` can be retrieved after the App Service has been created. More details are available below.

- `identity_ids` - (Optional) Specifies a list of user managed identity ids to be assigned. Required if `type` is `UserAssigned`.
-

A `logs` block supports the following:

- `application_logs` - (Optional) An `application_logs` block as defined below.
 - `http_logs` - (Optional) An `http_logs` block as defined below.
-

An `application_logs` block supports the following:

- `azure_blob_storage` - (Optional) An `azure_blob_storage` block as defined below.
-

An `http_logs` block supports *one* of the following:

- `file_system` - (Optional) A `file_system` block as defined below.
 - `azure_blob_storage` - (Optional) An `azure_blob_storage` block as defined below.
-

An `azure_blob_storage` block supports the following:

- `level` - (Required) The level at which to log. Possible values include `Error`, `Warning`, `Information`, `Verbose` and `Off`. **NOTE:** this field is not available for `http_logs`
 - `sas_url` - (Required) The URL to the storage container, with a Service SAS token appended. **NOTE:** there is currently no means of generating Service SAS tokens with the `azurerm` provider.
 - `retention_in_days` - (Required) The number of days to retain logs for.
-

A `file_system` block supports the following:

- `retention_in_days` - (Required) The number of days to retain logs for.
 - `retention_in_mb` - (Required) The maximum size in megabytes that http log files can use before being removed.
-

Elements of `ip_restriction` block (</docs/configuration/attr-as-blocks.html>) support:

- `ip_address` - (Required) The IP Address used for this IP Restriction.
- `subnet_mask` - (Optional) The Subnet mask used for this IP Restriction. Defaults to `255.255.255.255`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the App Service Slot.
- `default_site_hostname` - The Default Hostname associated with the App Service Slot - such as `mysite.azurewebsites.net`

- `site_credential` - A `site_credential` block as defined below, which contains the site-level credentials used to publish to this App Service.

`site_credential` exports the following:

- `username` - The username which can be used to publish to this App Service
- `password` - The password associated with the username, which can be used to publish to this App Service.

Import

App Service Slots can be imported using the `resource_id`, e.g.

```
terraform import azurerm_app_service_slot.instance1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/sites/website1/slots/instance1
```

azurerm_app_service_source_control_token

Manages an App Service source control token.

NOTE: Source Control Tokens are configured at the subscription level, not on each App Service - as such this can only be configured Subscription-wide

Example Usage

```
resource "azurerm_app_service_source_control_token" "test" {
  type = "GitHub"
  token = "7e57735e77e577e57"
}
```

Argument Reference

The following arguments are supported:

- `type` - (Required) The source control type. Possible values are `BitBucket`, `Dropbox`, `GitHub` and `OneDrive`.
- `token` - (Required) The OAuth access token.
- `token_secret` - (Optional) The OAuth access token secret.

Import

App Service source control tokens can be imported using the `type`, e.g.

```
terraform import azurerm_app_service_source_control_token.test GitHub
```


azurerm_automation_account

Manages a Automation Account.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name                = "automationAccount1"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku_name = "Basic"

  tags = {
    environment = "development"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Automation Account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Automation Account is created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Optional **Deprecated**) A `sku` block as described below.
- `sku_name` - (Optional) The SKU name of the account - only `Basic` is supported at this time.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The SKU name of the account - only `Basic` is supported at this time.

Attributes Reference

The following attributes are exported:

- `id` - The Automation Account ID.
- `dsc_server_endpoint` - The DSC Server Endpoint associated with this Automation Account.
- `dsc_primary_access_key` - The Primary Access Key for the DSC Endpoint associated with this Automation Account.
- `dsc_secondary_access_key` - The Secondary Access Key for the DSC Endpoint associated with this Automation Account.

Import

Automation Accounts can be imported using the `resource id`, e.g.

```
terraform import azurerm_automation_account.account1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/group1/providers/Microsoft.Automation/automationAccounts/account1
```

azurerm_automation_credential

Manages a Automation Credential.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name            = "account1"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_credential" "example" {
  name                = "credential1"
  resource_group_name = "${azurerm_resource_group.example.name}"
  account_name        = "${azurerm_automation_account.example.name}"
  username             = "example_user"
  password             = "example_pwd"
  description          = "This is an example credential"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Credential. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Credential is created. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the automation account in which the Credential is created. Changing this forces a new resource to be created.
- `username` - (Required) The username associated with this Automation Credential.
- `password` - (Required) The password associated with this Automation Credential.
- `description` - (Optional) The description associated with this Automation Credential.

Attributes Reference

The following attributes are exported:

- `id` - The Automation Credential ID.

Import

Automation Credentials can be imported using the `resource id`, e.g.

```
terraform import azurerm_automation_credential.credential1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Automation/automationAccounts/account1/credentials/credential1
```

azurerm_automation_dsc_configuration

Manages a Automation DSC Configuration.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name            = "account1"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_dsc_configuration" "example" {
  name                = "test"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  location            = "${azurerm_resource_group.example.location}"
  content_embedded    = "configuration test {}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the DSC Configuration. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the DSC Configuration is created. Changing this forces a new resource to be created.
- **automation_account_name** - (Required) The name of the automation account in which the DSC Configuration is created. Changing this forces a new resource to be created.
- **content_embedded** - (Required) The PowerShell DSC Configuration script.
- **location** - (Required) Must be the same location as the Automation Account.
- **log_verbose** - (Optional) Verbose log option.
- **description** - (Optional) Description to go with DSC Configuration.

Attributes Reference

The following attributes are exported:

- `id` - The DSC Configuration ID.

azurerm_automation_dsc_nodeconfiguration

Manages a Automation DSC Node Configuration.

Example Usage

```

resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name                = "account1"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_dsc_configuration" "example" {
  name                = "test"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  location            = "${azurerm_resource_group.example.location}"
  content_embedded    = "configuration test {}"
}

resource "azurerm_automation_dsc_nodeconfiguration" "example" {
  name                = "test.localhost"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  depends_on          = ["azurerm_automation_dsc_configuration.example"]

  content_embedded = <<mofcontent
instance of MSFT_FileDirectoryConfiguration as $MSFT_FileDirectoryConfiguration1ref
{
  ResourceID = "[File]bla";
  Ensure = "Present";
  Contents = "bogus Content";
  DestinationPath = "c:\\\\bogus.txt";
  ModuleName = "PSDesiredStateConfiguration";
  SourceInfo = "::-3::9::file";
  ModuleVersion = "1.0";
  ConfigurationName = "bla";
};
instance of OMI_ConfigurationDocument
{
  Version="2.0.0";
  MinimumCompatibleVersion = "1.0.0";
  CompatibleVersionAdditionalProperties= {"Omi_BaseResource:ConfigurationName"};
  Author="bogusAuthor";
  GenerationDate="06/15/2018 14:06:24";
  GenerationHost="bogusComputer";
  Name="test";
};
mofcontent
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the DSC Node Configuration. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the DSC Node Configuration is created. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the DSC Node Configuration is created. Changing this forces a new resource to be created.
- `content_embedded` - (Required) The PowerShell DSC Node Configuration (mof content).

Attributes Reference

The following attributes are exported:

- `id` - The DSC Node Configuration ID.

azurerm_automation_job_schedule

Links an Automation Runbook and Schedule.

Example Usage

This is an example of just the Job Schedule. A full example of the `azurerm_automation_job_schedule` resource can be found in the `./examples/automation-account` directory within the Github Repository ([https://github.com/terraform-provider-azurerm/tree/master/examples/automation-account](https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/automation-account))

```
resource "azurerm_automation_job_schedule" "example" {
  resource_group_name      = "tf-rgr-automation"
  automation_account_name = "tf-automation-account"
  schedule_name            = "hour"
  runbook_name             = "Get-VirtualMachine"

  parameters = {
    resourcegroup = "tf-rgr-vm"
    vmname        = "TF-VM-01"
  }
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which the Job Schedule is created. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the Automation Account in which the Job Schedule is created. Changing this forces a new resource to be created.
- `runbook_name` - (Required) The name of a Runbook to link to a Schedule. It needs to be in the same Automation Account as the Schedule and Job Schedule. Changing this forces a new resource to be created.
- `parameters` - (Optional) A map of key/value pairs corresponding to the arguments that can be passed to the Runbook. Changing this forces a new resource to be created.

NOTE: The parameter keys/names must strictly be in lowercase, even if this is not the case in the runbook. This is due to a limitation in Azure Automation where the parameter names are normalized. The values specified don't have this limitation.

- `run_on` - (Optional) Name of a Hybrid Worker Group the Runbook will be executed on. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Automation Job Schedule's full ID.
- `job_schedule_id` - The UUID identifying the Automation Job Schedule.

Import

Automation Job Schedule can be imported using the `resource id`, e.g.

```
terraform import azurerm_automation_job_schedule.example /subscriptions/00000000-0000-0000-0000-000000000000/000/resourceGroups/group1/providers/Microsoft.Automation/automationAccounts/account1/jobSchedules/10000000-1001-1001-1001-000000000001
```

azurerm_automation_module

Manages a Automation Module.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name                = "account1"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_module" "example" {
  name                = "xActiveDirectory"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"

  module_link {
    uri = "https://devopsgallerystorage.blob.core.windows.net/packages/xactivedirectory.2.19.0.nupkg"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Module. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Module is created. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the Module is created. Changing this forces a new resource to be created.
- `module_link` - (Required) The published Module link.

`module_link` supports the following:

- `uri` - (Required) The uri of the module content (zip or nupkg).

Attributes Reference

The following attributes are exported:

- `id` - The Automation Module ID.

azurerm_automation_runbook

Manages a Automation Runbook.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name                = "account1"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_runbook" "example" {
  name                = "Get-AzureVMTutorial"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  account_name        = "${azurerm_automation_account.example.name}"
  log_verbose         = "true"
  log_progress        = "true"
  description         = "This is an example runbook"
  runbook_type        = "PowerShellWorkflow"

  publish_content_link {
    uri = "https://raw.githubusercontent.com/Azure/azure-quickstart-templates/c4935ffb69246a6058eb24f54640f53f69d3ac9f/101-automation-runbook-getvms/Runbooks/Get-AzureVMTutorial.ps1"
  }
}
```

Example Usage - custom content

```

resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name            = "account1"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

data "local_file" "example" {
  filename = "${path.module}/example.ps1"
}

resource "azurerm_automation_runbook" "example" {
  name            = "Get-AzureVMTutorial"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  account_name    = "${azurerm_automation_account.example.name}"
  log_verbose     = "true"
  log_progress    = "true"
  description     = "This is an example runbook"
  runbook_type    = "PowerShell"

  publish_content_link {
    uri = "https://raw.githubusercontent.com/Azure/azure-quickstart-templates/c4935ffb69246a6058eb24f54640f53f69d3ac9f/101-automation-runbook-getvms/Runbooks/Get-AzureVMTutorial.ps1"
  }

  content = "${data.local_file.example.content}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Runbook. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Runbook is created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the automation account in which the Runbook is created. Changing this forces a new resource to be created.
- `runbook_type` - (Required) The type of the runbook - can be either `Graph`, `GraphPowerShell`, `GraphPowerShellWorkflow`, `PowerShellWorkflow`, `PowerShell` or `Script`.

- `log_progress` - (Required) Progress log option.
- `log_verbose` - (Required) Verbose log option.
- `publish_content_link` - (Required) The published runbook content link.
- `description` - (Optional) A description for this credential.
- `content` - (Optional) The desired content of the runbook.

NOTE The Azure API requires a `publish_content_link` to be supplied even when specifying your own `content` .

NOTE Setting `content` to an empty string will revert the runbook to the `publish_content_link` .

- `tags` - (Optional) A mapping of tags to assign to the resource.

`publish_content_link` supports the following:

- `uri` - (Required) The uri of the runbook content.

Attributes Reference

The following attributes are exported:

- `id` - The Automation Runbook ID.

Import

Automation Runbooks can be imported using the `resource id` , e.g.

```
terraform import azurerm_automation_runbook.Get-AzureVMTutorial /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Automation/automationAccounts/account1/runbooks/Get-AzureVMTutorial
```


azurerm_automation_schedule

Manages a Automation Schedule.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-automation-account"
  location  = "West Europe"
}

resource "azurerm_automation_account" "example" {
  name                = "tfex-automation-account"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_schedule" "example" {
  name                = "tfex-automation-schedule"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  frequency           = "Week"
  interval            = 1
  timezone            = "Central Europe Standard Time"
  start_time          = "2014-04-15T18:00:15+02:00"
  description         = "This is an example schedule"
  week_days           = ["Friday"]
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Schedule. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the Schedule is created. Changing this forces a new resource to be created.
- **automation_account_name** - (Required) The name of the automation account in which the Schedule is created. Changing this forces a new resource to be created.
- **frequency** - (Required) The frequency of the schedule. - can be either `OneTime`, `Day`, `Hour`, `Week`, or `Month`.
- **description** - (Optional) A description for this Schedule.
- **interval** - (Optional) The number of `frequency` s between runs. Only valid when frequency is `Day`, `Hour`, `Week`, or `Month` and defaults to `1`.

- `start_time` - (Optional) Start time of the schedule. Must be at least five minutes in the future. Defaults to seven minutes in the future from the time the resource is created.
- `expiry_time` - (Optional) The end time of the schedule.
- `timezone` - (Optional) The timezone of the start time. Defaults to `UTC` . For possible values see: [https://msdn.microsoft.com/en-us/library/ms912391\(v=winembedded.11\).aspx](https://msdn.microsoft.com/en-us/library/ms912391(v=winembedded.11).aspx) ([https://msdn.microsoft.com/en-us/library/ms912391\(v=winembedded.11\).aspx](https://msdn.microsoft.com/en-us/library/ms912391(v=winembedded.11).aspx))
- `week_days` - (Optional) List of days of the week that the job should execute on. Only valid when frequency is `Week` .
- `month_days` - (Optional) List of days of the month that the job should execute on. Must be between `1` and `31` . `-1` for last day of the month. Only valid when frequency is `Month` .
- `monthly_occurrence` - (Optional) List of occurrences of days within a month. Only valid when frequency is `Month` . The `monthly_occurrence` block supports fields documented below.

The `monthly_occurrence` block supports:

- `day` - (Required) Day of the occurrence. Must be one of `Monday` , `Tuesday` , `Wednesday` , `Thursday` , `Friday` , `Saturday` , `Sunday` .
- `occurrence` - (Required) Occurrence of the week within the month. Must be between `1` and `5` . `-1` for last week within the month.

Attributes Reference

The following attributes are exported:

- `id` - The Automation Schedule ID.

Import

Automation Schedule can be imported using the `resource_id` , e.g.

```
terraform import azurerm_automation_schedule.schedule1 /subscriptions/00000000-0000-0000-0000-000000000000/0/resourceGroups/group1/providers/Microsoft.Automation/automationAccounts/account1/schedules/schedule1
```

azurerm_automation_variable_bool

Manages a boolean variable in Azure Automation

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-example-rg"
  location  = "West US"
}

resource "azurerm_automation_account" "example" {
  name                = "tfex-example-account"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_variable_bool" "example" {
  name                = "tfex-example-var"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  value               = false
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Automation Variable. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the Variable is created. Changing this forces a new resource to be created.
- `description` - (Optional) The description of the Automation Variable.
- `encrypted` - (Optional) Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - (Optional) The value of the Automation Variable as a `boolean`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.

Import

Automation Bool Variable can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_automation_variable_bool.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/tfex-example-rg/providers/Microsoft.Automation/automationAccounts/tfex-example-account/variables/tfex-example-var
```

azurerm_automation_variable_datetime

Manages a date/time variable in Azure Automation

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-example-rg"
  location  = "West US"
}

resource "azurerm_automation_account" "example" {
  name                = "tfex-example-account"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_variable_datetime" "example" {
  name                = "tfex-example-var"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  value              = "2019-04-24T21:40:54.074Z"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Automation Variable. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the Variable is created. Changing this forces a new resource to be created.
- `description` - (Optional) The description of the Automation Variable.
- `encrypted` - (Optional) Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - (Optional) The value of the Automation Variable in the RFC3339 Section 5.6 Internet Date/Time Format (<https://tools.ietf.org/html/rfc3339#section-5.6>).

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.

Import

Automation Datetime Variable can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_automation_variable_datetime.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/tfex-example-rg/providers/Microsoft.Automation/automationAccounts/tfex-example-account/variables/tfex-example-var
```

azurerm_automation_variable_int

Manages a integer variable in Azure Automation

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-example-rg"
  location  = "West US"
}

resource "azurerm_automation_account" "example" {
  name                = "tfex-example-account"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_variable_int" "example" {
  name                = "tfex-example-var"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  value               = 1234
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Automation Variable. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the Variable is created. Changing this forces a new resource to be created.
- `description` - (Optional) The description of the Automation Variable.
- `encrypted` - (Optional) Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - (Optional) The value of the Automation Variable as a `integer`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.

Import

Automation Int Variable can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_automation_variable_int.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/tfex-example-rg/providers/Microsoft.Automation/automationAccounts/tfex-example-account/variables/tfex-example-var
```


azurerm_automation_variable_string

Manages a string variable in Azure Automation

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-example-rg"
  location  = "West US"
}

resource "azurerm_automation_account" "example" {
  name                = "tfex-example-account"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name = "Basic"
  }
}

resource "azurerm_automation_variable_string" "example" {
  name                = "tfex-example-var"
  resource_group_name = "${azurerm_resource_group.example.name}"
  automation_account_name = "${azurerm_automation_account.example.name}"
  value              = "Hello, Terraform Basic Test."
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Automation Variable. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Automation Variable. Changing this forces a new resource to be created.
- `automation_account_name` - (Required) The name of the automation account in which the Variable is created. Changing this forces a new resource to be created.
- `description` - (Optional) The description of the Automation Variable.
- `encrypted` - (Optional) Specifies if the Automation Variable is encrypted. Defaults to `false`.
- `value` - (Optional) The value of the Automation Variable as a `string`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Automation Variable.

Import

Automation String Variable can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_automation_variable_string.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/tfex-example-rg/providers/Microsoft.Automation/automationAccounts/tfex-example-account/variables/tfex-example-var
```

azurerm_autoscale_setting

Manages an AutoScale Setting which can be applied to Virtual Machine Scale Sets, App Services and other scalable resources.

NOTE: This resource has been deprecated in favour of the `azurerm_monitor_autoscale_setting` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "autoscalingTest"
  location  = "West US"
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_autoscale_setting" "test" {
  name                        = "myAutoscaleSetting"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  target_resource_id         = "${azurerm_virtual_machine_scale_set.test.id}"

  profile {
    name = "defaultProfile"

    capacity {
      default = 1
      minimum = 1
      maximum = 10
    }

    rule {
      metric_trigger {
        metric_name      = "Percentage CPU"
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
        time_grain       = "PT1M"
        statistic         = "Average"
        time_window       = "PT5M"
        time_aggregation = "Average"
        operator          = "GreaterThan"
        threshold         = 75
      }

      scale_action {
        direction = "Increase"
        type      = "ChangeCount"
        value     = "1"
        cooldown  = "PT1M"
      }
    }
  }
}
```

```

        cooldown = PT1M
    }
}

rule {
    metric_trigger {
        metric_name      = "Percentage CPU"
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
        time_grain       = "PT1M"
        statistic         = "Average"
        time_window       = "PT5M"
        time_aggregation  = "Average"
        operator          = "LessThan"
        threshold         = 25
    }

    scale_action {
        direction = "Decrease"
        type      = "ChangeCount"
        value     = "1"
        cooldown  = "PT1M"
    }
}

notification {
    email {
        send_to_subscription_administrator = true
        send_to_subscription_co_administrator = true
        custom_emails                      = ["admin@contoso.com"]
    }
}
}

```

Example Usage (repeating on weekends)

```

resource "azurerm_resource_group" "test" {
    name     = "autoscalingTest"
    location = "West US"
}

resource "azurerm_virtual_machine_scale_set" "test" {
    # ...
}

resource "azurerm_autoscale_setting" "test" {
    name                = "myAutoscaleSetting"
    resource_group_name = "${azurerm_resource_group.test.name}"
    location             = "${azurerm_resource_group.test.location}"
    target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"

    profile {
        name = "Weekends"

        capacity {
            default = 1
        }
    }
}

```



```
}  
}
```

Example Usage (for fixed dates)

```
resource "azurerm_resource_group" "test" {  
  name      = "autoscalingTest"  
  location = "West US"  
}  
  
resource "azurerm_virtual_machine_scale_set" "test" {  
  # ...  
}  
  
resource "azurerm_autoscale_setting" "test" {  
  name                = "myAutoscaleSetting"  
  enabled             = true  
  resource_group_name = "${azurerm_resource_group.test.name}"  
  location            = "${azurerm_resource_group.test.location}"  
  target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"  
  
  profile {  
    name = "forJuly"  
  
    capacity {  
      default = 1  
      minimum = 1  
      maximum = 10  
    }  
  
    rule {  
      metric_trigger {  
        metric_name      = "Percentage CPU"  
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"  
        time_grain       = "PT1M"  
        statistic        = "Average"  
        time_window      = "PT5M"  
        time_aggregation = "Average"  
        operator         = "GreaterThan"  
        threshold        = 90  
      }  
  
      scale_action {  
        direction = "Increase"  
        type      = "ChangeCount"  
        value     = "2"  
        cooldown  = "PT1M"  
      }  
    }  
  
    rule {  
      metric_trigger {  
        metric_name      = "Percentage CPU"  
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"  
        time_grain       = "PT1M"  
        statistic        = "Average"
```

```

    time_window      = "PT5M"
    time_aggregation = "Average"
    operator         = "LessThan"
    threshold        = 10
  }

  scale_action {
    direction = "Decrease"
    type      = "ChangeCount"
    value     = "2"
    cooldown  = "PT1M"
  }
}

fixed_date {
  timezone = "Pacific Standard Time"
  start    = "2020-07-01T00:00:00Z"
  end      = "2020-07-31T23:59:59Z"
}

notification {
  email {
    send_to_subscription_administrator = true
    send_to_subscription_co_administrator = true
    custom_emails                      = ["admin@contoso.com"]
  }
}
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the AutoScale Setting. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the Resource Group in the AutoScale Setting should be created. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the AutoScale Setting should exist. Changing this forces a new resource to be created.
- **profile** - (Required) Specifies one or more (up to 20) `profile` blocks as defined below.
- **target_resource_id** - (Required) Specifies the resource ID of the resource that the autoscale setting should be added to.
- **enabled** - (Optional) Specifies whether automatic scaling is enabled for the target resource. Defaults to `true`.
- **notification** - (Optional) Specifies a `notification` block as defined below.
- **tags** - (Optional) A mapping of tags to assign to the resource.

A `profile` block supports the following:

- `name` - (Required) Specifies the name of the profile.
- `capacity` - (Required) A `capacity` block as defined below.
- `rule` - (Required) One or more (up to 10) `rule` blocks as defined below.
- `fixed_date` - (Optional) A `fixed_date` block as defined below. This cannot be specified if a `recurrence` block is specified.
- `recurrence` - (Optional) A `recurrence` block as defined below. This cannot be specified if a `fixed_date` block is specified.

A `capacity` block supports the following:

- `default` - (Required) The number of instances that are available for scaling if metrics are not available for evaluation. The default is only used if the current instance count is lower than the default. Valid values are between 0 and 1000 .
- `maximum` - (Required) The maximum number of instances for this resource. Valid values are between 0 and 1000 .

NOTE: The maximum number of instances is also limited by the amount of Cores available in the subscription.

- `minimum` - (Required) The minimum number of instances for this resource. Valid values are between 0 and 1000 .

A `rule` block supports the following:

- `metric_trigger` - (Required) A `metric_trigger` block as defined below.
- `scale_action` - (Required) A `scale_action` block as defined below.

A `metric_trigger` block supports the following:

- `metric_name` - (Required) The name of the metric that defines what the rule monitors, such as `Percentage CPU` for `Virtual Machine Scale Sets` and `CpuPercentage` for `App Service Plan` .

NOTE: The allowed value of `metric_name` highly depends on the targeting resource type, please visit [Supported metrics with Azure Monitor \(https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported\)](https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported) for more details.

- `metric_resource_id` - (Required) The ID of the Resource which the Rule monitors.
- `operator` - (Required) Specifies the operator used to compare the metric data and threshold. Possible values are: `Equals` , `NotEquals` , `GreaterThan` , `GreaterThanOrEqual` , `LessThan` , `LessThanOrEqual` .
- `statistic` - (Required) Specifies how the metrics from multiple instances are combined. Possible values are `Average` , `Min` and `Max` .
- `time_aggregation` - (Required) Specifies how the data that's collected should be combined over time. Possible values include `Average` , `Count` , `Maximum` , `Minimum` , `Last` and `Total` . Defaults to `Average` .
- `time_grain` - (Required) Specifies the granularity of metrics that the rule monitors, which must be one of the pre-defined values returned from the metric definitions for the metric. This value must be between 1 minute and 12 hours and be formatted as an ISO 8601 string.

- `time_window` - (Required) Specifies the time range for which data is collected, which must be greater than the delay in metric collection (which varies from resource to resource). This value must be between 5 minutes and 12 hours and be formatted as an ISO 8601 string.
 - `threshold` - (Required) Specifies the threshold of the metric that triggers the scale action.
-

A `scale_action` block supports the following:

- `cooldown` - (Required) The amount of time to wait since the last scaling action before this action occurs. Must be between 1 minute and 1 week and formatted as a ISO 8601 string.
 - `direction` - (Required) The scale direction. Possible values are `Increase` and `Decrease`.
 - `type` - (Required) The type of action that should occur. Possible values are `ChangeCount`, `ExactCount` and `PercentChangeCount`.
 - `value` - (Required) The number of instances involved in the scaling action. Defaults to `1`.
-

A `fixed_date` block supports the following:

- `end` - (Required) Specifies the end date for the profile, formatted as an RFC3339 date string.
 - `start` - (Required) Specifies the start date for the profile, formatted as an RFC3339 date string.
 - `timezone` (Optional) The Time Zone of the `start` and `end` times. A list of possible values can be found here (<https://msdn.microsoft.com/en-us/library/azure/dn931928.aspx>). Defaults to `UTC`.
-

A `recurrence` block supports the following:

- `timezone` - (Required) The Time Zone used for the `hours` field. A list of possible values can be found here (<https://msdn.microsoft.com/en-us/library/azure/dn931928.aspx>). Defaults to `UTC`.
 - `days` - (Required) A list of days that this profile takes effect on. Possible values include `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday` and `Sunday`.
 - `hours` - (Required) A list containing a single item, which specifies the Hour interval at which this recurrence should be triggered (in 24-hour time). Possible values are from `0` to `23`.
 - `minutes` - (Required) A list containing a single item which specifies the Minute interval at which this recurrence should be triggered.
-

A `notification` block supports the following:

- `email` - (Required) A `email` block as defined below.
 - `webhook` - (Optional) One or more `webhook` blocks as defined below.
-

A `email` block supports the following:

- `send_to_subscription_administrator` - (Optional) Should email notifications be sent to the subscription administrator? Defaults to `false`.
- `send_to_subscription_co_administrator` - (Optional) Should email notifications be sent to the subscription co-

administrator? Defaults to `false`.

- `custom_emails` - (Optional) Specifies a list of custom email addresses to which the email notifications will be sent.

A `webhook` block supports the following:

- `service_uri` - (Required) The HTTPS URI which should receive scale notifications.
- `properties` - (Optional) A map of settings.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the AutoScale Setting.

Import

AutoScale Setting can be imported using the `resource id`, e.g.

```
terraform import azurerm_autoscale_setting.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/microsoft.insights/autoscalesettings/setting1
```

azurerm_availability_set

Manages an availability set for virtual machines.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_availability_set" "test" {
  name                = "acceptanceTestAvailabilitySet1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the availability set. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the availability set. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `platform_update_domain_count` - (Optional) Specifies the number of update domains that are used. Defaults to 5.

NOTE: The number of Update Domains varies depending on which Azure Region you're using - a list can be found here (<https://github.com/MicrosoftDocs/azure-docs/blob/master/includes/managed-disks-common-fault-domain-region-list.md>).

- `platform_fault_domain_count` - (Optional) Specifies the number of fault domains that are used. Defaults to 3.

NOTE: The number of Fault Domains varies depending on which Azure Region you're using - a list can be found here (<https://github.com/MicrosoftDocs/azure-docs/blob/master/includes/managed-disks-common-fault-domain-region-list.md>).

- `proximity_placement_group_id` - (Optional) The ID of the Proximity Placement Group to which this Virtual Machine should be assigned. Changing this forces a new resource to be created

- `managed` - (Optional) Specifies whether the availability set is managed or not. Possible values are `true` (to specify aligned) or `false` (to specify classic). Default is `false`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The virtual Availability Set ID.

Import

Availability Sets can be imported using the `resource id`, e.g.

```
terraform import azurerm_availability_set.group1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Compute/availabilitySets/webAvailSet
```

azurerm_azuread_application

Manages an Application within Azure Active Directory.

NOTE: The Azure Active Directory resources have been split out into a new AzureAD Provider (<http://terraform.io/docs/providers/azuread/index.html>) - as such the AzureAD resources within the AzureRM Provider are deprecated and will be removed in the next major version (2.0). Information on how to migrate from the existing resources to the new AzureAD Provider can be found here (</docs/providers/azurerm/guides/migrating-to-azuread.html>).

NOTE: If you're authenticating using a Service Principal then it must have permissions to both `Read` and `write` all applications and `Sign in and read user profile` within the legacy Windows Azure Active Directory API.

Example Usage

```
resource "azurerm_azuread_application" "test" {
  name                = "example"
  homepage            = "https://homepage"
  identifier_uris     = ["https://uri"]
  reply_urls         = ["https://replyurl"]
  available_to_other_tenants = false
  oauth2_allow_implicit_flow = true
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The display name for the application.
- `homepage` - (optional) The URL to the application's home page. If no homepage is specified this defaults to `https://{name}`.
- `identifier_uris` - (Optional) A list of user-defined URI(s) that uniquely identify a Web application within it's Azure AD tenant, or within a verified custom domain if the application is multi-tenant.
- `reply_urls` - (Optional) A list of URLs that user tokens are sent to for sign in, or the redirect URIs that OAuth 2.0 authorization codes and access tokens are sent to.
- `available_to_other_tenants` - (Optional) Is this Azure AD Application available to other tenants? Defaults to `false`.
- `oauth2_allow_implicit_flow` - (Optional) Does this Azure AD Application allow OAuth2.0 implicit flow tokens? Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `application_id` - The Application ID.

Import

Azure Active Directory Applications can be imported using the `object_id`, e.g.

```
terraform import azurerm_azuread_application.test 00000000-0000-0000-0000-000000000000
```

azurerm_azuread_service_principal

Manages a Service Principal associated with an Application within Azure Active Directory.

NOTE: The Azure Active Directory resources have been split out into a new AzureAD Provider (<http://terraform.io/docs/providers/azuread/index.html>) - as such the AzureAD resources within the AzureRM Provider are deprecated and will be removed in the next major version (2.0). Information on how to migrate from the existing resources to the new AzureAD Provider can be found here (</docs/providers/azurerm/guides/migrating-to-azuread.html>).

NOTE: If you're authenticating using a Service Principal then it must have permissions to both `Read` and `write` all applications and `Sign in and read user profile` within the Windows Azure Active Directory API.

Example Usage

```
resource "azurerm_azuread_application" "test" {
  name                = "example"
  homepage            = "http://homepage"
  identifier_uris     = ["http://uri"]
  reply_urls         = ["http://replyurl"]
  available_to_other_tenants = false
  oauth2_allow_implicit_flow = true
}

resource "azurerm_azuread_service_principal" "test" {
  application_id = "${azurerm_azuread_application.test.application_id}"
}
```

Argument Reference

The following arguments are supported:

- `application_id` - (Required) The ID of the Azure AD Application for which to create a Service Principal.

Attributes Reference

The following attributes are exported:

- `id` - The Object ID for the Service Principal.
- `display_name` - The Display Name of the Azure Active Directory Application associated with this Service Principal.

Import

Azure Active Directory Service Principals can be imported using the `object_id`, e.g.

```
terraform import azurerm_azuread_service_principal.test 00000000-0000-0000-0000-000000000000
```


azurerm_azuread_service_principal_password

Manages a Password associated with a Service Principal within Azure Active Directory.

NOTE: The Azure Active Directory resources have been split out into a new AzureAD Provider (<http://terraform.io/docs/providers/azuread/index.html>) - as such the AzureAD resources within the AzureRM Provider are deprecated and will be removed in the next major version (2.0). Information on how to migrate from the existing resources to the new AzureAD Provider can be found here (</docs/providers/azurerm/guides/migrating-to-azuread.html>).

NOTE: If you're authenticating using a Service Principal then it must have permissions to both `Read` and `write` all applications and `Sign in and read user profile` within the Windows Azure Active Directory API.

Example Usage

```
resource "azurerm_azuread_application" "test" {
  name                = "example"
  homepage             = "https://homepage"
  identifier_uris      = ["https://uri"]
  reply_urls           = ["https://replyurl"]
  available_to_other_tenants = false
  oauth2_allow_implicit_flow = true
}

resource "azurerm_azuread_service_principal" "test" {
  application_id = "${azurerm_azuread_application.test.application_id}"
}

resource "azurerm_azuread_service_principal_password" "test" {
  service_principal_id = "${azurerm_azuread_service_principal.test.id}"
  value                = "VT=uSgbTanZhyz@%nL9Hpd+Tfay_MRV#"
  end_date             = "2020-01-01T01:02:03Z"
}
```

Argument Reference

The following arguments are supported:

- `service_principal_id` - (Required) The ID of the Service Principal for which this password should be created. Changing this field forces a new resource to be created.
- `value` - (Required) The Password for this Service Principal.
- `end_date` - (Required) The End Date which the Password is valid until, formatted as a RFC3339 date string (e.g. `2018-01-01T01:02:03Z`). Changing this field forces a new resource to be created.

- `key_id` - (Optional) A GUID used to uniquely identify this Key. If not specified a GUID will be created. Changing this field forces a new resource to be created.
- `start_date` - (Optional) The Start Date which the Password is valid from, formatted as a RFC3339 date string (e.g. `2018-01-01T01:02:03Z`). If this isn't specified, the current date is used. Changing this field forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Key ID for the Service Principal Password.

Import

Service Principal Passwords can be imported using the `object_id`, e.g.

```
terraform import azurerm_azuread_service_principal_password.test 00000000-0000-0000-0000-000000000000/11111111-1111-1111-1111-111111111111
```

NOTE: This ID format is unique to Terraform and is composed of the Service Principal's Object ID and the Service Principal Password's Key ID in the format `{ServicePrincipalObjectId}/{ServicePrincipalPasswordKeyId}` .

azurerm_bastion_host

Manages a Bastion Host Instance.

Note: Bastion Host Instances are a preview feature in Azure, and therefore are only supported in a select number of regions. Read more (<https://docs.microsoft.com/en-us/azure/bastion/bastion-faq>).

Example Usage

This example deploys an Azure Bastion Host Instance to a target virtual network.

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "example" {
  name            = "examplevnet"
  address_space   = ["192.168.1.0/24"]
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_subnet" "example" {
  name                 = "AzureBastionSubnet"
  resource_group_name  = "${azurerm_resource_group.example.name}"
  virtual_network_name = "${azurerm_virtual_network.example.name}"
  address_prefix        = "192.168.1.224/27"
}

resource "azurerm_public_ip" "example" {
  name                = "examplepip"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  allocation_method   = "Static"
  sku                 = "Standard"
}

resource "azurerm_bastion_host" "example" {
  name            = "examplebastion"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  ip_configuration {
    name                = "configuration"
    subnet_id           = "${azurerm_subnet.example.id}"
    public_ip_address_id = "${azurerm_public_ip.example.id}"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Bastion Host. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Bastion Host.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `ip_configuration` - (Required) A `ip_configuration` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `ip_configuration` block supports the following:

- `name` - (Required) The name of the IP configuration.
- `subnet_id` - (Required) Reference to a subnet in which this Bastion Host has been created.
- `public_ip_address_id` (Required) Reference to a Public IP Address to associate with this Bastion Host.

Attributes Reference

The following attributes are exported:

- `dns_name` - The FQDN for the Azure Bastion Host.

Import

Bastion Hosts can be imported using the `resource_id`, e.g.

```
terraform import azurerm_bastion_host.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/bastionHosts/instance1
```

azurerm_batch_account

Manages an Azure Batch account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "testbatch"
  location  = "westeurope"
}

resource "azurerm_storage_account" "test" {
  name                        = "teststorage"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_batch_account" "test" {
  name                      = "testbatchaccount"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  location                  = "${azurerm_resource_group.test.location}"
  pool_allocation_mode      = "BatchService"
  storage_account_id        = "${azurerm_storage_account.test.id}"

  tags = {
    env = "test"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Batch account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Batch account. Changing this forces a new resource to be created.

NOTE: To work around a bug in the Azure API (<https://github.com/Azure/azure-rest-api-specs/issues/5574>) this property is currently treated as case-insensitive. A future version of Terraform will require that the casing is correct.

- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `pool_allocation_mode` - (Optional) Specifies the mode to use for pool allocation. Possible values are `BatchService` or `UserSubscription`. Defaults to `BatchService`.

NOTE: When using `UserSubscription` mode, an Azure KeyVault reference has to be specified. See `key_vault_reference` below.

NOTE: When using `UserSubscription` mode, the Microsoft Azure Batch service principal has to have `Contributor` role on your subscription scope, as documented here (<https://docs.microsoft.com/en-us/azure/batch/batch-account-create-portal#additional-configuration-for-user-subscription-mode>).

- `key_vault_reference` - (Optional) A `key_vault_reference` block that describes the Azure KeyVault reference to use when deploying the Azure Batch account using the `UserSubscription` pool allocation mode.
- `storage_account_id` - (Optional) Specifies the storage account to use for the Batch account. If not specified, Azure Batch will manage the storage.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `key_vault_reference` block supports the following:

- `id` - (Required) The Azure identifier of the Azure KeyVault to use.
- `url` - (Required) The HTTPS URL of the Azure KeyVault to use.

Attributes Reference

The following attributes are exported:

- `id` - The Batch account ID.
- `primary_access_key` - The Batch account primary access key.
- `secondary_access_key` - The Batch account secondary access key.
- `account_endpoint` - The account endpoint used to interact with the Batch service.

NOTE: Primary and secondary access keys are only available when `pool_allocation_mode` is set to `BatchService`. See documentation (<https://docs.microsoft.com/en-us/azure/batch/batch-api-basics>) for more information.

azurerm_batch_application

Manages Azure Batch Application instance.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example-rg"
  location = "West US"
}

resource "azurerm_storage_account" "example" {
  name                     = "examplesa"
  resource_group_name      = "${azurerm_resource_group.example.name}"
  location                 = "${azurerm_resource_group.example.location}"
  account_tier             = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_batch_account" "example" {
  name                = "exampleba"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"
  pool_allocation_mode = "BatchService"
  storage_account_id  = "${azurerm_storage_account.example.id}"
}

resource "azurerm_batch_application" "example" {
  name                = "example-batch-application"
  resource_group_name = "${azurerm_resource_group.example.name}"
  account_name        = "${azurerm_batch_account.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the application. This must be unique within the account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group that contains the Batch account. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the Batch account. Changing this forces a new resource to be created.
- `allow_updates` - (Optional) A value indicating whether packages within the application may be overwritten using the same version string. Defaults to `true`.
- `default_version` - (Optional) The package to use if a client requests the application but does not specify a version. This property can only be set to the name of an existing package.

- `display_name` - (Optional) The display name for the application.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the resource.

Import

Batch Application can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_batch_application.example /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/example-rg/providers/Microsoft.Batch/batchAccounts/exampleba/applications/example-batch-ap  
plication
```


azurerm_batch_certificate

Manages a certificate in an Azure Batch account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "testbatch"
  location  = "westeurope"
}

resource "azurerm_storage_account" "test" {
  name                        = "teststorage"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_batch_account" "test" {
  name                      = "testbatchaccount"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  location                  = "${azurerm_resource_group.test.location}"
  pool_allocation_mode      = "BatchService"
  storage_account_id        = "${azurerm_storage_account.test.id}"

  tags = {
    env = "test"
  }
}

resource "azurerm_batch_certificate" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  account_name        = "${azurerm_batch_account.test.name}"
  certificate          = "${filebase64("certificate.pfx")}"
  format              = "Pfx"
  password            = "terraform"
  thumbprint          = "42C107874FD0E4A9583292A2F1098E8FE4B2EDDA"
  thumbprint_algorithm = "SHA1"
}
```

Argument Reference

The following arguments are supported:

- `account_name` - (Required) Specifies the name of the Batch account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Batch account. Changing this forces a new resource to be created.

- `certificate` - (Required) The base64-encoded contents of the certificate.
- `format` - (Required) The format of the certificate. Possible values are `Cer` or `Pfx`.
- `password` - (Optional) The password to access the certificate's private key. This must and can only be specified when `format` is `Pfx`.
- `thumbprint` - (Required) The thumbprint of the certificate. At this time the only supported value is 'SHA1'.

Attributes Reference

The following attributes are exported:

- `id` - The Batch certificate ID.
- `name` - The generated name of the certificate.
- `public_data` - The public key of the certificate.

azurerm_batch_pool

Manages an Azure Batch pool.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "testaccbatch"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "testaccsa"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_batch_account" "test" {
  name                = "testaccbatch"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  pool_allocation_mode = "BatchService"
  storage_account_id  = "${azurerm_storage_account.test.id}"

  tags = {
    env = "test"
  }
}

resource "azurerm_batch_certificate" "testcer" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  account_name        = "${azurerm_batch_account.test.name}"
  certificate          = "${filebase64("certificate.cer")}"
  format              = "Cer"
  thumbprint          = "312d31a79fa0cef49c00f769afc2b73e9f4edf34"
  thumbprint_algorithm = "SHA1"
}

resource "azurerm_batch_pool" "test" {
  name                = "testaccpool"
  resource_group_name = "${azurerm_resource_group.test.name}"
  account_name        = "${azurerm_batch_account.test.name}"
  display_name        = "Test Acc Pool Auto"
  vm_size             = "Standard_A1"
  node_agent_sku_id   = "batch.node.ubuntu 16.04"

  auto_scale {
    evaluation_interval = "PT15M"

    formula = <<EOF
      startingNumberOfVMs = 1;
      maxNumberOfVMs = 25;
      pendingTaskSamplePercent = $PendingTasks.GetSamplePercent(180 * TimeInterval_Second);
    >>EOF
  }
}
```

```

        pendingTaskSamples = pendingTaskSamplePercent < 70 ? startingNumberOfVMs : avg($PendingTasks.GetSam
ple(180 *   TimeInterval_Second));
        $TargetDedicatedNodes=min(maxNumberOfVMs, pendingTaskSamples);
EOF
    }

    storage_image_reference {
        publisher = "microsoft-azure-batch"
        offer      = "ubuntu-server-container"
        sku        = "16-04-lts"
        version    = "latest"
    }

    container_configuration {
        type = "DockerCompatible"
        container_registries = [
            {
                registry_server = "docker.io"
                user_name       = "login"
                password        = "apassword"
            },
        ]
    }

    start_task {
        command_line      = "echo 'Hello World from $env'"
        max_task_retry_count = 1
        wait_for_success   = true

        environment = {
            env = "TEST"
        }

        user_identity {
            auto_user {
                elevation_level = "NonAdmin"
                scope           = "Task"
            }
        }
    }

    certificate {
        id          = "${azurerm_batch_certificate.testcer.id}"
        visibility = ["StartTask"]
    }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Batch pool. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Batch pool. Changing this forces a new resource to be created.

NOTE: To work around a bug in the Azure API (<https://github.com/Azure/azure-rest-api-specs/issues/5574>) this property is currently treated as case-insensitive. A future version of Terraform will require that the casing is correct.

- `account_name` - (Required) Specifies the name of the Batch account in which the pool will be created. Changing this forces a new resource to be created.
- `node_agent_sku_id` - (Required) Specifies the Sku of the node agents that will be created in the Batch pool.
- `vm_size` - (Required) Specifies the size of the VM created in the Batch pool.
- `storage_image_reference` - (Required) A `storage_image_reference` block for the virtual machines that will compose the Batch pool.
- `display_name` - (Optional) Specifies the display name of the Batch pool.
- `max_tasks_per_node` - (Optional) Specifies the maximum number of tasks that can run concurrently on a single compute node in the pool. Defaults to 1. Changing this forces a new resource to be created.
- `fixed_scale` - (Optional) A `fixed_scale` block that describes the scale settings when using fixed scale.
- `auto_scale` - (Optional) A `auto_scale` block that describes the scale settings when using auto scale.
- `start_task` - (Optional) A `start_task` block that describes the start task settings for the Batch pool.
- `certificate` - (Optional) One or more `certificate` blocks that describe the certificates to be installed on each compute node in the pool.
- `container_configuration` - (Optional) The container configuration used in the pool's VMs.

NOTE: For Windows compute nodes, the Batch service installs the certificates to the specified certificate store and location. For Linux compute nodes, the certificates are stored in a directory inside the task working directory and an environment variable `AZ_BATCH_CERTIFICATES_DIR` is supplied to the task to query for this location. For certificates with visibility of `remoteUser`, a `certs` directory is created in the user's home directory (e.g., `/home/{user-name}/certs`) and certificates are placed in that directory.

Please Note: `fixed_scale` and `auto_scale` blocks cannot be used both at the same time.

A `storage_image_reference` block supports the following:

This block provisions virtual machines in the Batch Pool from one of two sources: an Azure Platform Image (e.g. Ubuntu/Windows Server) or a Custom Image.

To provision from an Azure Platform Image, the following fields are applicable:

- `publisher` - (Required) Specifies the publisher of the image used to create the virtual machines. Changing this forces a new resource to be created.
- `offer` - (Required) Specifies the offer of the image used to create the virtual machines. Changing this forces a new resource to be created.
- `sku` - (Required) Specifies the SKU of the image used to create the virtual machines. Changing this forces a new resource to be created.

- `version` - (Optional) Specifies the version of the image used to create the virtual machines. Changing this forces a new resource to be created.

To provision a Custom Image, the following fields are applicable:

* `id` - (Required) Specifies the ID of the Custom Image which the virtual machines should be created from. Changing this forces a new resource to be created. See official documentation (<https://docs.microsoft.com/en-us/azure/batch/batch-custom-images>) for more details.

A `fixed_scale` block supports the following:

- `target_dedicated_nodes` - (Optional) The number of nodes in the Batch pool. Defaults to `1`.
 - `target_low_priority_nodes` - (Optional) The number of low priority nodes in the Batch pool. Defaults to `0`.
 - `resize_timeout` - (Optional) The timeout for resize operations. Defaults to `PT15M`.
-

A `auto_scale` block supports the following:

- `evaluation_interval` - (Optional) The interval to wait before evaluating if the pool needs to be scaled. Defaults to `PT15M`.
 - `formula` - (Required) The autoscale formula that needs to be used for scaling the Batch pool.
-

A `start_task` block supports the following:

- `command_line` - (Required) The command line executed by the start task.
 - `max_task_retry_count` - (Optional) The number of retry count. Defaults to `1`.
 - `wait_for_success` - (Optional) A flag that indicates if the Batch pool should wait for the start task to be completed. Default to `false`.
 - `environment` - (Optional) A map of strings (key,value) that represents the environment variables to set in the start task.
 - `user_identity` - (Required) A `user_identity` block that describes the user identity under which the start task runs.
 - `resource_file` - (Optional) One or more `resource_file` blocks that describe the files to be downloaded to a compute node.
-

A `user_identity` block supports the following:

- `user_name` - (Optional) The username to be used by the Batch pool start task.
- `auto_user` - (Optional) A `auto_user` block that describes the user identity under which the start task runs.

Please Note: `user_name` and `auto_user` blocks cannot be used both at the same time, but you need to define one or the other.

A `auto_user` block supports the following:

- `elevation_level` - (Optional) The elevation level of the user identity under which the start task runs. Possible values are `Admin` or `NonAdmin`. Defaults to `NonAdmin`.
- `scope` - (Optional) The scope of the user identity under which the start task runs. Possible values are `Task` or `Pool`. Defaults to `Task`.

A `certificate` block supports the following:

- `id` - (Required) The ID of the Batch Certificate to install on the Batch Pool, which must be inside the same Batch Account.
- `store_location` - (Required) The location of the certificate store on the compute node into which to install the certificate. Possible values are `CurrentUser` or `LocalMachine`.

NOTE: This property is applicable only for pools configured with Windows nodes (that is, created with `cloudServiceConfiguration`, or with `virtualMachineConfiguration` using a Windows image reference). For Linux compute nodes, the certificates are stored in a directory inside the task working directory and an environment variable `AZ_BATCH_CERTIFICATES_DIR` is supplied to the task to query for this location. For certificates with visibility of `remoteUser`, a 'certs' directory is created in the user's home directory (e.g., `/home/{user-name}/certs`) and certificates are placed in that directory.

- `store_name` - (Optional) The name of the certificate store on the compute node into which to install the certificate. This property is applicable only for pools configured with Windows nodes (that is, created with `cloudServiceConfiguration`, or with `virtualMachineConfiguration` using a Windows image reference). Common store names include: `My`, `Root`, `CA`, `Trust`, `Disallowed`, `TrustedPeople`, `TrustedPublisher`, `AuthRoot`, `AddressBook`, but any custom store name can also be used. The default value is `My`.
- `visibility` - (Optional) Which user accounts on the compute node should have access to the private data of the certificate.

A `container_configuration` block supports the following:

- `type` - (Optional) The type of container configuration. Possible value is `DockerCompatible`.
- `container_registries` - (Optional) Additional container registries from which container images can be pulled by the pool's VMs.

A `resource_file` block supports the following:

- `auto_storage_container_name` - (Optional) The storage container name in the auto storage account.

- `blob_prefix` - (Optional) The blob prefix to use when downloading blobs from an Azure Storage container. Only the blobs whose names begin with the specified prefix will be downloaded. The property is valid only when `auto_storage_container_name` or `storage_container_url` is used. This prefix can be a partial filename or a subdirectory. If a prefix is not specified, all the files in the container will be downloaded.
- `file_mode` - (Optional) The file permission mode represented as a string in octal format (e.g. `"0644"`). This property applies only to files being downloaded to Linux compute nodes. It will be ignored if it is specified for a `resource_file` which will be downloaded to a Windows node. If this property is not specified for a Linux node, then a default value of `0770` is applied to the file.
- `file_path` - (Optional) The location on the compute node to which to download the file, relative to the task's working directory. If the `http_url` property is specified, the `file_path` is required and describes the path which the file will be downloaded to, including the filename. Otherwise, if the `auto_storage_container_name` or `storage_container_url` property is specified, `file_path` is optional and is the directory to download the files to. In the case where `file_path` is used as a directory, any directory structure already associated with the input data will be retained in full and appended to the specified `filePath` directory. The specified relative path cannot break out of the task's working directory (for example by using `'..'`).
- `http_url` - (Optional) The URL of the file to download. If the URL is Azure Blob Storage, it must be readable using anonymous access; that is, the Batch service does not present any credentials when downloading the blob. There are two ways to get such a URL for a blob in Azure storage: include a Shared Access Signature (SAS) granting read permissions on the blob, or set the ACL for the blob or its container to allow public access.
- `storage_container_url` - (Optional) The URL of the blob container within Azure Blob Storage. This URL must be readable and listable using anonymous access; that is, the Batch service does not present any credentials when downloading the blob. There are two ways to get such a URL for a blob in Azure storage: include a Shared Access Signature (SAS) granting read and list permissions on the blob, or set the ACL for the blob or its container to allow public access.

Please Note: Exactly one of `auto_storage_container_name`, `storage_container_url` and `auto_user` must be specified.

A `container_registries` block supports the following:

- `registry_server` - (Optional) The container registry URL. The default is `"docker.io"`. Changing this forces a new resource to be created.
- `user_name` - (Optional) The user name to log into the registry server. Changing this forces a new resource to be created.
- `password` - (Optional) The password to log into the registry server. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Batch pool ID.

azurerm_bot_connection

Manages a Email integration for a Bot Channel

Note A bot can only have a single Email Channel associated with it.

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "example" {
  name       = "example"
  location   = "northeurope"
}

resource "azurerm_bot_channels_registration" "example" {
  name                = "example"
  location            = "global"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "F0"
  microsoft_app_id    = "${data.azurerm_client_config.current.service_principal_application_id}"
}

resource "azurerm_bot_channel_email" "example" {
  bot_name            = "${azurerm_bot_channels_registration.example.name}"
  location            = "${azurerm_bot_channels_registration.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  client_id           = "exampleId"
  client_secret       = "exampleSecret"
  verification_token  = "exampleVerificationToken"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which to create the Bot Channel. Changing this forces a new resource to be created.
- `location` - (Required) The supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `bot_name` - (Required) The name of the Bot Resource this channel will be associated with. Changing this forces a new resource to be created.
- `email_address` - (Required) The email address that the Bot will authenticate with.
- `email_password` - (Required) The email password that the the Bot will authenticate with.

Attributes Reference

The following attributes are exported:

- `id` - The Bot Channel ID.

Import

The Email Channel for a Bot can be imported using the `resource id`, e.g.

```
terraform import azurerm_bot_channel_email.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.BotService/botServices/example/channels/EmailChannel
```

azurerm_bot_connection

Manages a Slack integration for a Bot Channel

Note A bot can only have a single Slack Channel associated with it.

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "example" {
  name       = "example"
  location   = "northeurope"
}

resource "azurerm_bot_channels_registration" "example" {
  name                = "example"
  location            = "global"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "F0"
  microsoft_app_id    = "${data.azurerm_client_config.current.service_principal_application_id}"
}

resource "azurerm_bot_channel_slack" "example" {
  bot_name           = "${azurerm_bot_channels_registration.example.name}"
  location           = "${azurerm_bot_channels_registration.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  client_id          = "exampleId"
  client_secret       = "exampleSecret"
  verification_token  = "exampleVerificationToken"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which to create the Bot Channel. Changing this forces a new resource to be created.
- `location` - (Required) The supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `bot_name` - (Required) The name of the Bot Resource this channel will be associated with. Changing this forces a new resource to be created.
- `client_id` - (Required) The Client ID that will be used to authenticate with Slack.
- `client_secret` - (Required) The Client Secret that will be used to authenticate with Slack.
- `verification_token` - (Required) The Verification Token that will be used to authenticate with Slack.

- `landing_page_url` - (Optional) The Slack Landing Page URL.

Attributes Reference

The following attributes are exported:

- `id` - The Bot Channel ID.

Import

The Slack Channel for a Bot can be imported using the `resource id`, e.g.

```
terraform import azurerm_bot_channel_slack.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.BotService/botServices/example/channels/SlackChannel
```

azurerm_bot_channels_registration

Manages a Bot Channels Registration.

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "example" {
  name     = "example"
  location = "northeurope"
}

resource "azurerm_bot_channels_registration" "example" {
  name                = "example"
  location            = "global"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "F0"
  microsoft_app_id    = "${data.azurerm_client_config.current.service_principal_application_id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Bot Channels Registration. Changing this forces a new resource to be created. Must be globally unique.
- `resource_group_name` - (Required) The name of the resource group in which to create the Bot Channels Registration. Changing this forces a new resource to be created.
- `location` - (Required) The supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) The SKU of the Bot Channels Registration. Valid values include `F0` or `S1` . Changing this forces a new resource to be created.
- `microsoft_app_id` - (Required) The Microsoft Application ID for the Bot Channels Registration. Changing this forces a new resource to be created.
- `display_name` - (Optional) The name of the Bot Channels Registration will be displayed as. This defaults to `name` if not specified.
- `endpoint` - (Optional) The Bot Channels Registration endpoint.
- `developer_app_insights_key` - (Optional) The Application Insights Key to associate with the Bot Channels Registration.
- `developer_app_insights_api_key` - (Optional) The Application Insights API Key to associate with the Bot Channels Registration.

- `developer_app_insights_application_id` - (Optional) The Application Insights Application ID to associate with the Bot Channels Registration.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Bot Channels Registration ID.

Import

Bot Channels Registration can be imported using the `resource id`, e.g.

```
terraform import azurerm_bot_channels_registration.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.BotService/botServices/example
```

azurerm_bot_connection

Manages a Bot Connection.

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "example" {
  name     = "example"
  location = "northeurope"
}

resource "azurerm_bot_channels_registration" "example" {
  name                = "example"
  location            = "global"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "F0"
  microsoft_app_id    = "${data.azurerm_client_config.current.service_principal_application_id}"
}

resource "azurerm_bot_connection" "example" {
  name                = "example"
  bot_name            = "${azurerm_bot_channels_registration.example.name}"
  location            = "${azurerm_bot_channels_registration.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  service_provider_name = "box"
  client_id           = "exampleId"
  client_secret       = "exampleSecret"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Bot Connection. Changing this forces a new resource to be created. Must be globally unique.
- **resource_group_name** - (Required) The name of the resource group in which to create the Bot Connection. Changing this forces a new resource to be created.
- **location** - (Required) The supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **bot_name** - (Required) The name of the Bot Resource this connection will be associated with. Changing this forces a new resource to be created.
- **service_provider_name** - (Required) The name of the service provider that will be associated with this connection. Changing this forces a new resource to be created.
- **client_id** - (Required) The Client ID that will be used to authenticate with the service provider.

- `client_secret` - (Required) The Client Secret that will be used to authenticate with the service provider.
- `scopes` - (Optional) The Scopes at which the connection should be applied.
- `parameters` - (Optional) A map of additional parameters to apply to the connection.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Bot Connection ID.

Import

Bot Connection can be imported using the `resource id`, e.g.

```
terraform import azurerm_bot_connection.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.BotService/botServices/example/connections/example
```


azurerm_bot_web_app

Manages a Web App Bot.

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "example" {
  name     = "example"
  location = "northeurope"
}

resource "azurerm_bot_web_app" "example" {
  name                = "example"
  location            = "global"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "F0"
  microsoft_app_id    = "${data.azurerm_client_config.current.service_principal_application_id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Web App Bot. Changing this forces a new resource to be created. Must be globally unique.
- `resource_group_name` - (Required) The name of the resource group in which to create the Web App Bot. Changing this forces a new resource to be created.
- `location` - (Required) The supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) The SKU of the Web App Bot. Valid values include `F0` or `S1` . Changing this forces a new resource to be created.
- `microsoft_app_id` - (Required) The Microsoft Application ID for the Web App Bot. Changing this forces a new resource to be created.
- `display_name` - (Optional) The name of the Web App Bot will be displayed as. This defaults to `name` if not specified.
- `endpoint` - (Optional) The Web App Bot endpoint.
- `developer_app_insights_key` - (Optional) The Application Insights Key to associate with the Web App Bot.
- `developer_app_insights_api_key` - (Optional) The Application Insights API Key to associate with the Web App Bot.
- `developer_app_insights_application_id` - (Optional) The Application Insights Application ID to associate with the Web App Bot.

- `luis_app_ids` - (Optional) A list of LUIS App IDs to associate with the Web App Bot.
- `luis_key` - (Optional) The LUIS key to associate with the Web App Bot.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Web App Bot ID.

Import

Web App Bots can be imported using the `resource id`, e.g.

```
terraform import azurerm_bot_web_app.example /subscriptions/00000000-0000-0000-0000-000000000000/resource
Groups/example/providers/Microsoft.BotService/botServices/example
```

azurerm_cdn_endpoint

A CDN Endpoint is the entity within a CDN Profile containing configuration information regarding caching behaviors and origins. The CDN Endpoint is exposed using the URL format `.azureedge.net`.

Example Usage

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_cdn_profile" "test" {
  name                = "exampleCdnProfile"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Standard_Verizon"
}

resource "azurerm_cdn_endpoint" "test" {
  name                = "${random_id.server.hex}"
  profile_name        = "${azurerm_cdn_profile.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  origin {
    name     = "exampleCdnOrigin"
    host_name = "www.example.com"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the CDN Endpoint. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the CDN Endpoint.
- `profile_name` - (Required) The CDN Profile to which to attach the CDN Endpoint.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.

- `is_http_allowed` - (Optional) Defaults to `true` .
- `is_https_allowed` - (Optional) Defaults to `true` .
- `content_types_to_compress` - (Optional) An array of strings that indicates a content types on which compression will be applied. The value for the elements should be MIME types.
- `geo_filter` - (Optional) A set of Geo Filters for this CDN Endpoint. Each `geo_filter` block supports fields documented below.
- `is_compression_enabled` - (Optional) Indicates whether compression is to be enabled. Defaults to `false`.
- `querystring_caching_behaviour` - (Optional) Sets query string caching behavior. Allowed values are `IgnoreQueryString` , `BypassCaching` and `UseQueryString` . Defaults to `IgnoreQueryString` .
- `optimization_type` - (Optional) What types of optimization should this CDN Endpoint optimize for? Possible values include `DynamicSiteAcceleration` , `GeneralMediaStreaming` , `GeneralWebDelivery` , `LargeFileDownload` and `VideoOnDemandMediaStreaming` .
- `origin` - (Optional) The set of origins of the CDN endpoint. When multiple origins exist, the first origin will be used as primary and rest will be used as failover options. Each `origin` block supports fields documented below.
- `origin_host_header` - (Optional) The host header CDN provider will send along with content requests to origins. Defaults to the host name of the origin.
- `origin_path` - (Optional) The path used at for origin requests.
- `probe_path` - (Optional) the path to a file hosted on the origin which helps accelerate delivery of the dynamic content and calculate the most optimal routes for the CDN. This is relative to the `origin_path` .
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `origin` block supports:

- `name` - (Required) The name of the origin. This is an arbitrary value. However, this value needs to be unique under the endpoint. Changing this forces a new resource to be created.
- `host_name` - (Required) A string that determines the hostname/IP address of the origin server. This string can be a domain name, Storage Account endpoint, Web App endpoint, IPv4 address or IPv6 address. Changing this forces a new resource to be created.
- `http_port` - (Optional) The HTTP port of the origin. Defaults to `80` . Changing this forces a new resource to be created.
- `https_port` - (Optional) The HTTPS port of the origin. Defaults to `443` . Changing this forces a new resource to be created.

The `geo_filter` block supports:

- `relative_path` - (Required) The relative path applicable to geo filter.
- `action` - (Required) The Action of the Geo Filter. Possible values include `Allow` and `Block` .
- `country_codes` - (Required) A List of two letter country codes (e.g. `US` , `GB`) to be associated with this Geo Filter.

Attributes Reference

The following attributes are exported:

- `id` - The CDN Endpoint ID.

Import

CDN Endpoints can be imported using the `resource id`, e.g.

```
terraform import azurerm_cdn_endpoint.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Cdn/profiles/myprofile1/endpoints/myendpoint1
```

azurerm_cdn_profile

Manages a CDN Profile to create a collection of CDN Endpoints.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_cdn_profile" "test" {
  name                = "exampleCdnProfile"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Standard_Verizon"

  tags = {
    environment = "Production"
    cost_center = "MSFT"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the CDN Profile. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the CDN Profile.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) The pricing related information of current CDN profile. Accepted values are `Standard_Akamai`, `Standard_ChinaCdn`, `Standard_Microsoft`, `Standard_Verizon` or `Premium_Verizon`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The CDN Profile ID.

Import

CDN Profiles can be imported using the `resource_id`, e.g.

```
terraform import azurerm_cdn_profile.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Cdn/profiles/myprofile1
```

azurerm_cognitive_account

Manages a Cognitive Services Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_cognitive_account" "test" {
  name                = "example-account"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  kind                = "Face"

  sku {
    name = "S0"
    tier  = "Standard"
  }

  tags = {
    Acceptance = "Test"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cognitive Service Account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cognitive Service Account is created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `kind` - (Required) Specifies the type of Cognitive Service Account that should be created. Possible values are `Academic`, `Bing.Autosuggest`, `Bing.Autosuggest.v7`, `Bing.CustomSearch`, `Bing.Search`, `Bing.Search.v7`, `Bing.Speech`, `Bing.SpellCheck`, `Bing.SpellCheck.v7`, `CognitiveServices`, `ComputerVision`, `ContentModerator`, `CustomSpeech`, `CustomVision.Prediction`, `CustomVision.Training`, `Emotion`, `Face`, `LUIS`, `QnAMaker`, `Recommendations`, `SpeakerRecognition`, `Speech`, `SpeechServices`, `SpeechTranslation`, `TextAnalytics`, `TextTranslation` and `WebLM`. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) Specifies the Name of the Sku. Possible values are `F0` , `F1` , `S0` , `S1` , `S2` , `S3` , `S4` , `S5` , `S6` , `P0` , `P1` and `P2` .
- `tier` - (Required) Specifies the Tier of the Sku. Possible values include `Free` , `Standard` and `Premium` .

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Cognitive Service Account.
- `endpoint` - The endpoint used to connect to the Cognitive Service Account.
- `primary_access_key` - A primary access key which can be used to connect to the Cognitive Service Account.
- `secondary_access_key` - The secondary access key which can be used to connect to the Cognitive Service Account.

Import

Cognitive Service Accounts can be imported using the `resource id` , e.g.

```
terraform import azurerm_cognitive_account.account1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.CognitiveServices/accounts/account1
```

azurerm_connection_monitor

Configures a Connection Monitor to monitor communication between a Virtual Machine and an endpoint using a Network Watcher.

NOTE: This resource has been deprecated in favour of the `azurerm_network_connection_monitor` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "connection-monitor-rg"
  location = "West US"
}

resource "azurerm_network_watcher" "test" {
  name                = "network-watcher"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_virtual_network" "test" {
  name                = "production-network"
  address_space       = ["10.0.0.0/16"]
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "cmtest-nic"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id           = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "cmtest-vm"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
```

```

resource_group_name = ${azurerm_resource_group.test.name}
network_interface_ids = ["${azurerm_network_interface.test.id}"]
vm_size              = "Standard_F2"

storage_image_reference {
  publisher = "Canonical"
  offer     = "UbuntuServer"
  sku       = "16.04-LTS"
  version   = "latest"
}

storage_os_disk {
  name          = "osdisk"
  caching       = "ReadWrite"
  create_option = "FromImage"
  managed_disk_type = "Standard_LRS"
}

os_profile {
  computer_name = "cmtest-vm"
  admin_username = "testadmin"
  admin_password = "Password1234!"
}

os_profile_linux_config {
  disable_password_authentication = false
}
}

resource "azurerm_virtual_machine_extension" "test" {
  name                = "cmtest-vm-network-watcher"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_machine_name = "${azurerm_virtual_machine.test.name}"
  publisher            = "Microsoft.Azure.NetworkWatcher"
  type                 = "NetworkWatcherAgentLinux"
  type_handler_version = "1.4"
  auto_upgrade_minor_version = true
}

resource "azurerm_connection_monitor" "test" {
  name                = "cmtest-connectionmonitor"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_watcher_name = "${azurerm_network_watcher.test.name}"

  source {
    virtual_machine_id = "${azurerm_virtual_machine.test.id}"
  }

  destination {
    address = "terraform.io"
    port    = 80
  }

  depends_on = ["azurerm_virtual_machine_extension.test"]
}

```

NOTE: This Resource requires that the Network Watcher Agent Virtual Machine Extension (<https://docs.microsoft.com/en-us/azure/network-watcher/connection-monitor>) is installed on the Virtual Machine before monitoring can be started. The extension can be installed via the `azurerm_virtual_machine_extension` resource ([/docs/providers/azurerm/r/virtual_machine_extension.html](/docs/providers/azurerm/r/virtual_machine_extension.html)).

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Connection Monitor. Changing this forces a new resource to be created.
- `network_watcher_name` - (Required) The name of the Network Watcher. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Connection Monitor. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `auto_start` - (Optional) Specifies whether the connection monitor will start automatically once created. Defaults to `true`. Changing this forces a new resource to be created.
- `interval_in_seconds` - (Optional) Monitoring interval in seconds. Defaults to `60`.
- `source` - (Required) A `source` block as defined below.
- `destination` - (Required) A `destination` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `source` block contains:

- `virtual_machine_id` - (Required) The ID of the Virtual Machine to monitor connectivity from.
- `port` - (Optional) The port on the Virtual Machine to monitor connectivity from. Defaults to `0` (Dynamic Port Assignment).

A `destination` block contains:

- `virtual_machine_id` - (Optional) The ID of the Virtual Machine to monitor connectivity to.
- `address` - (Optional) IP address or domain name to monitor connectivity to.
- `port` - (Required) The port on the destination to monitor connectivity to.

NOTE: One of `virtual_machine_id` or `address` must be specified.

Attributes Reference

The following attributes are exported:

- `id` - The Connection Monitor ID.

Import

Connection Monitors can be imported using the `resource id`, e.g.

```
terraform import azurerm_connection_monitor.monitor1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.Network/networkWatchers/watcher1/connectionMonitors/monitor1
```

azurerm_container_group

Manages as an Azure Container Group instance.

Example Usage

This example provisions a Basic Container. Other examples of the `azurerm_container_group` resource can be found in the `./examples/container-instance` directory within the Github Repository ([https://github.com/terraform-provider-azurerm/tree/master/examples/container-instance](https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/container-instance)).

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_container_group" "example" {
  name                = "example-continst"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  ip_address_type     = "public"
  dns_name_label      = "aci-label"
  os_type             = "Linux"

  container {
    name   = "hello-world"
    image  = "microsoft/aci-helloworld:latest"
    cpu    = "0.5"
    memory = "1.5"

    ports {
      port      = 443
      protocol  = "TCP"
    }
  }

  container {
    name   = "sidecar"
    image  = "microsoft/aci-tutorial-sidecar"
    cpu    = "0.5"
    memory = "1.5"
  }

  tags = {
    environment = "testing"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Container Group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Container Group. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `identity` - (Optional) An `identity` block as defined below.

Note: managed identities are not supported for containers in virtual networks.

- `container` - (Required) The definition of a container that is part of the group as documented in the `container` block below. Changing this forces a new resource to be created.
- `os_type` - (Required) The OS for the container group. Allowed values are `Linux` and `Windows`. Changing this forces a new resource to be created.

Note: if `os_type` is set to `Windows` currently only a single `container` block is supported. Windows containers are not supported in virtual networks.

- `diagnostics` - (Optional) A `diagnostics` block as documented below.
- `dns_name_label` - (Optional) The DNS label/name for the container groups IP. Changing this forces a new resource to be created.

Note: DNS label/name is not supported when deploying to virtual networks.

- `ip_address_type` - (Optional) Specifies the ip address type of the container. `Public` or `Private`. Changing this forces a new resource to be created. If set to `Private`, `network_profile_id` also needs to be set.

Note: `dns_name_label`, `identity` and `os_type` set to `windows` are not compatible with `Private ip_address_type`

- `network_profile_id` - (Optional) Network profile ID for deploying to virtual network.
- `image_registry_credential` - (Optional) A `image_registry_credential` block as documented below. Changing this forces a new resource to be created.
- `restart_policy` - (Optional) Restart policy for the container group. Allowed values are `Always`, `Never`, `OnFailure`. Defaults to `Always`. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource. Changing this forces a new resource to be created.

An `identity` block supports the following:

- `type` - (Required) The Managed Service Identity Type of this container group. Possible values are `SystemAssigned` (where Azure will generate a Service Principal for you), `UserAssigned` where you can specify the Service Principal IDs in the `identity_ids` field, and `SystemAssigned, UserAssigned` which assigns both a system managed identity as

well as the specified user assigned identities. Changing this forces a new resource to be created.

NOTE: When `type` is set to `SystemAssigned`, identity the Principal ID can be retrieved after the container group has been created. See documentation (<https://docs.microsoft.com/en-us/azure/active-directory/managed-service-identity/overview>) for more information.

- `identity_ids` - (Optional) Specifies a list of user managed identity ids to be assigned. Required if `type` is `UserAssigned`. Changing this forces a new resource to be created.

A `container` block supports:

- `name` - (Required) Specifies the name of the Container. Changing this forces a new resource to be created.
- `image` - (Required) The container image name. Changing this forces a new resource to be created.
- `cpu` - (Required) The required number of CPU cores of the containers. Changing this forces a new resource to be created.
- `memory` - (Required) The required memory of the containers in GB. Changing this forces a new resource to be created.
- `gpu` - (Optional) A `gpu` block as defined below. Changing this forces a new resource to be created.

Note: Gpu resources are currently only supported in Linux containers.

- `ports` - (Optional) A set of public ports for the container. Changing this forces a new resource to be created. Set as documented in the `ports` block below.
- `environment_variables` - (Optional) A list of environment variables to be set on the container. Specified as a map of name/value pairs. Changing this forces a new resource to be created.
- `secure_environment_variables` - (Optional) A list of sensitive environment variables to be set on the container. Specified as a map of name/value pairs. Changing this forces a new resource to be created.
- `readiness_probe` - (Optional) The definition of a readiness probe for this container as documented in the `readiness_probe` block below. Changing this forces a new resource to be created.
- `liveness_probe` - (Optional) The definition of a readiness probe for this container as documented in the `liveness_probe` block below. Changing this forces a new resource to be created.
- `command` - (Optional) A command line to be run on the container. Changing this forces a new resource to be created.

NOTE: The field `command` has been deprecated in favor of `commands` to better match the API.

- `commands` - (Optional) A list of commands which should be run on the container. Changing this forces a new resource to be created.
- `volume` - (Optional) The definition of a volume mount for this container as documented in the `volume` block below. Changing this forces a new resource to be created.

A `diagnostics` block supports:

- `log_analytics` - (Required) A `log_analytics` block as defined below. Changing this forces a new resource to be

created.

A `image_registry_credential` block supports:

- `username` - (Required) The username with which to connect to the registry. Changing this forces a new resource to be created.
 - `password` - (Required) The password with which to connect to the registry. Changing this forces a new resource to be created.
 - `server` - (Required) The address to use to connect to the registry without protocol ("https"/"http"). For example: "myacr.acr.io". Changing this forces a new resource to be created.
-

A `log_analytics` block supports:

- `log_type` - (Optional) The log type which should be used. Possible values are `ContainerInsights` and `ContainerInstanceLogs`. Changing this forces a new resource to be created.
 - `workspace_id` - (Required) The Workspace ID of the Log Analytics Workspace. Changing this forces a new resource to be created.
 - `workspace_key` - (Required) The Workspace Key of the Log Analytics Workspace. Changing this forces a new resource to be created.
 - `metadata` - (Optional) Any metadata required for Log Analytics. Changing this forces a new resource to be created.
-

A `ports` block supports:

- `port` - (Required) The port number the container will expose. Changing this forces a new resource to be created.
- `protocol` - (Required) The network protocol associated with port. Possible values are `TCP` & `UDP`. Changing this forces a new resource to be created.

--

A `gpu` block supports:

- `count` - (Required) The number of GPUs which should be assigned to this container. Allowed values are `1`, `2`, or `4`. Changing this forces a new resource to be created.
 - `sku` - (Required) The Sku which should be used for the GPU. Possible values are `K80`, `P100`, or `V100`. Changing this forces a new resource to be created.
-

A `volume` block supports:

- `name` - (Required) The name of the volume mount. Changing this forces a new resource to be created.
- `mount_path` - (Required) The path on which this volume is to be mounted. Changing this forces a new resource to be created.
- `read_only` - (Optional) Specify if the volume is to be mounted as read only or not. The default value is `false`. Changing this forces a new resource to be created.

- `storage_account_name` - (Required) The Azure storage account from which the volume is to be mounted. Changing this forces a new resource to be created.
 - `storage_account_key` - (Required) The access key for the Azure Storage account specified as above. Changing this forces a new resource to be created.
 - `share_name` - (Required) The Azure storage share that is to be mounted as a volume. This must be created on the storage account specified as above. Changing this forces a new resource to be created.
-

The `readiness_probe` block supports:

- `exec` - (Optional) Commands to be run to validate container readiness. Changing this forces a new resource to be created.
 - `httpget` - (Optional) The definition of the `httpget` for this container as documented in the `httpget` block below. Changing this forces a new resource to be created.
 - `initial_delay_seconds` - (Optional) Number of seconds after the container has started before liveness or readiness probes are initiated. Changing this forces a new resource to be created.
 - `period_seconds` - (Optional) How often (in seconds) to perform the probe. The default value is `10` and the minimum value is `1`. Changing this forces a new resource to be created.
 - `failure_threshold` - (Optional) How many times to try the probe before restarting the container (liveness probe) or marking the container as unhealthy (readiness probe). The default value is `3` and the minimum value is `1`. Changing this forces a new resource to be created.
 - `success_threshold` - (Optional) Minimum consecutive successes for the probe to be considered successful after having failed. The default value is `1` and the minimum value is `1`. Changing this forces a new resource to be created.
 - `timeout_seconds` - (Optional) Number of seconds after which the probe times out. The default value is `1` and the minimum value is `1`. Changing this forces a new resource to be created.
-

The `liveness_probe` block supports:

- `exec` - (Optional) Commands to be run to validate container readiness. Changing this forces a new resource to be created.
- `httpget` - (Optional) The definition of the `httpget` for this container as documented in the `httpget` block below. Changing this forces a new resource to be created.
- `initial_delay_seconds` - (Optional) Number of seconds after the container has started before liveness or readiness probes are initiated. Changing this forces a new resource to be created.
- `period_seconds` - (Optional) How often (in seconds) to perform the probe. The default value is `10` and the minimum value is `1`. Changing this forces a new resource to be created.
- `failure_threshold` - (Optional) How many times to try the probe before restarting the container (liveness probe) or marking the container as unhealthy (readiness probe). The default value is `3` and the minimum value is `1`. Changing this forces a new resource to be created.
- `success_threshold` - (Optional) Minimum consecutive successes for the probe to be considered successful after having failed. The default value is `1` and the minimum value is `1`. Changing this forces a new resource to be created.

- `timeout_seconds` - (Optional) Number of seconds after which the probe times out. The default value is `1` and the minimum value is `1`. Changing this forces a new resource to be created.

The `httpget` block supports:

- `path` - (Optional) Path to access on the HTTP server. Changing this forces a new resource to be created.
- `port` - (Optional) Number of the port to access on the container. Changing this forces a new resource to be created.
- `scheme` - (Optional) Scheme to use for connecting to the host. Possible values are `Http` and `Https`. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The container group ID.
- `ip_address` - The IP address allocated to the container group.
- `fqdn` - The FQDN of the container group derived from `dns_name_label`.

Import

Container Group's can be imported using the `resource id`, e.g.

```
terraform import azurerm_container_group.containerGroup1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.ContainerInstance/containerGroups/myContainerGroup1
```

azurerm_container_registry

Manages an Azure Container Registry.

Note: All arguments including the access key will be stored in the raw state as plain-text. Read more about sensitive data in state (/docs/state/sensitive-data.html).

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name     = "resourceGroup1"
  location = "West US"
}

resource "azurerm_container_registry" "acr" {
  name                    = "containerRegistry1"
  resource_group_name     = "${azurerm_resource_group.rg.name}"
  location                = "${azurerm_resource_group.rg.location}"
  sku                     = "Premium"
  admin_enabled           = false
  georeplication_locations = ["East US", "West Europe"]
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Container Registry. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Container Registry. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **admin_enabled** - (Optional) Specifies whether the admin user is enabled. Defaults to `false`.
- **storage_account_id** - (Required for `Classic` Sku - Forbidden otherwise) The ID of a Storage Account which must be located in the same Azure Region as the Container Registry.
- **sku** - (Optional) The SKU name of the the container registry. Possible values are `Classic` (which was previously `Basic`), `Basic`, `Standard` and `Premium`.

NOTE: The `Classic` SKU is Deprecated and will no longer be available for new resources from the end of March 2019.

- **tags** - (Optional) A mapping of tags to assign to the resource.

- `georeplication_locations` - (Optional) A list of Azure locations where the container registry should be geo-replicated.
- `network_rule_set` - (Optional) A `network_rule_set` block as documented below.

`network_rule_set` supports the following:

- `default_action` - (Optional) The behaviour for requests matching no rules. Either `Allow` or `Deny`. Defaults to `Allow`.
- `ip_rule` - (Optional) One or more `ip_rule` blocks as defined below.
- `virtual_network` - (Optional) One or more `virtual_network` blocks as defined below.

NOTE: `network_rule_set` is only supported with the `Premium` SKU at this time.

`ip_rule` supports the following:

- `action` - (Required) The behaviour for requests matching this rule. At this time the only supported value is `Allow`.
- `ip_range` - (Required) The CIDR block from which requests will match the rule.

`virtual_network` supports the following:

- `action` - (Required) The behaviour for requests matching this rule. At this time the only supported value is `Allow`.
- `subnet_id` - (Required) The subnet id from which requests will match the rule.

Attributes Reference

The following attributes are exported:

- `id` - The Container Registry ID.
- `login_server` - The URL that can be used to log into the container registry.
- `admin_username` - The Username associated with the Container Registry Admin account - if the admin account is enabled.
- `admin_password` - The Password associated with the Container Registry Admin account - if the admin account is enabled.

Import

Container Registries can be imported using the `resource id`, e.g.

```
terraform import azurerm_container_registry.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/mygroup1/providers/Microsoft.ContainerRegistry/registries/myregistry1
```

azurerm_container_registry_webhook

Manages an Azure Container Registry Webhook.

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_container_registry" "acr" {
  name                = "containerRegistry1"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  location            = "${azurerm_resource_group.rg.location}"
  sku                 = "Standard"
  admin_enabled       = false
}

resource "azurerm_container_registry_webhook" "webhook" {
  name                = "mywebhook"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  registry_name       = "${azurerm_container_registry.acr.name}"
  location            = "${azurerm_resource_group.rg.location}"

  service_uri = "https://mywebhookreceiver.example/mytag"
  status      = "enabled"
  scope       = "mytag:*"
  actions     = ["push"]
  custom_headers = { "Content-Type" = "application/json" }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Container Registry Webhook. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Container Registry Webhook. Changing this forces a new resource to be created.
- **registry_name** - (Required) The Name of Container registry this Webhook belongs to. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **service_uri** - (Required) Specifies the service URI for the Webhook to post notifications.

- **actions** - (Required) A list of actions that trigger the Webhook to post notifications. At least one action needs to be specified. Valid values are: `push`, `delete`, `quarantine`, `chart_push`, `chart_delete`
 - **status** - (Optional) Specifies if this Webhook triggers notifications or not. Valid values: `enabled` and `disabled`. Default is `enabled`.
 - **scope** - (Optional) Specifies the scope of repositories that can trigger an event. For example, `'foo:*'` means events for all tags under repository `'foo'`. `'foo:bar'` means events for `'foo:bar'` only. `'foo'` is equivalent to `'foo:latest'`. Empty means all events.
 - **custom_headers** - (Optional) Custom headers that will be added to the webhook notifications request.
-

Attributes Reference

The following attributes are exported:

- **id** - The Container Registry Webhook ID.

Import

Container Registry Webhooks can be imported using the `resource id`, e.g.

```
terraform import azurerm_container_registry_webhook.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/mygroup1/providers/Microsoft.ContainerRegistry/registries/myregistry1/webhooks/mywebhook1
```

azurerm_container_service

Manages an Azure Container Service Instance

NOTE: All arguments including the client secret will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

DEPRECATED: Azure Container Service (ACS) has been deprecated by Azure in favour of Azure (Managed) Kubernetes Service (AKS) (<https://azure.microsoft.com/en-us/updates/azure-container-service-will-retire-on-january-31-2020/>). Support for ACS will be removed in the next major version of the AzureRM Provider (2.0) - and we **strongly recommend** you consider using Azure Kubernetes Service (AKS) ([/docs/providers/azurerm/r/kubernetes_cluster.html](/docs/providers/azurerm/r/kubernetes_cluster.html)) for new deployments.

Example Usage (DCOS)

```

resource "azurerm_resource_group" "test" {
  name      = "acctestRG1"
  location  = "West US"
}

resource "azurerm_container_service" "test" {
  name                        = "acctestcontservice1"
  location                   = "${azurerm_resource_group.test.location}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  orchestration_platform     = "DCOS"

  master_profile {
    count      = 1
    dns_prefix = "acctestmaster1"
  }

  linux_profile {
    admin_username = "acctestuser1"

    ssh_key {
      key_data = "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCaZoyiz1qbd0Q8xEf6uEu1cCwYowo5FHtsBhqLoDnnp7KUTE
BN+L2NxRIfQ781rxV6Iq5jSav6b2Q8z5Kise0lvKA/RF2wqU0UPYqQviQhLmW6THTpmrv/YkUCuzxDpsH7DUDhZcwySLKVVe0Qm3+5N2T
a6UYH3lSDf9R9wTP2K/+vAnflKebuypNlmocIvakFwoZda18F0ms0oIVXQ8HWFNCuw9ZCunMSN62QGamCe3dL5cXlkgHYv7ekJE15IA9a
OJcM7e90oeTqo+7HTcWfdu0qQqPWY5ujyMw/llas8tsXY85LFqRnr3gJ02bAscjc477+X+j/gkpFoN1QEmt terraform@demo.tld"
    }
  }

  agent_pool_profile {
    name      = "default"
    count     = 1
    dns_prefix = "acctestagent1"
    vm_size   = "Standard_F2"
  }

  diagnostics_profile {
    enabled = false
  }

  tags = {
    Environment = "Production"
  }
}

```

Example Usage (Kubernetes)

```

resource "azurerm_resource_group" "test" {
  name      = "acctestRG1"
  location  = "West US"
}

resource "azurerm_container_service" "test" {
  name                        = "acctestcontservice1"
  location                   = "${azurerm_resource_group.test.location}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  orchestration_platform     = "Kubernetes"

  master_profile {
    count      = 1
    dns_prefix = "acctestmaster1"
  }

  linux_profile {
    admin_username = "acctestuser1"

    ssh_key {
      key_data = "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCaZoyiz1qbd0Q8xEf6uEu1cCwYowo5FHtsBhqLoDnnp7KUTE
BN+L2NxRIfQ781rxV6Iq5jSav6b2Q8z5Kise0lvKA/RF2wqU0UPYqQviQhLmW6THTpmrv/YkUCuzxDpsH7DUDhZcwySLKVVe0Qm3+5N2T
a6UYH3lsDf9R9wTP2K/+vAnflKebuypNlmocIvakFWoZda18F0ms0oIVXQ8HWFNCuw9ZCunMSN62QGamCe3dL5cXlkgHYv7ekJE15IA9a
OJcM7e90oeTqo+7HTcWfdu0qQqPWY5ujyMw/llas8tsXY85LFqRnr3gJ02bAscjc477+X+j/gkpFoN1QEmt terraform@demo.tld"
    }
  }

  agent_pool_profile {
    name      = "default"
    count     = 1
    dns_prefix = "acctestagent1"
    vm_size   = "Standard_F2"
  }

  service_principal {
    client_id      = "00000000-0000-0000-0000-000000000000"
    client_secret = "00000000000000000000000000000000"
  }

  diagnostics_profile {
    enabled = false
  }

  tags = {
    Environment = "Production"
  }
}

```

Example Usage (Swarm)

```

resource "azurerm_resource_group" "test" {
  name      = "acctestRG1"
  location  = "West US"
}

resource "azurerm_container_service" "test" {
  name                        = "acctestcontservice1"
  location                   = "${azurerm_resource_group.test.location}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  orchestration_platform     = "Swarm"

  master_profile {
    count      = 1
    dns_prefix = "acctestmaster1"
  }

  linux_profile {
    admin_username = "acctestuser1"

    ssh_key {
      key_data = "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCaZoyiz1qbd0Q8xEf6uEu1cCwYowo5FHtsBhqLoDnnp7KUTE
BN+L2NxRIfQ781rxV6Iq5jSav6b2Q8z5Kise0lvKA/RF2wqU0UPYqQviQhLmW6THTpmrv/YkUCuzxDpsH7DUDhZcwySLKVVe0Qm3+5N2T
a6UYH3lsDf9R9wTP2K/+vAnflKebuypNlmocIvakFWoZda18F0ms0oIVXQ8HWFNCuw9ZCunMSN62QGamCe3dL5cXlkgHYv7ekJE15IA9a
OJcM7e90oeTqo+7HTcWfdu0qQqPWY5ujyMw/llas8tsXY85LFqRnr3gJ02bAscjc477+X+j/gkpFoN1QEmt terraform@demo.tld"
    }
  }

  agent_pool_profile {
    name      = "default"
    count     = 1
    dns_prefix = "acctestagent1"
    vm_size   = "Standard_F2"
  }

  diagnostics_profile {
    enabled = false
  }

  tags = {
    Environment = "Production"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Container Service instance to create. Changing this forces a new resource to be created.
- `location` - (Required) The location where the Container Service instance should be created. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `orchestration_platform` - (Required) Specifies the Container Orchestration Platform to use. Currently can be either `DCOS` , `Kubernetes` or `Swarm` . Changing this forces a new resource to be created.
- `master_profile` - (Required) A Master Profile block as documented below.
- `linux_profile` - (Required) A Linux Profile block as documented below.
- `agent_pool_profile` - (Required) A Agent Pool Profile's block as documented below.
- `service_principal` - (only Required when you're using `Kubernetes` as an Orchestration Platform) A Service Principal block as documented below.
- `diagnostics_profile` - (Required) A VM Diagnostics Profile block as documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`master_profile` supports the following:

- `count` - (Required) Number of masters (VMs) in the container service cluster. Allowed values are 1, 3, and 5. The default value is 1.
- `dns_prefix` - (Required) The DNS Prefix to use for the Container Service master nodes.

`linux_profile` supports the following:

- `admin_username` - (Required) The Admin Username for the Cluster.
- `ssh_key` - (Required) An SSH Key block as documented below.

`ssh_key` supports the following:

- `key_data` - (Required) The Public SSH Key used to access the cluster.

`agent_pool_profile` supports the following:

- `name` - (Required) Unique name of the agent pool profile in the context of the subscription and resource group.
- `count` - (Required) Number of agents (VMs) to host docker containers. Allowed values must be in the range of 1 to 100 (inclusive). The default value is 1.
- `dns_prefix` - (Required) The DNS Prefix given to Agents in this Agent Pool.
- `vm_size` - (Required) The VM Size of each of the Agent Pool VM's (e.g. `Standard_F1` / `Standard_D2v2`).

`service_principal` supports the following:

- `client_id` - (Required) The ID for the Service Principal.
- `client_secret` - (Required) The secret password associated with the service principal.

`diagnostics_profile` supports the following:

- `enabled` - (Required) Should VM Diagnostics be enabled for the Container Service VM's

Attributes Reference

The following attributes are exported:

- `id` - The Container Service ID.
- `master_profile.fqdn` - FQDN for the master.
- `agent_pool_profile.fqdn` - FQDN for the agent pool.
- `diagnostics_profile.storage_uri` - The URI of the storage account where diagnostics are stored.

azurerm_cosmosdb_account

Manages a CosmosDB (formerly DocumentDB) Account.

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "${var.resource_group_name}"
  location  = "${var.resource_group_location}"
}

resource "random_integer" "ri" {
  min = 10000
  max = 99999
}

resource "azurerm_cosmosdb_account" "db" {
  name                        = "tfex-cosmos-db-${random_integer.ri.result}"
  location                   = "${azurerm_resource_group.rg.location}"
  resource_group_name        = "${azurerm_resource_group.rg.name}"
  offer_type                 = "Standard"
  kind                       = "GlobalDocumentDB"

  enable_automatic_failover = true

  consistency_policy {
    consistency_level      = "BoundedStaleness"
    max_interval_in_seconds = 10
    max_staleness_prefix   = 200
  }

  geo_location {
    location              = "${var.failover_location}"
    failover_priority = 1
  }

  geo_location {
    prefix              = "tfex-cosmos-db-${random_integer.ri.result}-customid"
    location            = "${azurerm_resource_group.rg.location}"
    failover_priority = 0
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the CosmosDB Account. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the CosmosDB Account is created. Changing this forces a new resource to be created.

- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `offer_type` - (Required) Specifies the Offer Type to use for this CosmosDB Account - currently this can only be set to `Standard`.
- `kind` - (Optional) Specifies the Kind of CosmosDB to create - possible values are `GlobalDocumentDB` and `MongoDB`. Defaults to `GlobalDocumentDB`. Changing this forces a new resource to be created.
- `consistency_policy` - (Required) Specifies a `consistency_policy` resource, used to define the consistency policy for this CosmosDB account.
- `geo_location` - (Required) Specifies a `geo_location` resource, used to define where data should be replicated with the `failover_priority 0` specifying the primary location.
- `ip_range_filter` - (Optional) CosmosDB Firewall Support: This value specifies the set of IP addresses or IP address ranges in CIDR form to be included as the allowed list of client IP's for a given database account. IP addresses/ranges must be comma separated and must not contain any spaces.
- `enable_automatic_failover` - (Optional) Enable automatic fail over for this Cosmos DB account.
- `capabilities` - (Optional) The capabilities which should be enabled for this Cosmos DB account. Possible values are `EnableAggregationPipeline`, `EnableCassandra`, `EnableGremlin`, `EnableTable`, `MongoDBv3.4`, and `mongoEnableDocLevelTTL`.
- `is_virtual_network_filter_enabled` - (Optional) Enables virtual network filtering for this Cosmos DB account.
- `virtual_network_rule` - (Optional) Specifies a `virtual_network_rules` resource, used to define which subnets are allowed to access this CosmosDB account.
- `enable_multiple_write_locations` - (Optional) Enable multi-master support for this Cosmos DB account.

`consistency_policy` Configures the database consistency and supports the following:

- `consistency_level` - (Required) The Consistency Level to use for this CosmosDB Account - can be either `BoundedStaleness`, `Eventual`, `Session`, `Strong` or `ConsistentPrefix`.
- `max_interval_in_seconds` - (Optional) When used with the `Bounded Staleness` consistency level, this value represents the time amount of staleness (in seconds) tolerated. Accepted range for this value is `5 - 86400` (1 day). Defaults to `5`. Required when `consistency_level` is set to `BoundedStaleness`.
- `max_staleness_prefix` - (Optional) When used with the `Bounded Staleness` consistency level, this value represents the number of stale requests tolerated. Accepted range for this value is `10 - 2147483647`. Defaults to `100`. Required when `consistency_level` is set to `BoundedStaleness`.

Note: `max_interval_in_seconds` and `max_staleness_prefix` can only be set to custom values when `consistency_level` is set to `BoundedStaleness` - otherwise they will return the default values shown above.

`geo_location` Configures the geographic locations the data is replicated to and supports the following:

- `prefix` - (Optional) The string used to generate the document endpoints for this region. If not specified it defaults to `${cosmosdb_account.name}-${location}`. Changing this causes the location to be deleted and re-provisioned and cannot be changed for the location with failover priority `0`.

- `location` - (Required) The name of the Azure region to host replicated data.
- `failover_priority` - (Required) The failover priority of the region. A failover priority of `0` indicates a write region. The maximum value for a failover priority = (total number of regions - 1). Failover priority values must be unique for each of the regions in which the database account exists. Changing this causes the location to be re-provisioned and cannot be changed for the location with failover priority `0`.

`capabilities` Configures the capabilities to enable for this Cosmos DB account:

- `name` - (Required) The capability to enable - Possible values are `EnableTable`, `EnableCassandra`, and `EnableGremlin`.

NOTE: The `prefix` and `failover_priority` fields of a location cannot be changed for the location with a failover priority of `0`.

`virtual_network_rule` Configures the virtual network subnets allowed to access this Cosmos DB account and supports the following:

- `id` - (Required) The ID of the virtual network subnet.

Attributes Reference

The following attributes are exported:

- `id` - The CosmosDB Account ID.
- `endpoint` - The endpoint used to connect to the CosmosDB account.
- `read_endpoints` - A list of read endpoints available for this CosmosDB account.
- `write_endpoints` - A list of write endpoints available for this CosmosDB account.
- `primary_master_key` - The Primary master key for the CosmosDB Account.
- `secondary_master_key` - The Secondary master key for the CosmosDB Account.
- `primary_readonly_master_key` - The Primary read-only master Key for the CosmosDB Account.
- `secondary_readonly_master_key` - The Secondary read-only master key for the CosmosDB Account.
- `connection_strings` - A list of connection strings available for this CosmosDB account. If the kind is `GlobalDocumentDB`, this will be empty.

Import

CosmosDB Accounts can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_account.account1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DocumentDB/databaseAccounts/account1
```


azurerm_cosmosdb_cassandra_keyspace

Manages a Cassandra KeySpace within a Cosmos DB Account.

Example Usage

```
data "azurerm_cosmosdb_account" "example" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

resource "azurerm_cosmosdb_cassandra_keyspace" "example" {
  name                = "tfex-cosmos-cassandra-keyspace"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cosmos DB Cassandra KeySpace. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cosmos DB Cassandra KeySpace is created. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the Cosmos DB Cassandra KeySpace to create the table within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB Cassandra KeySpace ID.

Import

Cosmos Cassandra KeySpace can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_cassandra_keyspace.ks1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/cassandra/keyspaces/ks1
```

azurerm_cosmosdb_mongo_collection

Manages a Mongo Collection within a Cosmos DB Account.

Example Usage

```
data "azurerm_cosmosdb_account" "example" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

resource "azurerm_cosmosdb_mongo_database" "example" {
  name                = "tfex-cosmos-mongo-db"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
}

resource "azurerm_cosmosdb_mongo_collection" "example" {
  name                = "tfex-cosmos-mongo-db"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
  database_name       = "${azurerm_cosmosdb_mongo_database.example.name}"

  default_ttl_seconds = "777"
  shard_key           = "uniqueKey"
  throughput          = 400

  indexes {
    key    = "aKey"
    unique = false
  }

  indexes {
    key    = "uniqueKey"
    unique = true
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cosmos DB Mongo Collection. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cosmos DB Mongo Collection is created. Changing this forces a new resource to be created.
- `database_name` - (Required) The name of the Cosmos DB Mongo Database in which the Cosmos DB Mongo Collection is created. Changing this forces a new resource to be created.

- `default_ttl_seconds` - (Required) The default Time To Live in seconds. If the value is `-1` items are not automatically expired.
 - `shard_key` - (Required) The name of the key to partition on for sharding. There must not be any other unique index keys.
 - `throughput` - (Optional) The throughput of the MongoDB collection (RU/s). Must be set in increments of `100`. The default and minimum value is `400`.
 - `indexes` - (Optional) One or more `indexes` blocks as defined below.
-

An `indexes` block supports the following:

- `key` - (Required) The name of the key to use for this index.
- `unique` - (Required) Whether the index key should be unique.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB Mongo Collection ID.

Import

Cosmos DB Mongo Collection can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_mongo_collection.collection1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/mongodb/databases/db1/collections/collection1
```

azurerm_cosmosdb_mongo_database

Manages a Mongo Database within a Cosmos DB Account.

Example Usage

```
data "azurerm_cosmosdb_account" "example" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

resource "azurerm_cosmosdb_mongo_database" "example" {
  name                = "tfex-cosmos-mongo-db"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cosmos DB Mongo Database. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cosmos DB Mongo Database is created. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the Cosmos DB Mongo Database to create the table within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB Mongo Database ID.

Import

Cosmos Mongo Database can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_mongo_database.db1 /subscriptions/00000000-0000-0000-0000-000000000000/
resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/mongodb/databases/db1
```

azurerm_cosmosdb_sql_container

Manages a SQL Container within a Cosmos DB Account.

Example Usage

```
resource "azurerm_cosmosdb_sql_container" "example" {
  name                        = "example-container"
  resource_group_name        = "${azurerm_cosmosdb_account.example.resource_group_name}"
  account_name               = "${azurerm_cosmosdb_account.example.name}"
  database_name              = "${azurerm_cosmosdb_sql_database.example.name}"
  partition_key_path         = "/definition/id"

  unique_key {
    paths = ["/definition/idlong", "/definition/idshort"]
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Cosmos DB SQL Database. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the Cosmos DB SQL Database is created. Changing this forces a new resource to be created.
- **account_name** - (Required) The name of the Cosmos DB Account to create the container within. Changing this forces a new resource to be created.
- **database_name** - (Required) The name of the Cosmos DB SQL Database to create the container within. Changing this forces a new resource to be created.
- **partition_key_path** - (Optional) Define a partition key. Changing this forces a new resource to be created.
- **unique_key** - (Optional) One or more `unique_key` blocks as defined below. Changing this forces a new resource to be created.

A `unique_key` block supports the following:

- **paths** - (Required) A list of paths to use for this unique key.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB SQL Database ID.

Import

Cosmos SQL Database can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_sql_container.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/sql/databases/databases1/containers/example
```

azurerm_cosmosdb_sql_database

Manages a SQL Database within a Cosmos DB Account.

Example Usage

```
data "azurerm_cosmosdb_account" "example" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

resource "azurerm_cosmosdb_sql_database" "example" {
  name                = "tfex-cosmos-mongo-db"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cosmos DB SQL Database. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cosmos DB SQL Database is created. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the Cosmos DB SQL Database to create the table within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB SQL Database ID.

Import

Cosmos SQL Database can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_sql_database.db1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/sql/databases/db1
```

azurerm_cosmosdb_table

Manages a Table within a Cosmos DB Account.

Example Usage

```
data "azurerm_cosmosdb_account" "example" {
  name                = "tfex-cosmosdb-account"
  resource_group_name = "tfex-cosmosdb-account-rg"
}

resource "azurerm_cosmosdb_table" "example" {
  name                = "tfex-cosmos-table"
  resource_group_name = "${data.azurerm_cosmosdb_account.example.resource_group_name}"
  account_name        = "${data.azurerm_cosmosdb_account.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Cosmos DB Table. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Cosmos DB Table is created. Changing this forces a new resource to be created.
- `account_name` - (Required) The name of the Cosmos DB Table to create the table within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - the Cosmos DB Table ID.

Import

Cosmos Tables can be imported using the `resource id`, e.g.

```
terraform import azurerm_cosmosdb_table.t1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.DocumentDB/databaseAccounts/account1/apis/table/tables/t1
```


azurerm_dashboard

Manages a shared dashboard in the Azure Portal.

Example Usage

```
variable "md_content" {
  description = "Content for the MD tile"
  default     = "# Hello all :)"
}

variable "video_link" {
  description = "Link to a video"
  default     = "https://www.youtube.com/watch?v=....."
}

data "azurerm_subscription" "current" {}

resource "azurerm_resource_group" "my-group" {
  name       = "mygroup"
  location   = "uksouth"
}

resource "azurerm_dashboard" "my-board" {
  name                = "my-cool-dashboard"
  resource_group_name = azurerm_resource_group.my-group.name
  location             = azurerm_resource_group.my-group.location
  tags = {
    source = "terraform"
  }
  dashboard_properties = <<DASH
{
  "lenses": {
    "0": {
      "order": 0,
      "parts": {
        "0": {
          "position": {
            "x": 0,
            "y": 0,
            "rowSpan": 2,
            "colSpan": 3
          },
          "metadata": {
            "inputs": [],
            "type": "Extension/HubsExtension/PartType/MarkdownPart",
            "settings": {
              "content": {
                "settings": {
                  "content": "${var.md_content}",
                  "subtitle": "",
                  "title": ""
                }
              }
            }
          }
        }
      }
    }
  }
}
```

```

    }
  },
  "1": {
    "position": {
      "x": 5,
      "y": 0,
      "rowSpan": 4,
      "colSpan": 6
    },
    "metadata": {
      "inputs": [],
      "type": "Extension/HubsExtension/PartType/VideoPart",
      "settings": {
        "content": {
          "settings": {
            "title": "Important Information",
            "subtitle": "",
            "src": "${var.video_link}",
            "autoplay": true
          }
        }
      }
    }
  },
  "2": {
    "position": {
      "x": 0,
      "y": 4,
      "rowSpan": 4,
      "colSpan": 6
    },
    "metadata": {
      "inputs": [
        {
          "name": "ComponentId",
          "value": "/subscriptions/${data.azurerm_subscription.current.subscription_id}/resourceGroups/myRG/providers/microsoft.insights/components/myWebApp"
        }
      ],
      "type": "Extension/AppInsightsExtension/PartType/AppMapGalPt",
      "settings": {},
      "asset": {
        "idInputName": "ComponentId",
        "type": "ApplicationInsights"
      }
    }
  },
  "model": {
    "timeRange": {
      "value": {
        "relative": {
          "duration": 24,
          "timeUnit": 1
        }
      }
    }
  },

```

```

        "type": "MsPortalFx.Composition.Configuration.ValueTypes.TimeRange"
    },
    "filterLocale": {
        "value": "en-us"
    },
    "filters": {
        "value": {
            "MsPortalFx_TimeRange": {
                "model": {
                    "format": "utc",
                    "granularity": "auto",
                    "relative": "24h"
                },
                "displayCache": {
                    "name": "UTC Time",
                    "value": "Past 24 hours"
                },
                "filteredPartIds": [
                    "StartboardPart-UnboundPart-ae44fef5-76b8-46b0-86f0-2b3f47bad1c7"
                ]
            }
        }
    }
}
}
}
}
}
DASH
}

```

It is recommended to follow the steps outlined here (<https://docs.microsoft.com/en-us/azure/azure-portal/azure-portal-dashboards-create-programmatically#fetch-the-json-representation-of-the-dashboard>) to create a Dashboard in the Portal and extract the relevant JSON to use in this resource. From the extracted JSON, the contents of the `properties: {}` object can be used. Variables can be injected as needed - see above example.

Using a `template_file` data source or the `templatefile` function

Since the contents of the dashboard JSON can be quite lengthy, use a template file to improve readability:

`dash.tpl` :

```

{
  "lenses": {
    "0": {
      "order": 0,
      "parts": {
        "0": {
          "position": {
            "x": 0,
            "y": 0,
            "rowSpan": 2,
            "colSpan": 3
          },
          "metadata": {
            "inputs": [],
            "type": "Extension/HubsExtension/PartType/MarkdownPart",
            "settings": {
              "content": {
                "settings": {
                  "content": "${md_content}", // <-- note the 'var.' is dropped
                  "subtitle": "",
                  "title": ""
                }
              }
            }
          }
        }
      }
    },
    ...
    ...
  }
}

```

This is then referenced in the `.tf` file by using a `template_file`

(<https://www.terraform.io/docs/providers/template/d/file.html>) data source (terraform 0.11 or earlier), or the

`templatefile` (<https://www.terraform.io/docs/configuration/functions/templatefile.html>) function (terraform 0.12+).

`main.tf` (terraform 0.11 or earlier):

```

data "template_file" "dash-template" {
  template = "${file("${path.module}/dash.tpl")}"
  vars = {
    md_content = "Variable content here!"
    video_link = "https://www.youtube.com/watch?v=....."
    sub_id      = data.azurerm_subscription.current.subscription_id
  }
}

...

resource "azurerm_dashboard" "my-board" {
  name                        = "my-cool-dashboard"
  resource_group_name = azurerm_resource_group.my-group.name
  location                 = azurerm_resource_group.my-group.location
  tags = {
    source = "terraform"
  }
  dashboard_properties = data.template_file.dash-template.rendered
}

```

main.tf (terraform 0.12+)

```

...

resource "azurerm_dashboard" "my-board" {
  name                        = "my-cool-dashboard"
  resource_group_name = azurerm_resource_group.my-group.name
  location                 = azurerm_resource_group.my-group.location
  tags = {
    source = "terraform"
  }
  dashboard_properties = templatefile("dash.tpl",
  {
    md_content = "Variable content here!",
    video_link = "https://www.youtube.com/watch?v=.....",
    sub_id      = data.azurerm_subscription.current.subscription_id
  })
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Shared Dashboard. This should be be 64 chars max, only alphanumeric and hyphens (no spaces). For a more friendly display name, add the `hidden-title` tag.
- **resource_group_name** - (Required) The name of the resource group in which to create the dashboard.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.

- `dashboard_properties` - (Required) JSON data representing dashboard body. See above for details on how to obtain this from the Portal.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Dashboard ID.

Import

Dashboards can be imported using the `resource id`, e.g.

```
terraform import azurerm_dashboard.my-board /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/rg1/providers/Microsoft.Portal/dashboards/00000000-0000-0000-0000-000000000000
```

Note the URI in the above sample can be found using the Resource Explorer tool in the Azure Portal.

azurerm_databricks_workspace

Manages a Databricks Workspace

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_databricks_workspace" "test" {
  name                        = "databricks-test"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  sku                        = "standard"

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Databricks Workspace resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Databricks Workspace should exist. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource has to be created. Changing this forces a new resource to be created.
- `sku` - (Required) The `sku` to use for the Databricks Workspace. Possible values are `standard` or `premium`. Changing this forces a new resource to be created.
- `managed_resource_group_name` - (Optional) The name of the resource group where Azure should place the managed Databricks resources. Changing this forces a new resource to be created.

NOTE Azure requires that this Resource Group does not exist in this Subscription (and that the Azure API creates it) - otherwise the deployment will fail.

- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Databricks Workspace.
- `managed_resource_group_id` - The ID of the Managed Resource Group created by the Databricks Workspace.

Import

Databrick Workspaces can be imported using the `resource id`, e.g.

```
terraform import azurerm_databricks_workspace.workspace1 /subscriptions/00000000-0000-0000-0000-000000000000/000/resourceGroups/group1/providers/Microsoft.Databricks/workspaces/workspace1
```


azurerm_data_factory_dataset_mysql

Manages a MySQL Dataset inside a Azure Data Factory.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_mysql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  connection_string   = "Server=test;Port=3306;Database=test;User=test;SSLMode=1;UseSystemTrustStore=0;Password=test"
}

resource "azurerm_data_factory_dataset_mysql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  linked_service_name = "${azurerm_data_factory_linked_service_mysql.test.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Dataset MySQL. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Dataset MySQL. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Dataset with. Changing this forces a new resource.
- **linked_service_name** - (Required) The Data Factory Linked Service name in which to associate the Dataset with.
- **table_name** - (Optional) The table name of the Data Factory Dataset MySQL.
- **folder** - (Optional) The folder that this Dataset is in. If not specified, the Dataset will appear at the root level.

- `schema_column` - (Optional) A `schema_column` block as defined below.
- `description` - (Optional) The description for the Data Factory Dataset MySQL.
- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Dataset MySQL.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Dataset MySQL.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Dataset MySQL.

A `schema_column` block supports the following:

- `name` - (Required) The name of the column.
- `type` - (Optional) Type of the column. Valid values are `Byte`, `Byte[]`, `Boolean`, `Date`, `DateTime`, `DateTimeOffset`, `Decimal`, `Double`, `Guid`, `Int16`, `Int32`, `Int64`, `Single`, `String`, `TimeSpan`. Please note these values are case sensitive.
- `description` - (Optional) The description of the column.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Dataset.

Import

Data Factory Dataset MySQL can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_dataset_mysql.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/datasets/example
```

azurerm_data_factory_dataset_postgresql

Manages a PostgreSQL Dataset inside a Azure Data Factory.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_postgresql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  connection_string   = "Host=example;Port=5432;Database=example;UID=example;EncryptionMethod=0;Password=example"
}

resource "azurerm_data_factory_dataset_postgresql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  linked_service_name = "${azurerm_data_factory_linked_service_postgresql.test.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Dataset PostgreSQL. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Dataset PostgreSQL. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Dataset with. Changing this forces a new resource.
- **linked_service_name** - (Required) The Data Factory Linked Service name in which to associate the Dataset with.
- **table_name** - (Optional) The table name of the Data Factory Dataset PostgreSQL.
- **folder** - (Optional) The folder that this Dataset is in. If not specified, the Dataset will appear at the root level.

- `schema_column` - (Optional) A `schema_column` block as defined below.
- `description` - (Optional) The description for the Data Factory Dataset PostgreSQL.
- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Dataset PostgreSQL.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Dataset PostgreSQL.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Dataset PostgreSQL.

A `schema_column` block supports the following:

- `name` - (Required) The name of the column.
- `type` - (Optional) Type of the column. Valid values are `Byte`, `Byte[]`, `Boolean`, `Date`, `DateTime`, `DateTimeOffset`, `Decimal`, `Double`, `Guid`, `Int16`, `Int32`, `Int64`, `Single`, `String`, `TimeSpan`. Please note these values are case sensitive.
- `description` - (Optional) The description of the column.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Dataset.

Import

Data Factory Dataset PostgreSQL can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_dataset_postgresql.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/datasets/example
```

azurerm_data_factory_dataset_sql_server

Manages a SQL Server Table Dataset inside a Azure Data Factory.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_sql_server" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  connection_string   = "Integrated Security=False;Data Source=test;Initial Catalog=test;User ID=test;Password=test"
}

resource "azurerm_data_factory_dataset_sql_server_table" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  linked_service_name = "${azurerm_data_factory_linked_service_sql_server.test.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Dataset SQL Server Table. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Dataset SQL Server Table. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Dataset with. Changing this forces a new resource.
- **linked_service_name** - (Required) The Data Factory Linked Service name in which to associate the Dataset with.
- **table_name** - (Optional) The table name of the Data Factory Dataset SQL Server Table.
- **folder** - (Optional) The folder that this Dataset is in. If not specified, the Dataset will appear at the root level.

- `schema_column` - (Optional) A `schema_column` block as defined below.
- `description` - (Optional) The description for the Data Factory Dataset SQL Server Table.
- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Dataset SQL Server Table.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Dataset SQL Server Table.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Dataset SQL Server Table.

A `schema_column` block supports the following:

- `name` - (Required) The name of the column.
- `type` - (Optional) Type of the column. Valid values are `Byte`, `Byte[]`, `Boolean`, `Date`, `DateTime`, `DateTimeOffset`, `Decimal`, `Double`, `Guid`, `Int16`, `Int32`, `Int64`, `Single`, `String`, `TimeSpan`. Please note these values are case sensitive.
- `description` - (Optional) The description of the column.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Dataset.

Import

Data Factory Dataset SQL Server Table can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_dataset_sql_server_table.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/datasets/example
```

azurerm_data_factory

Manages an Azure Data Factory (Version 2).

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example"
  location = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Factory. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- `resource_group_name` - (Required) The name of the resource group in which to create the Data Factory.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `github_configuration` - (Optional) A `github_configuration` block as defined below.
- `identity` - (Optional) An `identity` block as defined below.
- `vsts_configuration` - (Optional) A `vsts_configuration` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `github_configuration` block supports the following:

- `account_name` - (Required) Specifies the GitHub account name.
- `branch_name` - (Required) Specifies the branch of the repository to get code from.
- `git_url` - (Required) Specifies the GitHub Enterprise host name. For example: <https://github.mydomain.com> (<https://github.mydomain.com>). Use <https://github.com> (<https://github.com>) for open source repositories.
- `repository_name` - (Required) Specifies the name of the git repository.
- `root_folder` - (Required) Specifies the root folder within the repository. Set to `/` for the top level.

Note: You must log in to the Data Factory management UI to complete the authentication to the GitHub repository.

A `identity` block supports the following:

- `type` - (Required) Specifies the identity type of the Data Factory. At this time the only allowed value is `SystemAssigned`.

A `vsts_configuration` block supports the following:

- `account_name` - (Required) Specifies the VSTS account name.
- `branch_name` - (Required) Specifies the branch of the repository to get code from.
- `project_name` - (Required) Specifies the name of the VSTS project.
- `repository_name` - (Required) Specifies the name of the git repository.
- `root_folder` - (Required) Specifies the root folder within the repository. Set to `/` for the top level.
- `tenant_id` - (Required) Specifies the Tenant ID associated with the VSTS account.

Attributes Reference

The following attributes are exported:

- `id` - The Data Factory ID.
- `identity` - An `identity` block as defined below.

The `identity` block exports the following:

- `principal_id` - The ID of the Principal (Client) in Azure Active Directory
- `tenant_id` - The ID of the Azure Active Directory Tenant.

Import

Data Factory can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example
```


azurerm_data_factory_integration_runtime_managed

Manages an Azure Data Factory Managed Integration Runtime.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_integration_runtime_managed" "example" {
  name                = "example"
  data_factory_name   = "${azurerm_data_factory.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location             = "${azurerm_resource_group.example.location}"

  node_size = "Standard_D8_v3"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Managed Integration Runtime. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **data_factory_name** - (Required) Specifies the name of the Data Factory the Managed Integration Runtime belongs to. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Managed Integration Runtime. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **node_size** - (Required) The size of the nodes on which the Managed Integration Runtime runs. Valid values are: `Standard_D2_v3`, `Standard_D4_v3`, `Standard_D8_v3`, `Standard_D16_v3`, `Standard_D32_v3`, `Standard_D64_v3`, `Standard_E2_v3`, `Standard_E4_v3`, `Standard_E8_v3`, `Standard_E16_v3`, `Standard_E32_v3`, `Standard_E64_v3`, `Standard_D1_v2`, `Standard_D2_v2`, `Standard_D3_v2`, `Standard_D4_v2`, `Standard_A4_v2` and `Standard_A8_v2`
- **number_of_nodes** - (Optional) Number of nodes for the Managed Integration Runtime. Max is 10 . Defaults to 1 .

- `max_parallel_executions_per_node` - (Optional) Defines the maximum parallel executions per node. Defaults to 1. Max is 16.
 - `edition` - (Optional) The Managed Integration Runtime edition. Valid values are `Standard` and `Enterprise`. Defaults to `Standard`.
 - `license_type` - (Optional) The type of the license that is used. Valid values are `LicenseIncluded` and `BasePrize`. Defaults to `LicenseIncluded`.
 - `catalog_info` - (Optional) A `catalog_info` block as defined below.
 - `custom_setup_script` - (Optional) A `custom_setup_script` block as defined below.
 - `vnet_integration` - (Optional) A `vnet_integration` block as defined below.
-

A `catalog_info` block supports the following:

- `server_endpoint` - (Required) The endpoint of an Azure SQL Server that will be used to host the SSIS catalog.
 - `administrator_login` - (Required) Administrator login name for the SQL Server.
 - `administrator_password` - (Required) Administrator login password for the SQL Server.
 - `pricing_tier` - (Required) Pricing tier for the database that will be created for the SSIS catalog. Valid values are: `Basic`, `Standard`, `Premium` and `PremiumRS`.
-

A `custom_setup_script` block supports the following:

- `blob_container_uri` - (Required) The blob endpoint for the container which contains a custom setup script that will be run on every node on startup. See <https://docs.microsoft.com/en-us/azure/data-factory/how-to-configure-azure-ssis-ir-custom-setup> (<https://docs.microsoft.com/en-us/azure/data-factory/how-to-configure-azure-ssis-ir-custom-setup>) for more information.
 - `sas_token` - (Required) A container SAS token that gives access to the files. See <https://docs.microsoft.com/en-us/azure/data-factory/how-to-configure-azure-ssis-ir-custom-setup> (<https://docs.microsoft.com/en-us/azure/data-factory/how-to-configure-azure-ssis-ir-custom-setup>) for more information.
-

A `vnet_integration` block supports the following:

- `vnet_id` - (Required) ID of the virtual network to which the nodes of the Managed Integration Runtime will be added.
- `subnet_name` - (Required) Name of the subnet to which the nodes of the Managed Integration Runtime will be added.

Attributes Reference

The following attributes are exported:

- `id` - The Managed Integration Runtime ID.

Import

Managed Integration Runtime can be imported using the `resource_id`, e.g.

```
terraform import azurerm_data_factory.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/integrationruntimes/example
```

azurerm_data_factory_linked_service_data_lake_storage_gen2

Manages a Linked Service (connection) between Data Lake Storage Gen2 and Azure Data Factory.

Note: All arguments including the `service_principal_key` will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

data "azurerm_client_config" "current" {}

resource "azurerm_data_factory_linked_service_data_lake_storage_gen2" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.test.name}"
  data_factory_name   = "${azurerm_data_factory.test.name}"
  service_principal_id = "${data.azurerm_client_config.current.client_id}"
  service_principal_key = "exampleKey"
  tenant              = "11111111-1111-1111-1111-111111111111"
  url                 = "https://datalakestoragegen2"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Factory Linked Service MySQL. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- `resource_group_name` - (Required) The name of the resource group in which to create the Data Factory Linked Service MySQL. Changing this forces a new resource
- `data_factory_name` - (Required) The Data Factory name in which to associate the Linked Service with. Changing this forces a new resource.
- `url` - (Required) The endpoint for the Azure Data Lake Storage Gen2 service.
- `service_principal_id` - (Required) The service principal id in which to authenticate against the Azure Data Lake Storage Gen2 account.

- `service_principal_key` - (Required) The service principal key in which to authenticate against the Azure Data Lake Storage Gen2 account.
- `tenant` - (Required) The tenant id or name in which to authenticate against the Azure Data Lake Storage Gen2 account.
- `description` - (Optional) The description for the Data Factory Linked Service MySQL.
- `integration_runtime_name` - (Optional) The integration runtime reference to associate with the Data Factory Linked Service MySQL.
- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Linked Service MySQL.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Linked Service MySQL.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Linked Service MySQL.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Linked Service.

Import

Data Factory Linked Service MySQL can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_linked_service_mysql.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/linkedservices/example
```

azurerm_data_factory_linked_service_mysql

Manages a Linked Service (connection) between MySQL and Azure Data Factory.

Note: All arguments including the connection_string will be stored in the raw state as plain-text. Read more about sensitive data in state (/docs/state/sensitive-data.html).

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_mysql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.test.name}"
  data_factory_name   = "${azurerm_data_factory.test.name}"
  connection_string    = "Server=test;Port=3306;Database=test;User=test;SSLMode=1;UseSystemTrustStore=0;Password=test"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Linked Service MySQL. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Linked Service MySQL. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Linked Service with. Changing this forces a new resource.
- **connection_string** - (Required) The connection string in which to authenticate with MySQL.
- **description** - (Optional) The description for the Data Factory Linked Service MySQL.
- **integration_runtime_name** - (Optional) The integration runtime reference to associate with the Data Factory Linked Service MySQL.

- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Linked Service MySQL.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Linked Service MySQL.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Linked Service MySQL.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Linked Service.

Import

Data Factory Linked Service MySQL can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_linked_service_mysql.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/linkedservices/example
```

azurerm_data_factory_linked_service_postgresql

Manages a Linked Service (connection) between PostgreSQL and Azure Data Factory.

Note: All arguments including the `connection_string` will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_postgresql" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.test.name}"
  data_factory_name   = "${azurerm_data_factory.test.name}"
  connection_string    = "Host=example;Port=5432;Database=example;UID=example;EncryptionMethod=0;Password=example"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Factory Linked Service PostgreSQL. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- `resource_group_name` - (Required) The name of the resource group in which to create the Data Factory Linked Service PostgreSQL. Changing this forces a new resource
- `data_factory_name` - (Required) The Data Factory name in which to associate the Linked Service with. Changing this forces a new resource.
- `connection_string` - (Required) The connection string in which to authenticate with PostgreSQL.
- `description` - (Optional) The description for the Data Factory Linked Service PostgreSQL.
- `integration_runtime_name` - (Optional) The integration runtime reference to associate with the Data Factory Linked Service PostgreSQL.

- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Linked Service PostgreSQL.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Linked Service PostgreSQL.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Linked Service PostgreSQL.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Linked Service.

Import

Data Factory Linked Service PostgreSQL can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_linked_service_postgresql.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/linkedservices/example
```

azurerm_data_factory_linked_service_sql_server

Manages a Linked Service (connection) between a SQL Server and Azure Data Factory.

Note: All arguments including the client secret will be stored in the raw state as plain-text. Read more about sensitive data in state (/docs/state/sensitive-data.html).

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_linked_service_sql_server" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.test.name}"
  data_factory_name   = "${azurerm_data_factory.test.name}"
  connection_string    = "Integrated Security=False;Data Source=test;Initial Catalog=test;User ID=test;Password=test"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Linked Service SQL Server. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Linked Service SQL Server. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Linked Service with. Changing this forces a new resource.
- **connection_string** - (Required) The connection string in which to authenticate with the SQL Server.
- **description** - (Optional) The description for the Data Factory Linked Service SQL Server.
- **integration_runtime_name** - (Optional) The integration runtime reference to associate with the Data Factory Linked Service SQL Server.

- `annotations` - (Optional) List of tags that can be used for describing the Data Factory Linked Service SQL Server.
- `parameters` - (Optional) A map of parameters to associate with the Data Factory Linked Service SQL Server.
- `additional_properties` - (Optional) A map of additional properties to associate with the Data Factory Linked Service SQL Server.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Linked Service.

Import

Data Factory Linked Service SQL Server can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_linked_service_sql_server.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/linkedservices/example
```

azurerm_data_factory_pipeline

Manages a Pipeline inside a Azure Data Factory.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_factory" "example" {
  name                = "example"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_data_factory_pipeline" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  data_factory_name   = "${azurerm_data_factory.example.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Factory Pipeline. Changing this forces a new resource to be created. Must be globally unique. See the Microsoft documentation (<https://docs.microsoft.com/en-us/azure/data-factory/naming-rules>) for all restrictions.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Factory Pipeline. Changing this forces a new resource
- **data_factory_name** - (Required) The Data Factory name in which to associate the Pipeline with. Changing this forces a new resource.
- **description** - (Optional) The description for the Data Factory Pipeline.
- **annotations** - (Optional) List of tags that can be used for describing the Data Factory Pipeline.
- **parameters** - (Optional) A map of parameters to associate with the Data Factory Pipeline.
- **variables** - (Optional) A map of variables to associate with the Data Factory Pipeline.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Data Factory Pipeline.

Import

Data Factory Pipeline can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_factory_pipeline.example /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/example/providers/Microsoft.DataFactory/factories/example/pipelines/example
```

azurerm_data_lake_analytics_account

Manages an Azure Data Lake Analytics Account.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-datalake-account"
  location  = "northeurope"
}

resource "azurerm_data_lake_store" "example" {
  name                = "tfexdatalakestore"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"
}

resource "azurerm_data_lake_analytics_account" "example" {
  name                = "tfexdatalakeaccount"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"

  default_store_account_name = "${azurerm_data_lake_store.example.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Lake Analytics Account. Changing this forces a new resource to be created. Has to be between 3 to 24 characters.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Lake Analytics Account.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **default_store_account_name** - (Required) Specifies the data lake store to use by default. Changing this forces a new resource to be created.
- **tier** - (Optional) The monthly commitment tier for Data Lake Analytics Account. Accepted values are `Consumption`, `Commitment_100000AUHours`, `Commitment_10000AUHours`, `Commitment_1000AUHours`, `Commitment_100AUHours`, `Commitment_500000AUHours`, `Commitment_50000AUHours`, `Commitment_5000AUHours`, or `Commitment_500AUHours`.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Date Lake Store ID.

Import

Date Lake Analytics Account can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_lake_analytics_account.test /subscriptions/00000000-0000-0000-0000-000000000000/0000/resourceGroups/mygroup1/providers/Microsoft.DataLakeAnalytics/accounts/mydatalakeaccount
```

azurerm_data_lake_analytics_firewall_rule

Manages a Azure Data Lake Analytics Firewall Rule.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex_datalake_fw_rule"
  location  = "northeurope"
}

resource "azurerm_data_lake_store" "example" {
  name                  = "tfexdatalakestore"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"
}

resource "azurerm_data_lake_analytics_account" "example" {
  name                  = "tfexdatalakeaccount"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"

  default_store_account_name = "${azurerm_data_lake_store.example.name}"
}

resource "azurerm_data_lake_analytics_firewall_rule" "example" {
  name              = "office-ip-range"
  account_name      = "${azurerm_data_lake_analytics.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  start_ip_address  = "1.2.3.4"
  end_ip_address    = "2.3.4.5"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Data Lake Analytics. Changing this forces a new resource to be created. Has to be between 3 to 24 characters.
- **resource_group_name** - (Required) The name of the resource group in which to create the Data Lake Analytics.
- **account_name** - (Required) Specifies the name of the Data Lake Analytics for which the Firewall Rule should take effect.
- **start_ip_address** - (Required) The Start IP address for the firewall rule.
- **end_ip_address** - (Required) The End IP Address for the firewall rule.

Attributes Reference

The following attributes are exported:

- `id` - The Date Lake Store Firewall Rule ID.

Import

Date Lake Store Firewall Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_lake_analytics_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DataLakeAnalytics/accounts/mydatalakeaccount/firewallRules/rule1
```

azurerm_data_lake_store_file

Manages a Azure Data Lake Store File.

Note: If you want to change the data in the remote file without changing the `local_file_path`, then taint the resource so the `azurerm_data_lake_store_file` gets recreated with the new data.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_lake_store" "example" {
  name                  = "consumptiondatalake"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"
}

resource "azurerm_data_lake_store_file" "example" {
  resource_group_name = "${azurerm_resource_group.example.name}"
  local_file_path     = "/path/to/local/file"
  remote_file_path    = "/path/created/for/remote/file"
}
```

Argument Reference

The following arguments are supported:

- `account_name` - (Required) Specifies the name of the Data Lake Store for which the File should created.
- `local_file_path` - (Required) The path to the local file to be added to the Data Lake Store.
- `remote_file_path` - (Required) The path created for the file on the Data Lake Store.

Import

Date Lake Store File's can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_lake_store_file.test example.azuredatalakestore.net/test/example.txt
```

azurerm_data_lake_store_firewall_rule

Manages a Azure Data Lake Store Firewall Rule.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example"
  location = "northeurope"
}

resource "azurerm_data_lake_store" "example" {
  name                = "consumptiondatalake"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location             = "${azurerm_resource_group.example.location}"
}

resource "azurerm_data_lake_store_firewall_rule" "example" {
  name                = "office-ip-range"
  account_name        = "${azurerm_data_lake_store.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  start_ip_address    = "1.2.3.4"
  end_ip_address      = "2.3.4.5"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Lake Store. Changing this forces a new resource to be created. Has to be between 3 to 24 characters.
- `resource_group_name` - (Required) The name of the resource group in which to create the Data Lake Store.
- `account_name` - (Required) Specifies the name of the Data Lake Store for which the Firewall Rule should take effect.
- `start_ip_address` - (Required) The Start IP address for the firewall rule.
- `end_ip_address` - (Required) The End IP Address for the firewall rule.

Attributes Reference

The following attributes are exported:

- `id` - The Date Lake Store Firewall Rule ID.

Import

Data Lake Store Firewall Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_lake_store_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DataLakeStore/accounts/mydatalakeaccount/firewallRules/rule1
```

azurerm_data_lake_store

Manages an Azure Data Lake Store.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "northeurope"
}

resource "azurerm_data_lake_store" "example" {
  name                  = "consumptiondatalake"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"
  encryption_state     = "Enabled"
  encryption_type      = "ServiceManaged"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Data Lake Store. Changing this forces a new resource to be created. Has to be between 3 to 24 characters.
- `resource_group_name` - (Required) The name of the resource group in which to create the Data Lake Store.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tier` - (Optional) The monthly commitment tier for Data Lake Store. Accepted values are `Consumption`, `Commitment_1TB`, `Commitment_10TB`, `Commitment_100TB`, `Commitment_500TB`, `Commitment_1PB` or `Commitment_5PB`.
- `encryption_state` - (Optional) Is Encryption enabled on this Data Lake Store Account? Possible values are `Enabled` or `Disabled`. Defaults to `Enabled`.
- `encryption_type` - (Optional) The Encryption Type used for this Data Lake Store Account. Currently can be set to `ServiceManaged` when `encryption_state` is `Enabled` - and must be a blank string when it's `Disabled`.

NOTE: Support for User Managed encryption will be supported in the future once a bug in the API is fixed.

- `firewall_allow_azure_ips` - are Azure Service IP's allowed through the firewall? Possible values are `Enabled` and `Disabled`. Defaults to `Enabled`.
- `firewall_state` - the state of the Firewall. Possible values are `Enabled` and `Disabled`. Defaults to `Enabled`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Date Lake Store ID.
- `endpoint` - The Endpoint for the Data Lake Store.

Import

Date Lake Store can be imported using the `resource id`, e.g.

```
terraform import azurerm_data_lake_store.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DataLakeStore/accounts/mydatalakeaccount
```

azurerm_ddos_protection_plan

Manages an Azure DDoS Protection Plan.

NOTE Azure only allow one DDoS Protection Plan per region.

NOTE: This resource has been deprecated in favour of the `azurerm_network_ddos_protection_plan` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_ddos_protection_plan" "test" {
  name                = "example-protection-plan"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the DDoS Protection Plan. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the DDoS Protection Plan
- `virtual_network_ids` - The Resource ID list of the Virtual Networks associated with DDoS Protection Plan.

Import

Azure DDoS Protection Plan can be imported using the `resource_id`, e.g.

```
terraform import azurerm_ddos_protection_plan.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/ddosProtectionPlans/testddosplan
```


azurerm_devspace_controller

Manages a DevSpace Controller.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example_resources"
  location  = "West Europe"
}

resource "azurerm_kubernetes_cluster" "test" {
  name                = "acctestaks1"
  location            = azurerm_resource_group.example.location
  resource_group_name = azurerm_resource_group.example.name
  dns_prefix          = "acctestaks1"

  agent_pool_profile {
    name     = "default"
    count    = "1"
    vm_size  = "Standard_DS2_v2"
  }

  service_principal {
    client_id     = "00000000-0000-0000-0000-000000000000"
    client_secret = "00000000000000000000000000000000"
  }
}

resource "azurerm_devspace_controller" "test" {
  name                = "acctestdsc1"
  location            = azurerm_resource_group.test.location
  resource_group_name = azurerm_resource_group.test.name

  sku {
    name = "S1"
    tier  = "Standard"
  }

  host_suffix          = "suffix"
  target_container_host_resource_id = "${azurerm_kubernetes_cluster.test.id}"
  target_container_host_credentials_base64 = "${base64encode(azurerm_kubernetes_cluster.test.kube_config_raw)}"

  tags = {
    Environment = "Testing"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the DevSpace Controller. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group under which the DevSpace Controller resource has to be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported location where the DevSpace Controller should exist. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as documented below. Changing this forces a new resource to be created.
- `target_container_host_resource_id` - (Required) The resource id of Azure Kubernetes Service cluster. Changing this forces a new resource to be created.
- `target_container_host_credentials_base64` - (Required) Base64 encoding of `kube_config_raw` of Azure Kubernetes Service cluster. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The name of the SKU for DevSpace Controller. Currently the only supported value is `S1` . Changing this forces a new resource to be created.
- `tier` - (Required) The tier of the SKU for DevSpace Controller. Currently the only supported value is `Standard` . Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the DevSpace Controller.
- `data_plane_fqdn` - DNS name for accessing DataPlane services.
- `host_suffix` - The host suffix for the DevSpace Controller.

Import

DevSpace Controller's can be imported using the `resource id` , e.g.

```
terraform import azurerm_devspace_controller.controller1 /subscriptions/00000000-0000-0000-0000-000000000000/000/resourceGroups/group1/providers/Microsoft.DevSpaces/controllers/controller1Name
```

azurerm_dev_test_lab

Manages a Dev Test Lab.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "test" {
  name                = "example-devtestlab"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    "Sydney" = "Australia"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Dev Test Lab. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group under which the Dev Test Lab resource has to be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Dev Test Lab should exist. Changing this forces a new resource to be created.
- `storage_type` - (Optional) The type of storage used by the Dev Test Lab. Possible values are `Standard` and `Premium`. Defaults to `Premium`. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Dev Test Lab.
- `artifacts_storage_account_id` - The ID of the Storage Account used for Artifact Storage.
- `default_storage_account_id` - The ID of the Default Storage Account for this Dev Test Lab.
- `default_premium_storage_account_id` - The ID of the Default Premium Storage Account for this Dev Test Lab.

- `key_vault_id` - The ID of the Key used for this Dev Test Lab.
- `premium_data_disk_storage_account_id` - The ID of the Storage Account used for Storage of Premium Data Disk.
- `unique_identifier` - The unique immutable identifier of the Dev Test Lab.

Import

Dev Test Labs can be imported using the `resource_id`, e.g.

```
terraform import azurerm_dev_test_lab.lab1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DevTestLab/labs/lab1
```

azurerm_dev_test_linux_virtual_machine

Manages a Linux Virtual Machine within a Dev Test Lab.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "test" {
  name                = "example-devtestlab"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    "Sydney" = "Australia"
  }
}

resource "azurerm_dev_test_virtual_network" "test" {
  name                = "example-network"
  lab_name            = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  subnet {
    use_public_ip_address      = "Allow"
    use_in_virtual_machine_creation = "Allow"
  }
}

resource "azurerm_dev_test_linux_virtual_machine" "test" {
  name                = "example-vm03"
  lab_name            = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  size               = "Standard_DS2"
  username            = "exampleuser99"
  ssh_key             = "${file("~/ssh/id_rsa.pub")}"
  lab_virtual_network_id = "${azurerm_dev_test_virtual_network.test.id}"
  lab_subnet_name     = "${azurerm_dev_test_virtual_network.test.subnet.0.name}"
  storage_type        = "Premium"
  notes               = "Some notes about this Virtual Machine."

  gallery_image_reference {
    offer      = "UbuntuServer"
    publisher  = "Canonical"
    sku        = "18.04-LTS"
    version    = "latest"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Dev Test Machine. Changing this forces a new resource to be created.

NOTE: The validation requirements for the Name change based on the `os_type` used in this Virtual Machine. For a Linux VM the name must be between 1-62 characters, and for a Windows VM the name must be between 1-15 characters. It must begin and end with a letter or number, and cannot be all numbers.

- `lab_name` - (Required) Specifies the name of the Dev Test Lab in which the Virtual Machine should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Dev Test Lab resource exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Dev Test Lab exists. Changing this forces a new resource to be created.
- `gallery_image_reference` - (Required) A `gallery_image_reference` block as defined below.
- `lab_subnet_name` - (Required) The name of a Subnet within the Dev Test Virtual Network where this machine should exist. Changing this forces a new resource to be created.
- `lab_virtual_network_id` - (Required) The ID of the Dev Test Virtual Network where this Virtual Machine should be created. Changing this forces a new resource to be created.
- `size` - (Required) The Machine Size to use for this Virtual Machine, such as `Standard_F2` . Changing this forces a new resource to be created.
- `storage_type` - (Required) The type of Storage to use on this Virtual Machine. Possible values are `Standard` and `Premium` .
- `username` - (Required) The Username associated with the local administrator on this Virtual Machine. Changing this forces a new resource to be created.

- `allow_claim` - (Optional) Can this Virtual Machine be claimed by users? Defaults to `true` .
- `disallow_public_ip_address` - (Optional) Should the Virtual Machine be created without a Public IP Address? Changing this forces a new resource to be created.
- `inbound_nat_rule` - (Optional) One or more `inbound_nat_rule` blocks as defined below. Changing this forces a new resource to be created.

NOTE: If any `inbound_nat_rule` blocks are specified then `disallow_public_ip_address` must be set to `true` .

- `notes` - (Optional) Any notes about the Virtual Machine.
- `password` - (Optional) The Password associated with the `username` used to login to this Virtual Machine. Changing this forces a new resource to be created.
- `ssh_key` - (Optional) The SSH Key associated with the `username` used to login to this Virtual Machine. Changing this forces a new resource to be created.

NOTE: One or either `password` or `ssh_key` must be specified.

- `tags` - (Optional) A mapping of tags to assign to the resource.

A `gallery_image_reference` block supports the following:

- `offer` - (Required) The Offer of the Gallery Image. Changing this forces a new resource to be created.
- `publisher` - (Required) The Publisher of the Gallery Image. Changing this forces a new resource to be created.
- `sku` - (Required) The SKU of the Gallery Image. Changing this forces a new resource to be created.
- `version` - (Required) The Version of the Gallery Image. Changing this forces a new resource to be created.

A `inbound_nat_rule` block supports the following:

- `protocol` - (Required) The Protocol used for this NAT Rule. Possible values are `Tcp` and `Udp`. Changing this forces a new resource to be created.
- `backend_port` - (Required) The Backend Port associated with this NAT Rule. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual Machine.
- `fqdn` - The FQDN of the Virtual Machine.
- `inbound_nat_rule` - One or more `inbound_nat_rule` blocks as defined below.
- `unique_identifier` - The unique immutable identifier of the Virtual Machine.

A `inbound_nat_rule` block exports the following:

- `frontend_port` - The frontend port associated with this Inbound NAT Rule.

Import

Dev Test Linux Virtual Machines can be imported using the `resource id`, e.g.

```
terraform import azurerm_dev_test_linux_virtual_machine.machine1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DevTestLab/labs/lab1/virtualmachines/machine1
```

azurerm_dev_test_policy

Manages a Policy within a Dev Test Policy Set.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "test" {
  name                = "example-devtestlab"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    "Sydney" = "Australia"
  }
}

resource "azurerm_dev_test_policy" "test" {
  name                = "LabVmCount"
  policy_set_name     = "default"
  lab_name            = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  fact_data           = ""
  threshold           = "999"
  evaluator_type       = "MaxValuePolicy"

  tags = {
    "Acceptance" = "Test"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Dev Test Policy. Possible values are `GalleryImage` , `LabPremiumVmCount` , `LabTargetCost` , `LabVmCount` , `LabVmSize` , `UserOwnedLabPremiumVmCount` , `UserOwnedLabVmCount` and `UserOwnedLabVmCountInSubnet` . Changing this forces a new resource to be created.
- `policy_set_name` - (Required) Specifies the name of the Policy Set within the Dev Test Lab where this policy should be created. Changing this forces a new resource to be created.
- `lab_name` - (Required) Specifies the name of the Dev Test Lab in which the Policy should be created. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which the Dev Test Lab resource exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Dev Test Lab exists. Changing this forces a new resource to be created.
- `description` - (Optional) A description for the Policy.
- `evaluator_type` - (Required) The Evaluation Type used for this Policy. Possible values include: 'AllowedValuesPolicy', 'MaxValuePolicy'. Changing this forces a new resource to be created.
- `threshold` - (Required) The Threshold for this Policy.
- `fact_data` - (Optional) The Fact Data for this Policy.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Dev Test Policy.

Import

Dev Test Policies can be imported using the `resource id`, e.g.

```
terraform import azurerm_dev_test_policy.policy1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DevTestLab/labs/lab1/policysets/default/policies/policy1
```

azurerm_dev_test_schedule

Manages automated startup and shutdown schedules for Azure Dev Test Lab.

Example Usage

```
resource "azurerm_resource_group" "sample" {
  name      = "acctestRG"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "sample" {
  name                = "YourDevTestLab"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dev_test_schedule" "sample" {
  name                = "LabVmAutoStart"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  lab_name            = "${azurerm_dev_test_lab.test.name}"

  weekly_recurrence {
    time      = "1100"
    week_days = ["Monday", "Tuesday"]
  }

  time_zone_id = "Pacific Standard Time"
  task_type    = "LabVmsStartupTask"

  notification_settings {
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the dev test lab schedule. Valid value for name depends on the `task_type`. For instance for `task_type LabVmsStartupTask` the name needs to be `LabVmAutoStart`.
- `location` - (Required) The location where the schedule is created. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the dev test lab schedule. Changing this forces a new resource to be created.
 - `lab_name` - (Required) The name of the dev test lab. Changing this forces a new resource to be created.
 - `status` - The status of this schedule. Possible values are `Enabled` and `Disabled`. Defaults to `Disabled`.
 - `task_type` - (Required) The task type of the schedule. Possible values include `LabVmsShutdownTask` and `LabVmAutoStart`.
 - `time_zone_id` - (Required) The time zone ID (e.g. Pacific Standard time).
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

A `weekly_recurrence` - block supports the following:

- `time` - The time when the schedule takes effect.
 - `week_days` - A list of days that this schedule takes effect. Possible values include `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday` and `Sunday`.
-

A `daily_recurrence` - block supports the following:

- `time` - The time each day when the schedule takes effect.
-

A `notification_settings` - (Required) - block supports the following:

- `status` - The status of the notification. Possible values are `Enabled` and `Disabled`. Defaults to `Disabled`.
- `time_in_minutes` - Time in minutes before event at which notification will be sent.
- `webhook_url` - The webhook URL to which the notification will be sent.

Attributes Reference

The following attributes are exported:

- `id` - The Dev Test Lab Schedule ID.

Import

Dev Test Schedule can be imported using the `resource id`, e.g.

```
terraform import azurerm_dev_test_schedule.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DevTestLab/labs/myDevTestLab/schedules/labvmautostart
```

azurerm_dev_test_virtual_network

Manages a Virtual Network within a Dev Test Lab.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "test" {
  name                = "example-devtestlab"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    "Sydney" = "Australia"
  }
}

resource "azurerm_dev_test_virtual_network" "test" {
  name                = "example-network"
  lab_name             = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  subnet {
    use_public_ip_address      = "Allow"
    use_in_virtual_machine_creation = "Allow"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Dev Test Virtual Network. Changing this forces a new resource to be created.
 - `lab_name` - (Required) Specifies the name of the Dev Test Lab in which the Virtual Network should be created. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the resource group in which the Dev Test Lab resource exists. Changing this forces a new resource to be created.
 - `description` - (Optional) A description for the Virtual Network.
 - `subnet` - (Optional) A `subnet` block as defined below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

A `subnet` block supports the following:

- `use_public_ip_address` - (Required) Can Virtual Machines in this Subnet use Public IP Addresses? Possible values are `Allow`, `Default` and `Deny`.
- `use_in_virtual_machine_creation` - (Required) Can this subnet be used for creating Virtual Machines? Possible values are `Allow`, `Default` and `Deny`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Dev Test Virtual Network.
- `subnet` - A `subnet` block as defined below.
- `unique_identifier` - The unique immutable identifier of the Dev Test Virtual Network.

A `subnet` block exports the following:

- `name` - The name of the Subnet for this Virtual Network.

Import

Dev Test Virtual Networks can be imported using the `resource id`, e.g.

```
terraform import azurerm_dev_test_virtual_network.network1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DevTestLab/labs/lab1/virtualnetworks/network1
```

azurerm_dev_test_windows_virtual_machine

Manages a Windows Virtual Machine within a Dev Test Lab.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_dev_test_lab" "test" {
  name                = "example-devtestlab"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    "Sydney" = "Australia"
  }
}

resource "azurerm_dev_test_virtual_network" "test" {
  name                = "example-network"
  lab_name             = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  subnet {
    use_public_ip_address      = "Allow"
    use_in_virtual_machine_creation = "Allow"
  }
}

resource "azurerm_dev_test_windows_virtual_machine" "test" {
  name                = "example-vm03"
  lab_name            = "${azurerm_dev_test_lab.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  size                = "Standard_DS2"
  username            = "exampleuser99"
  password            = "Pa$$w0rd1234!"
  lab_virtual_network_id = "${azurerm_dev_test_virtual_network.test.id}"
  lab_subnet_name      = "${azurerm_dev_test_virtual_network.test.subnet.0.name}"
  storage_type         = "Premium"
  notes               = "Some notes about this Virtual Machine."

  gallery_image_reference {
    offer      = "UbuntuServer"
    publisher  = "Canonical"
    sku        = "18.04-LTS"
    version    = "latest"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Dev Test Machine. Changing this forces a new resource to be created.

NOTE: The validation requirements for the Name change based on the `os_type` used in this Virtual Machine. For a Linux VM the name must be between 1-62 characters, and for a Windows VM the name must be between 1-15 characters. It must begin and end with a letter or number, and cannot be all numbers.

- `lab_name` - (Required) Specifies the name of the Dev Test Lab in which the Virtual Machine should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Dev Test Lab resource exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Dev Test Lab exists. Changing this forces a new resource to be created.
- `gallery_image_reference` - (Required) A `gallery_image_reference` block as defined below.
- `lab_subnet_name` - (Required) The name of a Subnet within the Dev Test Virtual Network where this machine should exist. Changing this forces a new resource to be created.
- `lab_virtual_network_id` - (Required) The ID of the Dev Test Virtual Network where this Virtual Machine should be created. Changing this forces a new resource to be created.
- `password` - (Required) The Password associated with the `username` used to login to this Virtual Machine. Changing this forces a new resource to be created.
- `size` - (Required) The Machine Size to use for this Virtual Machine, such as `Standard_F2` . Changing this forces a new resource to be created.
- `storage_type` - (Required) The type of Storage to use on this Virtual Machine. Possible values are `Standard` and `Premium` .
- `username` - (Required) The Username associated with the local administrator on this Virtual Machine. Changing this forces a new resource to be created.

- `allow_claim` - (Optional) Can this Virtual Machine be claimed by users? Defaults to `true` .
- `disallow_public_ip_address` - (Optional) Should the Virtual Machine be created without a Public IP Address? Changing this forces a new resource to be created.
- `inbound_nat_rule` - (Optional) One or more `inbound_nat_rule` blocks as defined below. Changing this forces a new resource to be created.

NOTE: If any `inbound_nat_rule` blocks are specified then `disallow_public_ip_address` must be set to `true` .

- `notes` - (Optional) Any notes about the Virtual Machine.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

A `gallery_image_reference` block supports the following:

- `offer` - (Required) The Offer of the Gallery Image. Changing this forces a new resource to be created.
 - `publisher` - (Required) The Publisher of the Gallery Image. Changing this forces a new resource to be created.
 - `sku` - (Required) The SKU of the Gallery Image. Changing this forces a new resource to be created.
 - `version` - (Required) The Version of the Gallery Image. Changing this forces a new resource to be created.
-

A `inbound_nat_rule` block supports the following:

- `protocol` - (Required) The Protocol used for this NAT Rule. Possible values are `Tcp` and `Udp`. Changing this forces a new resource to be created.
- `backend_port` - (Required) The Backend Port associated with this NAT Rule. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual Machine.
 - `fqdn` - The FQDN of the Virtual Machine.
 - `inbound_nat_rule` - One or more `inbound_nat_rule` blocks as defined below.
 - `unique_identifier` - The unique immutable identifier of the Virtual Machine.
-

A `inbound_nat_rule` block exports the following:

- `frontend_port` - The frontend port associated with this Inbound NAT Rule.

Import

Dev Test Windows Virtual Machines can be imported using the `resource id`, e.g.

```
terraform import azurerm_dev_test_windows_virtual_machine.machine1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.DevTestLab/labs/lab1/virtualmachines/machine1
```


azurerm_dns_aaaa_record

Enables you to manage DNS AAAA Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_aaaa_record" "test" {
  name           = "test"
  zone_name      = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl            = 300
  records        = ["2607:f8b0:4009:1803::1005"]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS AAAA Record.
- `resource_group_name` - (Required) Specifies the resource group where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the resource exists. Changing this forces a new resource to be created.
- `TTL` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `records` - (Required) List of IPv6 Addresses.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The DNS AAAA Record ID.

Import

AAAA records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_dns_aaaa_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1/AAAA/myrecord1
```

azurerm_dns_a_record

Enables you to manage DNS A Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_dns_zone" "test" {
  name            = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_a_record" "test" {
  name            = "test"
  zone_name       = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl             = 300
  records         = ["10.0.180.17"]
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the DNS A Record.
- **resource_group_name** - (Required) Specifies the resource group where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- **zone_name** - (Required) Specifies the DNS Zone where the resource exists. Changing this forces a new resource to be created.
- **TTL** - (Required) The Time To Live (TTL) of the DNS record in seconds.
- **records** - (Required) List of IPv4 Addresses.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- **id** - The DNS A Record ID.

Import

A records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_dns_a_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1/A/myrecord1
```

azurerm_dns_caa_record

Enables you to manage DNS CAA Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_caa_record" "test" {
  name                = "test"
  zone_name           = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl                 = 300

  record {
    flags = 0
    tag   = "issue"
    value = "example.com"
  }

  record {
    flags = 0
    tag   = "issue"
    value = "example.net"
  }

  record {
    flags = 0
    tag   = "issuewild"
    value = ";"
  }

  record {
    flags = 0
    tag   = "iodef"
    value = "mailto:terraform@nonexisting.tld"
  }

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS CAA Record.
- `resource_group_name` - (Required) Specifies the resource group where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the resource exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `record` - (Required) A list of values that make up the CAA record. Each `record` block supports fields documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `record` block supports:

- `flags` - (Required) Extensible CAA flags, currently only 1 is implemented to set the issuer critical flag.
- `tag` - (Required) A property tag, options are `issue`, `issuewild` and `iodef`.
- `value` - (Required) A property value such as a registrar domain.

Attributes Reference

The following attributes are exported:

- `id` - The DNS CAA Record ID.

Import

CAA records can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_caa_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resource
Groups/mygroup1/providers/Microsoft.Network/dnszones/zone1/CAA/myrecord1
```

azurerm_dns_cname_record

Enables you to manage DNS CNAME Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_cname_record" "test" {
  name           = "test"
  zone_name      = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl            = 300
  record         = "contoso.com"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the DNS CNAME Record.
- **resource_group_name** - (Required) Specifies the resource group where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- **zone_name** - (Required) Specifies the DNS Zone where the resource exists. Changing this forces a new resource to be created.
- **TTL** - (Required) The Time To Live (TTL) of the DNS record in seconds.
- **record** - (Required) The target of the CNAME.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- **id** - The DNS CName Record ID.

Import

CNAME records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_dns_cname_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1/CNAME/myrecord1
```


azurerm_dns_mx_record

Enables you to manage DNS MX Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_mx_record" "test" {
  name                = "test"
  zone_name           = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl                 = 300

  record {
    preference = 10
    exchange   = "mail1.contoso.com"
  }

  record {
    preference = 20
    exchange   = "mail2.contoso.com"
  }

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS MX Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.

- `record` - (Required) A list of values that make up the MX record. Each `record` block supports fields documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `record` block supports:

- `preference` - (Required) String representing the "preference" value of the MX records. Records with lower preference value take priority.
- `exchange` - (Required) The mail server responsible for the domain covered by the MX record.

Attributes Reference

The following attributes are exported:

- `id` - The DNS MX Record ID.

Import

MX records can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_mx_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1/MX/myrecord1
```

azurerm_dns_ns_record

Enables you to manage DNS NS Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_ns_record" "test" {
  name           = "test"
  zone_name      = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl            = 300

  records = ["ns1.contoso.com", "ns2.contoso.com"]

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS NS Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `records` - (Optional) A list of values that make up the NS record. *WARNING:* Either `records` or `record` is required.
- `record` - (Optional) A list of values that make up the NS record. Each `record` block supports fields documented below. This field has been deprecated and will be removed in a future release.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `record` block supports:

- `nsdname` - (Required) The value of the record.

Attributes Reference

The following attributes are exported:

- `id` - The DNS NS Record ID.

Import

NS records can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_ns_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1/NS/myrecord1
```

azurerm_dns_ptr_record

Enables you to manage DNS PTR Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_ptr_record" "test" {
  name              = "test"
  zone_name         = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl               = 300
  records           = ["yourdomain.com"]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS PTR Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `records` - (Required) List of Fully Qualified Domain Names.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The DNS PTR Record ID.

Import

PTR records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_dns_ptr_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resource  
Groups/mygroup1/providers/Microsoft.Network/dnszones/zone1/PTR/myrecord1
```

azurerm_dns_srv_record

Enables you to manage DNS SRV Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_srv_record" "test" {
  name                = "test"
  zone_name           = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl                 = 300

  record {
    priority = 1
    weight   = 5
    port     = 8080
    target    = "target1.contoso.com"
  }

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS SRV Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `record` - (Required) A list of values that make up the SRV record. Each `record` block supports fields documented below.

- `tags` - (Optional) A mapping of tags to assign to the resource.

The `record` block supports:

- `priority` - (Required) Priority of the SRV record.
- `weight` - (Required) Weight of the SRV record.
- `port` - (Required) Port the service is listening on.
- `target` - (Required) FQDN of the service.

Attributes Reference

The following attributes are exported:

- `id` - The DNS SRV Record ID.

Import

SRV records can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_srv_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resource
Groups/mygroup1/providers/Microsoft.Network/dnszones/zone1/SRV/myrecord1
```


azurerm_dns_txt_record

Enables you to manage DNS TXT Records within Azure DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_dns_txt_record" "test" {
  name                = "test"
  zone_name           = "${azurerm_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl                 = 300

  record {
    value = "google-site-authenticator"
  }

  record {
    value = "more site information here"
  }

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS TXT Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the DNS Zone where the DNS Zone (parent resource) exists. Changing this forces a new resource to be created.
- `ttl` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `record` - (Required) A list of values that make up the txt record. Each `record` block supports fields documented below.

- `tags` - (Optional) A mapping of tags to assign to the resource.

The `record` block supports:

- `value` - (Required) The value of the record.

Attributes Reference

The following attributes are exported:

- `id` - The DNS TXT Record ID.

Import

TXT records can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_txt_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resource
Groups/mygroup1/providers/Microsoft.Network/dnszones/zone1/TXT/myrecord1
```

azurerm_dns_zone

Enables you to manage DNS zones within Azure DNS. These zones are hosted on Azure's name servers to which you can delegate the zone from the parent domain.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
  zone_type           = "Public"
}

resource "azurerm_dns_zone" "test_private" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
  zone_type           = "Private"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS Zone. Must be a valid domain name.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `registration_virtual_network_ids` - (Optional) A list of Virtual Network ID's that register hostnames in this DNS zone. This field can only be set when `zone_type` is set to `Private`.
- `resolution_virtual_network_ids` - (Optional) A list of Virtual Network ID's that resolve records in this DNS zone. This field can only be set when `zone_type` is set to `Private`.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zone_type` - (Optional / **Deprecated**) Specifies the type of this DNS zone. Possible values are `Public` or `Private` (Defaults to `Public`).

NOTE: This field was part of the initial Preview for Private DNS Zones - which has been replaced by the separate resource `azurerm_private_dns_zone` ([/docs/providers/azurerm/r/private_dns_zone.html](/docs/providers/azurerm/r/private_dns_zone.html)) and will be removed in v2.0 of the Azure Provider.

Attributes Reference

The following attributes are exported:

- `id` - The DNS Zone ID.
- `max_number_of_record_sets` - (Optional) Maximum number of Records in the zone. Defaults to `1000`.
- `number_of_record_sets` - (Optional) The number of records already in the zone.
- `name_servers` - (Optional) A list of values that make up the NS record for the zone.

Import

DNS Zones can be imported using the `resource id`, e.g.

```
terraform import azurerm_dns_zone.zone1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/dnszones/zone1
```

azurerm_eventgrid_domain

Manages an EventGrid Domain

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "resourceGroup1"
  location = "West US 2"
}

resource "azurerm_eventgrid_domain" "test" {
  name                = "my-eventgrid-domain"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the EventGrid Domain resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the EventGrid Domain exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `input_schema` - (Optional) Specifies the schema in which incoming events will be published to this domain. Allowed values are `cloudeventv01schema`, `customeventschema`, or `eventgridschema`. Defaults to `eventgridschema`. Changing this forces a new resource to be created.
- `input_mapping_fields` - (Optional) A `input_mapping_fields` block as defined below.
- `input_mapping_default_values` - (Optional) A `input_mapping_default_values` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `input_mapping_fields` supports the following:

- `id` - (Optional) Specifies the id of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.

- `topic` - (Optional) Specifies the topic of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
 - `event_type` - (Optional) Specifies the event type of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
 - `event_time` - (Optional) Specifies the event time of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
 - `data_version` - (Optional) Specifies the data version of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
 - `subject` - (Optional) Specifies the subject of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
-

A `input_mapping_default_values` supports the following:

- `event_type` - (Optional) Specifies the default event type of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
- `data_version` - (Optional) Specifies the default data version of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.
- `subject` - (Optional) Specifies the default subject of the EventGrid Event to associate with the domain. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the EventGrid Domain.
- `endpoint` - The Endpoint associated with the EventGrid Domain.

Import

EventGrid Domain's can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventgrid_domain.domain1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventGrid/domains/domain1
```

azurerm_eventgrid_event_subscription

Manages an EventGrid Event Subscription

Example Usage

```
resource "azurerm_resource_group" "default" {
  name     = "defaultResourceGroup"
  location = "West US 2"
}

resource "azurerm_storage_account" "default" {
  name                        = "defaultStorageAccount"
  resource_group_name        = "${azurerm_resource_group.default.name}"
  location                   = "${azurerm_resource_group.default.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"

  tags = {
    environment = "staging"
  }
}

resource "azurerm_storage_queue" "default" {
  name                = "defaultStorageQueue"
  resource_group_name = "${azurerm_resource_group.default.name}"
  storage_account_name = "${azurerm_storage_account.default.name}"
}

resource "azurerm_eventgrid_event_subscription" "default" {
  name = "defaultEventSubscription"
  scope = "${azurerm_resource_group.default.id}"

  storage_queue_endpoint {
    storage_account_id = "${azurerm_storage_account.default.id}"
    queue_name         = "${azurerm_storage_queue.default.name}"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the EventGrid Event Subscription resource. Changing this forces a new resource to be created.
- **scope** - (Required) Specifies the scope at which the EventGrid Event Subscription should be created. Changing this forces a new resource to be created.
- **event_delivery_schema** - (Optional) Specifies the event delivery schema for the event subscription. Possible values include: `EventGridSchema`, `CloudEventV01Schema`, `CustomInputSchema`.

- `topic_name` - (Optional) Specifies the name of the topic to associate with the event subscription.
- `storage_queue_endpoint` - (Optional) A `storage_queue_endpoint` block as defined below.
- `eventhub_endpoint` - (Optional) A `eventhub_endpoint` block as defined below.
- `hybrid_connection_endpoint` - (Optional) A `hybrid_connection_endpoint` block as defined below.
- `webhook_endpoint` - (Optional) A `webhook_endpoint` block as defined below.

NOTE: One of `storage_queue_endpoint`, `eventhub_endpoint`, `hybrid_connection_endpoint` or `webhook_endpoint` must be specified.

- `included_event_types` - (Optional) A list of applicable event types that need to be part of the event subscription.
- `subject_filter` - (Optional) A `subject_filter` block as defined below.
- `storage_blob_dead_letter_destination` - (Optional) A `storage_blob_dead_letter_destination` block as defined below.
- `retry_policy` - (Optional) A `retry_policy` block as defined below.
- `labels` - (Optional) A list of labels to assign to the event subscription.

A `storage_queue_endpoint` supports the following:

- `storage_account_id` - (Required) Specifies the id of the storage account id where the storage queue is located.
- `queue_name` - (Required) Specifies the name of the storage queue where the Event Subscription will receive events.

A `eventhub_endpoint` supports the following:

- `eventhub_id` - (Required) Specifies the id of the eventhub where the Event Subscription will receive events.

A `hybrid_connection_endpoint` supports the following:

- `hybrid_connection_id` - (Required) Specifies the id of the hybrid connection where the Event Subscription will receive events.

A `webhook_endpoint` supports the following:

- `url` - (Required) Specifies the url of the webhook where the Event Subscription will receive events.

A `subject_filter` supports the following:

- `subject_begins_with` - (Optional) A string to filter events for an event subscription based on a resource path prefix.
- `subject_ends_with` - (Optional) A string to filter events for an event subscription based on a resource path suffix.
- `case_sensitive` - (Optional) Specifies if `subject_begins_with` and `subject_ends_with` case sensitive. This value defaults to `false`.

A `storage_blob_dead_letter_destination` supports the following:

- `storage_account_id` - (Required) Specifies the id of the storage account id where the storage blob is located.
 - `storage_blob_container_name` - (Required) Specifies the name of the Storage blob container that is the destination of the deadletter events
-

A `retry_policy` supports the following:

- `max_delivery_attempts` - (Required) Specifies the maximum number of delivery retry attempts for events.
- `event_time_to_live` - (Required) Specifies the time to live (in minutes) for events.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the EventGrid Event Subscription.

Import

EventGrid Domain's can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventgrid_event_subscription.eventSubscription1
/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventGrid/eventSubscriptions/eventSubscription1
```

azurerm_eventgrid_topic

Manages an EventGrid Topic

Note: at this time EventGrid Topic's are only available in a limited number of regions.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US 2"
}

resource "azurerm_eventgrid_topic" "test" {
  name                = "my-eventgrid-topic"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the EventGrid Topic resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the EventGrid Topic exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The EventGrid Topic ID.
- `endpoint` - The Endpoint associated with the EventGrid Topic.
- `primary_access_key` - The Primary Shared Access Key associated with the EventGrid Topic.

- `secondary_access_key` - The Secondary Shared Access Key associated with the EventGrid Topic.

Import

EventGrid Topic's can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventgrid_topic.topic1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventGrid/topics/topic1
```

azurerm_eventhub_authorization_rule

Manages a Event Hubs authorization Rule within an Event Hub.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_eventhub_namespace" "test" {
  name                = "acceptanceTestEventHubNamespace"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Basic"
  capacity            = 2

  tags = {
    environment = "Production"
  }
}

resource "azurerm_eventhub" "test" {
  name                = "acceptanceTestEventHub"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  partition_count     = 2
  message_retention   = 2
}

resource "azurerm_eventhub_authorization_rule" "test" {
  name                = "navi"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  eventhub_name       = "${azurerm_eventhub.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  listen              = true
  send                = false
  manage              = false
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the EventHub Authorization Rule resource. Changing this forces a new resource to be created.
- **namespace_name** - (Required) Specifies the name of the grandparent EventHub Namespace. Changing this forces a new resource to be created.

- `eventhub_name` - (Required) Specifies the name of the EventHub. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the EventHub Namespace exists. Changing this forces a new resource to be created.

NOTE At least one of the 3 permissions below needs to be set.

- `listen` - (Optional) Does this Authorization Rule have permissions to Listen to the Event Hub? Defaults to `false`.
- `send` - (Optional) Does this Authorization Rule have permissions to Send to the Event Hub? Defaults to `false`.
- `manage` - (Optional) Does this Authorization Rule have permissions to Manage to the Event Hub? When this property is `true` - both `listen` and `send` must be too. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub ID.
- `primary_key` - The Primary Key for the Event Hubs authorization Rule.
- `primary_connection_string` - The Primary Connection String for the Event Hubs authorization Rule.
- `secondary_key` - The Secondary Key for the Event Hubs authorization Rule.
- `secondary_connection_string` - The Secondary Connection String for the Event Hubs authorization Rule.

Import

EventHubs can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventhub_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1/eventhubs/eventhub1/authorizationRules/rule1
```

azurerm_eventhub_consumer_group

Manages a Event Hubs Consumer Group as a nested resource within an Event Hub.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_eventhub_namespace" "test" {
  name                = "acceptanceTestEventHubNamespace"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Basic"
  capacity            = 2

  tags = {
    environment = "Production"
  }
}

resource "azurerm_eventhub" "test" {
  name                = "acceptanceTestEventHub"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  partition_count     = 2
  message_retention   = 2
}

resource "azurerm_eventhub_consumer_group" "test" {
  name                = "acceptanceTestEventHubConsumerGroup"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  eventhub_name       = "${azurerm_eventhub.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  user_metadata       = "some-meta-data"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the EventHub Consumer Group resource. Changing this forces a new resource to be created.
- **namespace_name** - (Required) Specifies the name of the grandparent EventHub Namespace. Changing this forces a new resource to be created.
- **eventhub_name** - (Required) Specifies the name of the EventHub. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which the EventHub Consumer Group's grandparent Namespace exists. Changing this forces a new resource to be created.
- `user_metadata` - (Optional) Specifies the user metadata.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub Consumer Group ID.

Import

EventHub Consumer Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventhub_consumer_group.consumerGroup1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1/eventhubs/eventhub1/consumergroups/consumerGroup1
```

azurerm_eventhub

Manages a Event Hubs as a nested resource within a Event Hubs namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_eventhub_namespace" "test" {
  name                = "acceptanceTestEventHubNamespace"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Standard"
  capacity            = 1
  kafka_enabled       = false

  tags = {
    environment = "Production"
  }
}

resource "azurerm_eventhub" "test" {
  name                = "acceptanceTestEventHub"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  partition_count     = 2
  message_retention   = 1
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the EventHub Namespace resource. Changing this forces a new resource to be created.
- **namespace_name** - (Required) Specifies the name of the EventHub Namespace. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the EventHub's parent Namespace exists. Changing this forces a new resource to be created.
- **partition_count** - (Required) Specifies the current number of shards on the Event Hub. Changing this forces a new resource to be created.
- **message_retention** - (Required) Specifies the number of days to retain the events for this Event Hub. Needs to be between 1 and 7 days; or 1 day when using a Basic SKU for the parent EventHub Namespace.

- `capture_description` - (Optional) A `capture_description` block as defined below.

A `capture_description` block supports the following:

- `enabled` - (Required) Specifies if the Capture Description is Enabled.
- `encoding` - (Required) Specifies the Encoding used for the Capture Description. Possible values are `Avro` and `AvroDeflate`.
- `interval_in_seconds` - (Optional) Specifies the time interval in seconds at which the capture will happen. Values can be between `60` and `900` seconds. Defaults to `300` seconds.
- `size_limit_in_bytes` - (Optional) Specifies the amount of data built up in your EventHub before a Capture Operation occurs. Value should be between `10485760` and `524288000` bytes. Defaults to `314572800` bytes.
- `skip_empty_archives` - (Optional) Specifies if empty files should not be emitted if no events occur during the Capture time window. Defaults to `false`.
- `destination` - (Required) A `destination` block as defined below.

A `destination` block supports the following:

- `name` - (Required) The Name of the Destination where the capture should take place. At this time the only supported value is `EventHubArchive.AzureBlockBlob`.

At this time it's only possible to Capture EventHub messages to Blob Storage. There's a Feature Request for the Azure SDK to add support for Capturing messages to Azure Data Lake here (<https://github.com/Azure/azure-rest-api-specs/issues/2255>).

- `archive_name_format` - The Blob naming convention for archiving. e.g. `{Namespace}/{EventHub}/{PartitionId}/{Year}/{Month}/{Day}/{Hour}/{Minute}/{Second}`. Here all the parameters (Namespace, EventHub .. etc) are mandatory irrespective of order
- `blob_container_name` - (Required) The name of the Container within the Blob Storage Account where messages should be archived.
- `storage_account_id` - (Required) The ID of the Blob Storage Account where messages should be archived.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub ID.
- `partition_ids` - The identifiers for partitions created for Event Hubs.

Import

EventHubs can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventhub.eventhub1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1/eventhubs/eventhub1
```

azurerm_eventhub_namespace_authorization_rule

Manages an Authorization Rule for an Event Hub Namespace.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "resourcegroup"
  location = "West US"
}

resource "azurerm_eventhub_namespace" "example" {
  name                = "acceptanceTestEventHubNamespace"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Basic"
  capacity             = 2

  tags = {
    environment = "Production"
  }
}

resource "azurerm_eventhub_namespace_authorization_rule" "example" {
  name                = "navi"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  listen = true
  send   = false
  manage = false
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Authorization Rule. Changing this forces a new resource to be created.
- `namespace_name` - (Required) Specifies the name of the EventHub Namespace. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the EventHub Namespace exists. Changing this forces a new resource to be created.

NOTE At least one of the 3 permissions below needs to be set.

- `listen` - (Optional) Grants listen access to this this Authorization Rule. Defaults to `false`.
- `send` - (Optional) Grants send access to this this Authorization Rule. Defaults to `false`.

- `manage` - (Optional) Grants manage access to this this Authorization Rule. When this property is `true` - both `listen` and `send` must be too. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub Namespace Authorization Rule ID.
- `primary_key` - The Primary Key for the Authorization Rule.
- `primary_connection_string` - The Primary Connection String for the Authorization Rule.
- `secondary_key` - The Secondary Key for the Authorization Rule.
- `secondary_connection_string` - The Secondary Connection String for the Authorization Rule.

Import

EventHubs can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventhub_namespace_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1/authorizationRules/rule1
```

azurerm_eventhub_namespace_disaster_recovery_config

Manages an Disaster Recovery Config for an Event Hub Namespace.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "eventhub-replication"
  location = "West Europe"
}

resource "azurerm_eventhub_namespace" "primary" {
  name            = "eventhub-primary"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"
}

resource "azurerm_eventhub_namespace" "secondary" {
  name            = "eventhub-secondary"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"
}

resource "azurerm_eventhub_namespace_disaster_recovery_config" "example" {
  name                = "replicate-eventhub"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_eventhub_namespace.primary.name}"
  partner_namespace_id = "${azurerm_eventhub_namespace.secondary.id}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Disaster Recovery Config. Changing this forces a new resource to be created.
- **namespace_name** - (Required) Specifies the name of the primary EventHub Namespace to replicate. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the Disaster Recovery Config exists. Changing this forces a new resource to be created.
- **partner_namespace_id** - (Optional) The ID of the EventHub Namespace to replicate to.
- **alternate_name** - (Optional) An alternate name to use when the Disaster Recovery Config's name is the same as the replicated namespace's name.
- **wait_for_replication** - (Optional) Should the resource wait for replication upon creation? Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub Namespace Disaster Recovery Config ID.

Import

EventHubs can be imported using the `resource id`, e.g.

```
terraform import azurerm_eventhub_namespace_disaster_recovery_config.config1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1/disasterRecoveryConfigs/config1
```

azurerm_eventhub_namespace

Manages an EventHub Namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_eventhub_namespace" "test" {
  name                = "example-namespace"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "Standard"
  capacity            = 2

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the EventHub Namespace resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the namespace. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) Defines which tier to use. Valid options are `Basic` and `Standard`.
- `capacity` - (Optional) Specifies the Capacity / Throughput Units for a `Standard` SKU namespace. Valid values range from `1` - `20`.
- `auto_inflate_enabled` - (Optional) Is Auto Inflate enabled for the EventHub Namespace?
- `maximum_throughput_units` - (Optional) Specifies the maximum number of throughput units when Auto Inflate is Enabled. Valid values range from `1` - `20`.
- `kafka_enabled` - (Optional / **Deprecated**) Is Kafka enabled for the EventHub Namespace? Defaults to `false`.

NOTE: `kafka_enabled` is now configured depending on the `sku` being provisioned, where this is Disabled for a `Basic` sku and Enabled for a `Standard` sku.

- `tags` - (Optional) A mapping of tags to assign to the resource.
- `network_rulesets` - (Optional) A `network_rulesets` block as defined below.

A `network_rulesets` block supports the following:

- `default_action` - (Required) The default action to take when a rule is not matched. Possible values are `Allow` and `Deny`.
- `virtual_network_rule` - (Optional) One or more `virtual_network_rule` blocks as defined below.
- `ip_rule` - (Optional) One or more `ip_rule` blocks as defined below.

A `virtual_network_rule` block supports the following:

- `subnet_id` - (Required) The id of the subnet to match on.
- `ignore_missing_virtual_network_service_endpoint` - (Optional) Are missing virtual network service endpoints ignored? Defaults to `false`.

A `ip_rule` block supports the following:

- `ip_mask` - (Required) The ip mask to match on.
- `action` - (Optional) The action to take when the rule is matched. Possible values are `Allow`.

Attributes Reference

The following attributes are exported:

- `id` - The EventHub Namespace ID.

The following attributes are exported only if there is an authorization rule named `RootManageSharedAccessKey` which is created automatically by Azure.

- `default_primary_connection_string` - The primary connection string for the authorization rule `RootManageSharedAccessKey`.
- `default_secondary_connection_string` - The secondary connection string for the authorization rule `RootManageSharedAccessKey`.
- `default_primary_key` - The primary access key for the authorization rule `RootManageSharedAccessKey`.
- `default_secondary_key` - The secondary access key for the authorization rule `RootManageSharedAccessKey`.

Import

EventHub Namespaces can be imported using the `resource id`, e.g.


```
terraform import azurerm_eventhub_namespace.namespace1 /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/group1/providers/Microsoft.EventHub/namespaces/namespace1
```

azurerm_express_route_circuit_authorization

Manages an ExpressRoute Circuit Authorization.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "exprtTest"
  location  = "West US"
}

resource "azurerm_express_route_circuit" "test" {
  name                        = "expressRoute1"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  service_provider_name      = "Equinix"
  peering_location           = "Silicon Valley"
  bandwidth_in_mbps          = 50

  sku {
    tier      = "Standard"
    family    = "MeteredData"
  }

  allow_classic_operations = false

  tags = {
    environment = "Production"
  }
}

resource "azurerm_express_route_circuit_authorization" "test" {
  name                        = "exampleERCAuth"
  express_route_circuit_name = "${azurerm_express_route_circuit.test.name}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the ExpressRoute circuit. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the ExpressRoute circuit. Changing this forces a new resource to be created.
- `express_route_circuit_name` - (Required) The name of the Express Route Circuit in which to create the Authorization.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the ExpressRoute Circuit Authorization.
- `authorization_key` - The Authorization Key.
- `authorization_use_status` - The authorization use status.

Import

ExpressRoute Circuit Authorizations can be imported using the `resource id`, e.g.

```
terraform import azurerm_express_route_circuit_authorization.auth1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/expressRouteCircuits/myExpressRoute/authorizations/auth1
```

azurerm_express_route_circuit

Manages an ExpressRoute circuit.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "exprtTest"
  location = "West US"
}

resource "azurerm_express_route_circuit" "test" {
  name                       = "expressRoute1"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  service_provider_name     = "Equinix"
  peering_location          = "Silicon Valley"
  bandwidth_in_mbps         = 50

  sku {
    tier     = "Standard"
    family  = "MeteredData"
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the ExpressRoute circuit. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the ExpressRoute circuit. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `service_provider_name` - (Required) The name of the ExpressRoute Service Provider.
- `peering_location` - (Required) The name of the peering location and **not** the Azure resource location.
- `bandwidth_in_mbps` - (Required) The bandwidth in Mbps of the circuit being created.

NOTE: Once you increase your bandwidth, you will not be able to decrease it to it's previous value.

- `sku` - (Required) A `sku` block for the ExpressRoute circuit as documented below.

- `allow_classic_operations` - (Optional) Allow the circuit to interact with classic (RDFE) resources. The default value is `false`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`sku` supports the following:

- `tier` - (Required) The service tier. Possible values are `Standard` or `Premium`.
- `family` - (Required) The billing mode for bandwidth. Possible values are `MeteredData` or `UnlimitedData`.

NOTE: You can migrate from `MeteredData` to `UnlimitedData`, but not the other way around.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the ExpressRoute circuit.
- `service_provider_provisioning_state` - The ExpressRoute circuit provisioning state from your chosen service provider. Possible values are "NotProvisioned", "Provisioning", "Provisioned", and "Deprovisioning".
- `service_key` - The string needed by the service provider to provision the ExpressRoute circuit.

Import

ExpressRoute circuits can be imported using the `resource id`, e.g.

```
terraform import azurerm_express_route_circuit.myExpressRoute /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/expressRouteCircuits/myExpressRoute
```

azurerm_express_route_circuit_peering

Manages an ExpressRoute Circuit Peering.

Example Usage (Creating a Microsoft Peering)

```
resource "azurerm_resource_group" "test" {
  name      = "expTest"
  location  = "West US"
}

resource "azurerm_express_route_circuit" "test" {
  name                        = "expressRoute1"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  service_provider_name      = "Equinix"
  peering_location           = "Silicon Valley"
  bandwidth_in_mbps          = 50

  sku {
    tier      = "Standard"
    family    = "MeteredData"
  }

  allow_classic_operations = false

  tags = {
    environment = "Production"
  }
}

resource "azurerm_express_route_circuit_peering" "test" {
  peering_type                = "MicrosoftPeering"
  express_route_circuit_name  = "${azurerm_express_route_circuit.test.name}"
  resource_group_name         = "${azurerm_resource_group.test.name}"
  peer_asn                    = 100
  primary_peer_address_prefix = "123.0.0.0/30"
  secondary_peer_address_prefix = "123.0.0.4/30"
  vlan_id                     = 300

  microsoft_peering_config {
    advertised_public_prefixes = ["123.1.0.0/24"]
  }
}
```

Argument Reference

The following arguments are supported:

- `peering_type` - (Required) The type of the ExpressRoute Circuit Peering. Acceptable values include `AzurePrivatePeering`, `AzurePublicPeering` and `MicrosoftPeering`. Changing this forces a new resource to be

created.

NOTE: only one Peering of each Type can be created. Attempting to create multiple peerings of the same type will overwrite the original peering.

- `express_route_circuit_name` - (Required) The name of the ExpressRoute Circuit in which to create the Peering.
- `resource_group_name` - (Required) The name of the resource group in which to create the Express Route Circuit Peering. Changing this forces a new resource to be created.
- `primary_peer_address_prefix` - (Optional) A /30 subnet for the primary link.
- `secondary_peer_address_prefix` - (Optional) A /30 subnet for the secondary link.
- `vlan_id` - (Optional) A valid VLAN ID to establish this peering on.
- `shared_key` - (Optional) The shared key. Can be a maximum of 25 characters.
- `peer_asn` - (Optional) The Either a 16-bit or a 32-bit ASN. Can either be public or private..
- `microsoft_peering_config` - (Optional) A `microsoft_peering_config` block as defined below. Required when `peering_type` is set to `MicrosoftPeering`.

A `microsoft_peering_config` block contains:

- `advertised_public_prefixes` - (Required) A list of Advertised Public Prefixes

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the ExpressRoute Circuit Peering.
- `azure_asn` - The ASN used by Azure.
- `primary_azure_port` - The Primary Port used by Azure for this Peering.
- `secondary_azure_port` - The Secondary Port used by Azure for this Peering.

Import

ExpressRoute Circuit Peerings can be imported using the `resource id`, e.g.

```
terraform import azurerm_express_route_circuit_peering.peering1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/expressRouteCircuits/myExpressRoute/peerings/peering1
```

azurerm_firewall_application_rule_collection

Manages an Application Rule Collection within an Azure Firewall.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "North Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "testvnet"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "AzureFirewallSubnet"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "test" {
  name             = "testpip"
  location         = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
  sku              = "Standard"
}

resource "azurerm_firewall" "test" {
  name            = "testfirewall"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name            = "configuration"
    subnet_id       = "${azurerm_subnet.test.id}"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_firewall_application_rule_collection" "test" {
  name            = "testcollection"
  azure_firewall_name = "${azurerm_firewall.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  priority        = 100
  action          = "Allow"

  rule {
    name = "testrule"
  }
}
```



```

source_addresses = [
    "10.0.0.0/16",
]

target_fqdns = [
    "*.google.com",
]

protocol {
    port = "443"
    type = "Https"
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Application Rule Collection which must be unique within the Firewall. Changing this forces a new resource to be created.
- `azure_firewall_name` - (Required) Specifies the name of the Firewall in which the Application Rule Collection should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which the Firewall exists. Changing this forces a new resource to be created.
- `priority` - (Required) Specifies the priority of the rule collection. Possible values are between `100` - `65000`.
- `action` - (Required) Specifies the action the rule will apply to matching traffic. Possible values are `Allow` and `Deny`.
- `rule` - (Required) One or more `rule` blocks as defined below.

A `rule` block supports the following:

- `name` - (Required) Specifies the name of the rule.
- `description` - (Optional) Specifies a description for the rule.
- `source_addresses` - (Required) A list of source IP addresses and/or IP ranges.
- `fqdn_tags` - (Optional) A list of FQDN tags. Possible values are `AppServiceEnvironment`, `AzureBackup`, `MicrosoftActiveProtectionService`, `WindowsDiagnostics` and `WindowsUpdate`.
- `target_fqdns` - (Optional) A list of FQDNs.
- `protocol` - (Optional) One or more `protocol` blocks as defined below.

A `protocol` block supports the following:

- `port` - (Optional) Specify a port for the connection.
- `type` - (Required) Specifies the type of connection. Possible values are `Http` or `Https`.

Import

Azure Firewall Application Rule Collections can be imported using the `resource_id`, e.g.

```
terraform import azurerm_firewall_application_rule_collection.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/azureFirewalls/myfirewall/applicationRuleCollections/mycollection
```

azurerm_firewall

Manages an Azure Firewall.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "North Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "testvnet"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "AzureFirewallSubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "test" {
  name             = "testpip"
  location         = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
  sku              = "Standard"
}

resource "azurerm_firewall" "test" {
  name            = "testfirewall"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name            = "configuration"
    subnet_id       = "${azurerm_subnet.test.id}"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Firewall. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the resource. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `ip_configuration` - (Required) A `ip_configuration` block as documented below.
- `zones` - (Optional) Specifies the availability zones in which the Azure Firewall should be created.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

- `tags` - (Optional) A mapping of tags to assign to the resource.

A `ip_configuration` block supports the following:

- `name` - (Required) Specifies the name of the IP Configuration.
- `subnet_id` - (Required) Reference to the subnet associated with the IP Configuration.

NOTE The Subnet used for the Firewall must have the name `AzureFirewallSubnet` and the subnet mask must be at least `/26`.

- `public_ip_address_id` - (Required) The Resource ID of the Public IP Address associated with the firewall.

NOTE The Public IP must have a `Static` allocation and `Standard` sku.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Azure Firewall.
- `ip_configuration` - A `ip_configuration` block as defined below.

A `ip_configuration` block exports the following:

- `private_ip_address` - The private IP address of the Azure Firewall.

Import

Azure Firewalls can be imported using the `resource id`, e.g.

```
terraform import azurerm_firewall.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/azureFirewalls/testfirewall
```

azurerm_firewall_nat_rule_collection

Manages a NAT Rule Collection within an Azure Firewall.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "North Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "testvnet"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "AzureFirewallSubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "test" {
  name             = "testpip"
  location         = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
  sku              = "Standard"
}

resource "azurerm_firewall" "test" {
  name            = "testfirewall"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name            = "configuration"
    subnet_id       = "${azurerm_subnet.test.id}"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_firewall_nat_rule_collection" "test" {
  name            = "testcollection"
  azure_firewall_name = "${azurerm_firewall.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  priority        = 100
  action          = "Dnat"

  rule {
    name = "testrule"
  }
}
```

```
source_addresses = [
    "10.0.0.0/16",
]

destination_ports = [
    "53",
]

destination_addresses = [
    "8.8.8.8",
    "8.8.4.4",
]

protocols = [
    "TCP",
    "UDP",
]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the NAT Rule Collection which must be unique within the Firewall. Changing this forces a new resource to be created.
- `azure_firewall_name` - (Required) Specifies the name of the Firewall in which the NAT Rule Collection should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which the Firewall exists. Changing this forces a new resource to be created.
- `priority` - (Required) Specifies the priority of the rule collection. Possible values are between `100` - `65000`.
- `action` - (Required) Specifies the action the rule will apply to matching traffic. Possible values are `Dnat` and `Snat`.
- `rule` - (Required) One or more `rule` blocks as defined below.

A `rule` block supports the following:

- `name` - (Required) Specifies the name of the rule.
- `description` - (Optional) Specifies a description for the rule.
- `destination_addresses` - (Required) A list of destination IP addresses and/or IP ranges.
- `destination_ports` - (Required) A list of destination ports.
- `protocols` - (Required) A list of protocols. Possible values are `Any`, `ICMP`, `TCP` and `UDP`. If `action` is `Dnat`, `protocols` can only be `TCP` and `UDP`.
- `source_addresses` - (Required) A list of source IP addresses and/or IP ranges.
- `translated_address` - (Required) The address of the service behind the Firewall.

- `translated_port` - (Required) The port of the service behind the Firewall.

Import

Azure Firewall NAT Rule Collections can be imported using the `resource id`, e.g.

```
terraform import azurerm_firewall_nat_rule_collection.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/azureFirewalls/myfirewall/natRuleCollections/mycollection
```

azurerm_firewall_network_rule_collection

Manages a Network Rule Collection within an Azure Firewall.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "North Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "testvnet"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "AzureFirewallSubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "test" {
  name             = "testpip"
  location         = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
  sku              = "Standard"
}

resource "azurerm_firewall" "test" {
  name            = "testfirewall"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name            = "configuration"
    subnet_id       = "${azurerm_subnet.test.id}"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_firewall_network_rule_collection" "test" {
  name            = "testcollection"
  azure_firewall_name = "${azurerm_firewall.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  priority        = 100
  action          = "Allow"

  rule {
    name = "testrule"
  }
}
```



```
source_addresses = [
    "10.0.0.0/16",
]

destination_ports = [
    "53",
]

destination_addresses = [
    "8.8.8.8",
    "8.8.4.4",
]

protocols = [
    "TCP",
    "UDP",
]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Network Rule Collection which must be unique within the Firewall. Changing this forces a new resource to be created.
- `azure_firewall_name` - (Required) Specifies the name of the Firewall in which the Network Rule Collection should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which the Firewall exists. Changing this forces a new resource to be created.
- `priority` - (Required) Specifies the priority of the rule collection. Possible values are between `100` - `65000` .
- `action` - (Required) Specifies the action the rule will apply to matching traffic. Possible values are `Allow` and `Deny` .
- `rule` - (Required) One or more `rule` blocks as defined below.

A `rule` block supports the following:

- `name` - (Required) Specifies the name of the rule.
- `description` - (Optional) Specifies a description for the rule.
- `source_addresses` - (Required) A list of source IP addresses and/or IP ranges.
- `destination_addresses` - (Required) A list of destination IP addresses and/or IP ranges.
- `destination_ports` - (Required) A list of destination ports.
- `protocols` - (Required) A list of protocols. Possible values are `Any` , `ICMP` , `TCP` and `UDP` .

Import

Azure Firewall Network Rule Collections can be imported using the `resource_id`, e.g.

```
terraform import azurerm_firewall_network_rule_collection.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/azureFirewalls/myfirewall/networkRuleCollections/mycollection
```

azurerm_frontdoor_firewall_policy

Manages an Azure Front Door Web Application Firewall Policy instance.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-rg"
  location  = "West US 2"
}

resource "azurerm_frontdoor_firewall_policy" "example" {
  name                        = "example-fdwafpolicy"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  enabled                    = true
  mode                       = "Prevention"
  redirect_url                = "https://www.contoso.com"
  custom_block_response_status_code = 403
  custom_block_response_body    = "PGh0bWw+CjxoZWFKZXI+PHRpdGxlpkhlbGxvPC90aXRsZT48L2hlYWRLcj4KPGJvZH
k+CkhlbGxvIHdvcmxkCjwvYm9keT4KPC9odG1sPg=="

  custom_rule {
    name                = "Rule1"
    enabled              = true
    priority             = 1
    rate_limit_duration_in_minutes = 1
    rate_limit_threshold = 10
    type                = "MatchRule"
    action              = "Block"

    match_condition {
      match_variable = "RemoteAddr"
      operator       = "IPMatch"
      negation_condition = false
      match_values   = ["192.168.1.0/24", "10.0.0.0/24"]
    }
  }
}

custom_rules {
  name                = "Rule2"
  enabled              = true
  priority             = 2
  rate_limit_duration_in_minutes = 1
  rate_limit_threshold = 10
  type                = "MatchRule"
  action              = "Block"

  match_condition {
    match_variable = "RemoteAddr"
    operator       = "IPMatch"
    negation_condition = false
    match_values   = ["192.168.1.0/24"]
  }
}

match_condition {
```

```

    match_variable      = "RequestHeader"
    selector            = "UserAgent"
    operator            = "Contains"
    negation_condition  = false
    match_values        = ["windows"]
    transforms          = ["Lowercase", "Trim"]
  }
}

managed_rule {
  type      = "DefaultRuleSet"
  version   = "preview-0.1"

  override {
    rule_group_name = "PHP"

    rule {
      rule_id = "933111"
      enabled = false
      action  = "Block"
    }
  }
}

managed_rule {
  type      = "BotProtection"
  version   = "preview-0.1"
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the policy. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group. Changing this forces a new resource to be created.
- `enabled` - (Optional) Is the policy a enabled state or disabled state. Defaults to `true`.
- `mode` - (Optional) The firewall policy mode. Possible values are `Detection`, `Prevention` and defaults to `Prevention`.
- `redirect_url` - (Optional) If action type is `redirect`, this field represents redirect URL for the client.
- `custom_rule` - (Optional) One or more `custom_rule` blocks as defined below.
- `custom_block_response_status_code` - (Optional) If a `custom_rule` block's action type is `block`, this is the response status code. Possible values are `200`, `403`, `405`, `406`, or `429`.
- `custom_block_response_body` - (Optional) If a `custom_rule` block's action type is `block`, this is the response body. The body must be specified in base64 encoding.
- `managed_rule` - (Optional) One or more `managed_rule` blocks as defined below.

- `tags` - (Optional) A mapping of tags to assign to the Web Application Firewall Policy.
-

The `custom_rule` block supports the following:

- `name` - (Required) Gets name of the resource that is unique within a policy. This name can be used to access the resource.
 - `action` - (Required) The action to perform when the rule is matched. Possible values are `Allow`, `Block`, `Log`, or `Redirect`.
 - `enabled` - (Optional) Is the rule is enabled or disabled? Defaults to `true`.
 - `priority` - (Required) The priority of the rule. Rules with a lower value will be evaluated before rules with a higher value. Defaults to `1`.
 - `type` - (Required) The the type of rule. Possible values are `MatchRule` or `RateLimitRule`.
 - `match_condition` - (Required) One or more `match_condition` block defined below.
 - `rate_limit_duration_in_minutes` - (Optional) The rate limit duration in minutes. Defaults to `1`.
 - `rate_limit_threshold` - (Optional) The rate limit threshold. Defaults to `10`.
-

The `match_condition` block supports the following:

- `match_variable` - (Required) The request variable to compare with. Possible values are `Cookies`, `PostArgs`, `QueryString`, `RemoteAddr`, `RequestBody`, `RequestHeader`, `RequestMethod`, or `RequestURI`.
 - `match_values` - (Required) Up to `100` possible values to match.
 - `operator` - (Required) Comparison type to use for matching with the variable value. Possible values are `Any`, `BeginsWith`, `Contains`, `EndsWith`, `Equal`, `GeoMatch`, `GreaterThan`, `GreaterThanOrEqual`, `IPMatch`, `LessThan`, `LessThanOrEqual` or `Regex`.
 - `selector` - (Optional) Match against a specific key if the `match_variable` is `QueryString`, `PostArgs`, `RequestHeader` or `Cookies`.
 - `negation_condition` - (Optional) Should the result of the condition be negated.
 - `transforms` - (Optional) Up to `5` transforms to apply. Possible values are `Lowercase`, `RemoveNulls`, `Trim`, `Uppercase`, `URLDecode` or `URLEncode`.
-

The `managed_rule` block supports the following:

- `type` - (Required) The name of the managed rule to use with this resource.
 - `version` - (Required) The version on the managed rule to use with this resource.
 - `override` - (Optional) One or more `override` blocks as defined below.
-

The `override` block supports the following:

- `rule_group_name` - (Required) The managed rule group to override.

- `rule` - (Optional) One or more `rule` blocks as defined below. If none are specified, all of the rules in the group will be disabled.

The `rule` block supports the following:

- `rule_id` - (Required) Identifier for the managed rule.
- `action` - (Required) The action to be applied when the rule matches. Possible values are `Allow`, `Block`, `Log`, or `Redirect`.
- `enabled` - (Optional) Is the managed rule override enabled or disabled. Defaults to `false`

Attributes Reference

The following attributes are exported:

- `id` - Resource ID.
- `location` - Resource location.
- `frontend_endpoint_ids` - the Frontend Endpoints associated with this Front Door Web Application Firewall policy.

Import

Front Door Web Application Firewall Policy can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_frontdoor_firewall_policy.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-rg/providers/Microsoft.Network/frontdoorwebapplicationfirewallpolicies/example-fdwafpolicy
```

azurerem_front_door

Manages an Azure Front Door instance.

Azure Front Door Service is Microsoft's highly available and scalable web application acceleration platform and global HTTP(s) load balancer. It provides built-in DDoS protection and application layer security and caching. Front Door enables you to build applications that maximize and automate high-availability and performance for your end-users. Use Front Door with Azure services including Web/Mobile Apps, Cloud Services and Virtual Machines – or combine it with on-premises services for hybrid deployments and smooth cloud migration.

Below are some of the key scenarios that Azure Front Door Service addresses: * Use Front Door to improve application scale and availability with instant multi-region failover * Use Front Door to improve application performance with SSL offload and routing requests to the fastest available application backend. * Use Front Door for application layer security and DDoS protection for your application.

Example Usage

```

resource "azurerm_front_door" "example" {
  name                                = "example-FrontDoor"
  location                            = "${azurerm_resource_group.example.location}"
  resource_group_name                 = "${azurerm_resource_group.example.name}"
  enforce_backend_pools_certificate_name_check = false

  routing_rule {
    name                = "exampleRoutingRule1"
    accepted_protocols = ["Http", "Https"]
    patterns_to_match   = ["/*"]
    frontend_endpoints = ["exampleFrontendEndpoint1"]
    forwarding_configuration {
      forwarding_protocol = "MatchRequest"
      backend_pool_name   = "exampleBackendBing"
    }
  }

  backend_pool_load_balancing {
    name = "exampleLoadBalancingSettings1"
  }

  backend_pool_health_probe {
    name = "exampleHealthProbeSetting1"
  }

  backend_pool {
    name = "exampleBackendBing"
    backend {
      host_header = "www.bing.com"
      address     = "www.bing.com"
      http_port   = 80
      https_port  = 443
    }

    load_balancing_name = "exampleLoadBalancingSettings1"
    health_probe_name   = "exampleHealthProbeSetting1"
  }

  frontend_endpoint {
    name                = "exampleFrontendEndpoint1"
    host_name           = "example-FrontDoor.azurefd.net"
    custom_https_provisioning_enabled = false
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Name of the Front Door which is globally unique. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) Name of the Resource group within the Azure subscription. Changing this forces a new resource to be created.

- `location` - (Required) Resource location. Changing this forces a new resource to be created.
 - `backend_pool` - (Required) A `backend_pool` block as defined below.
 - `backend_pool_health_probe` - (Required) A `backend_pool_health_probe` block as defined below.
 - `backend_pool_load_balancing` - (Required) A `backend_pool_load_balancing` block as defined below.
 - `enforce_backend_pools_certificate_name_check` - (Required) Whether to enforce certificate name check on HTTPS requests to all backend pools. No effect on non-HTTPS requests. Permitted values are `true` or `false`.
 - `load_balancer_enabled` - (Optional) Operational status of the Front Door load balancer. Permitted values are `true` or `false`. Defaults to `true`.
 - `friendly_name` - (Optional) A friendly name for the Front Door service.
 - `frontend_endpoint` - (Required) A `frontend_endpoint` block as defined below.
 - `routing_rule` - (Required) A `routing_rule` block as defined below.
 - `tags` - (Optional) Resource tags.
-

The `backend_pool` block supports the following:

- `name` - (Required) The name of the Backend Pool.
 - `backend` - (Required) A `backend` block as defined below.
 - `load_balancing_name` - (Required) The name property of the `backend_pool_load_balancing` block within this resource to use for the Backend Pool.
 - `health_probe_name` - (Required) The name property of a `backend_pool_health_probe` block within this resource to use for the Backend Pool.
-

The `backend` block supports the following:

- `address` - (Required) Location of the backend (IP address or FQDN)
 - `host_header` - (Required) The value to use as the host header sent to the backend.
 - `http_port` - (Required) The HTTP TCP port number. Possible values are between 1 - 65535.
 - `https_port` - (Required) The HTTPS TCP port number. Possible values are between 1 - 65535.
 - `priority` - (Optional) Priority to use for load balancing. Higher priorities will not be used for load balancing if any lower priority backend is healthy. Defaults to 1.
 - `weight` - (Optional) Weight of this endpoint for load balancing purposes. Defaults to 50.
-

The `frontend_endpoint` block supports the following:

- `name` - (Required) The name of the Frontend Endpoint.
- `host_name` - (Required) The host name of the Frontend Endpoint. Must be a domain name.

- `custom_https_provisioning_enabled` - (Required) Whether to allow HTTPS protocol for a custom domain that's associated with Front Door to ensure sensitive data is delivered securely via TLS/SSL encryption when sent across the internet. Valid options are `true` or `false`.
 - `session_affinity_enabled` - (Optional) Whether to allow session affinity on this host. Valid options are `true` or `false`. Defaults to `false`.
 - `session_affinity_ttl_seconds` - (Optional) The TTL to use in seconds for session affinity, if applicable. Defaults to `0`.
 - `web_application_firewall_policy_link_id` - (Optional) Defines the Web Application Firewall policy ID for each host.
-

The `backend_pool_health_probe` block supports the following:

- `name` - (Required) The name of the Azure Front Door Backend Health Probe.
 - `path` - (Optional) The path to use for the Backend Health Probe. Default is `/`.
 - `protocol` - (Optional) Protocol scheme to use for the Backend Health Probe. Defaults to `Http`.
 - `interval_in_seconds` - (Optional) The number of seconds between health probes. Defaults to `120`.
-

The `backend_pool_load_balancing` block supports the following:

- `name` - (Required) The name of the Azure Front Door Backend Load Balancer.
 - `sample_size` - (Optional) The number of samples to consider for load balancing decisions. Defaults to `4`.
 - `successful_samples_required` - (Optional) The number of samples within the sample period that must succeed. Defaults to `2`.
 - `additional_latency_milliseconds` - (Optional) The additional latency in milliseconds for probes to fall into the lowest latency bucket. Defaults to `0`.
-

The `routing_rule` block supports the following:

- `name` - (Required) The name of the Front Door Backend Routing Rule.
 - `frontend_endpoints` - (Required) The names of the `frontend_endpoint` blocks within this resource to associate with this `routing_rule`.
 - `accepted_protocols` - (Optional) Protocol schemes to match for the Backend Routing Rule. Defaults to `Http`.
 - `patterns_to_match` - (Optional) The route patterns for the Backend Routing Rule. Defaults to `/*`.
 - `enabled` - (Optional) Enable or Disable use of this Backend Routing Rule. Permitted values are `true` or `false`. Defaults to `true`.
 - `forwarding_configuration` - (Optional) A `forwarding_configuration` block as defined below.
 - `redirect_configuration` - (Optional) A `redirect_configuration` block as defined below.
-

The `forwarding_configuration` block supports the following:

- `backend_pool_name` - (Required) The name of the Front Door Backend Pool.
 - `cache_use_dynamic_compression` - (Optional) Whether to use dynamic compression when caching. Valid options are `true` or `false`. Defaults to `true`.
 - `cache_query_parameter_strip_directive` - (Optional) Defines cache behavior in relation to query string parameters. Valid options are `StripAll` or `StripNone`. Defaults to `StripNone`.
 - `custom_forwarding_path` - (Optional) Path to use when constructing the request to forward to the backend. This functions as a URL Rewrite. Default behavior preserves the URL path.
 - `forwarding_protocol` - (Optional) Protocol to use when redirecting. Valid options are `HTTPOnly`, `HTTPSOnly`, or `MatchRequest`. Defaults to `MatchRequest`.
-

The `redirect_configuration` block supports the following:

- `custom_host` - (Optional) Set this to change the URL for the redirection.
 - `redirect_protocol` - (Optional) Protocol to use when redirecting. Valid options are `HTTPOnly`, `HTTPSOnly`, `MatchRequest`. Defaults to `MatchRequest`.
 - `redirect_type` - (Optional) Status code for the redirect. Valid options are `Moved`, `Found`, `TemporaryRedirect`, `PermanentRedirect`. Defaults to `Found`.
 - `custom_fragment` - (Optional) The destination fragment in the portion of URL after '#'. Set this to add a fragment to the redirect URL.
 - `custom_path` - (Optional) The path to retain as per the incoming request, or update in the URL for the redirection.
 - `custom_query_string` - (Optional) Replace any existing query string from the incoming request URL.
-

The `custom_https_configuration` block supports the following:

- `certificate_source` - (Optional) Certificate source to encrypted HTTPS traffic with. Allowed values are `FrontDoor` or `AzureKeyVault`. Defaults to `FrontDoor`.

The following attributes are only valid if `certificate_source` is set to `AzureKeyVault`:

- `azure_key_vault_certificate_vault_id` - (Required) The `id` of the Key Vault containing the SSL certificate.
- `azure_key_vault_certificate_secret_name` - (Required) The name of the Key Vault secret representing the full certificate PFX.
- `azure_key_vault_certificate_secret_version` - (Required) The version of the Key Vault secret representing the full certificate PFX.

Note: In order to enable the use of your own custom HTTPS certificate you must grant Azure Front Door Service access to your key vault. For instructions on how to configure your Key Vault correctly please refer to the product documentation (<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-custom-domain-https#option-2-use-your-own-certificate>).

Attributes Reference

`backend_pool` exports the following:

- `id` - The Resource ID of the Azure Front Door Backend Pool.

`backend` exports the following:

- `id` - The Resource ID of the Azure Front Door Backend.

`frontend_endpoint` exports the following:

- `id` - The Resource ID of the Azure Front Door Frontend Endpoint.
- `provisioning_state` - Provisioning state of the Front Door.
- `provisioning_substate` - Provisioning substate of the Front Door

`backend_pool_health_probe` exports the following:

- `id` - The Resource ID of the Azure Front Door Backend Health Probe.

`backend_pool_load_balancing` exports the following:

- `id` - The Resource ID of the Azure Front Door Backend Load Balancer.

`routing_rule` exports the following:

- `id` - The Resource ID of the Azure Front Door Backend Routing Rule.

The following attributes are exported:

- `cname` - The host that each frontendEndpoint must CNAME to.
- `id` - Resource ID.

Import

Front Doors can be imported using the `resource id`, e.g.

```
terraform import azurerm_frontdoor.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/mygroup1/providers/Microsoft.Network/frontdoors/frontdoor1
```

azurerm_function_app

Manages a Function App.

Example Usage (with App Service Plan)

```
resource "azurerm_resource_group" "test" {
  name      = "azure-functions-test-rg"
  location  = "westus2"
}

resource "azurerm_storage_account" "test" {
  name                        = "functionsapptests"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_app_service_plan" "test" {
  name                = "azure-functions-test-service-plan"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    tier = "Standard"
    size = "S1"
  }
}

resource "azurerm_function_app" "test" {
  name                = "test-azure-functions"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"
  storage_connection_string = "${azurerm_storage_account.test.primary_connection_string}"
}
```

Example Usage (in a Consumption Plan)

```

resource "azurerm_resource_group" "test" {
  name      = "azure-functions-cptest-rg"
  location  = "westus2"
}

resource "azurerm_storage_account" "test" {
  name                        = "functionsapptestsa"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_app_service_plan" "test" {
  name                = "azure-functions-test-service-plan"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  kind               = "FunctionApp"

  sku {
    tier = "Dynamic"
    size = "Y1"
  }
}

resource "azurerm_function_app" "test" {
  name                = "test-azure-functions"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  app_service_plan_id = "${azurerm_app_service_plan.test.id}"
  storage_connection_string = "${azurerm_storage_account.test.primary_connection_string}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Function App. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Function App.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `app_service_plan_id` - (Required) The ID of the App Service Plan within which to create this Function App.
- `storage_connection_string` - (Required) The connection string of the backend storage account which will be used by this Function App (such as the dashboard, logs).
- `app_settings` - (Optional) A key-value pair of App Settings.

Note: When integrating a CI/CD pipeline and expecting to run from a deployed package in Azure you must seed your `app_settings` as part of terraform code for function app to be successfully deployed. Important Default key pairs: (`"WEBSITE_RUN_FROM_PACKAGE"` = `""`, `"FUNCTIONS_WORKER_RUNTIME"` = `"node"` (or `python`, etc),

```
"WEBSITE_NODE_DEFAULT_VERSION" = "10.14.1", "APPINSIGHTS_INSTRUMENTATIONKEY" = "" ).
```

Note: When using an App Service Plan in the Free or Shared Tiers `use_32_bit_worker_process` must be set to `true`.

- `auth_settings` - (Optional) A `auth_settings` block as defined below.
- `enable_built_in_logging` - (Optional) Should the built-in logging of this Function App be enabled? Defaults to `true`.
- `connection_string` - (Optional) An `connection_string` block as defined below.
- `client_affinity_enabled` - (Optional) Should the Function App send session affinity cookies, which route client requests in the same session to the same instance?
- `enabled` - (Optional) Is the Function App enabled?
- `https_only` - (Optional) Can the Function App only be accessed via HTTPS? Defaults to `false`.
- `version` - (Optional) The runtime version associated with the Function App. Defaults to `~1`.
- `site_config` - (Optional) A `site_config` object as defined below.
- `identity` - (Optional) An `identity` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`connection_string` supports the following:

- `name` - (Required) The name of the Connection String.
- `type` - (Required) The type of the Connection String. Possible values are `APIHub`, `Custom`, `DocDb`, `EventHub`, `MySQL`, `NotificationHub`, `PostgreSQL`, `RedisCache`, `ServiceBus`, `SQLAzure` and `SQLServer`.
- `value` - (Required) The value for the Connection String.

`site_config` supports the following:

- `always_on` - (Optional) Should the Function App be loaded at all times? Defaults to `false`.
- `use_32_bit_worker_process` - (Optional) Should the Function App run in 32 bit mode, rather than 64 bit mode? Defaults to `true`.

Note: when using an App Service Plan in the Free or Shared Tiers `use_32_bit_worker_process` must be set to `true`.

- `websockets_enabled` - (Optional) Should WebSockets be enabled?
- `virtual_network_name` - (Optional) The name of the Virtual Network which this App Service should be attached to.
- `linux_fx_version` - (Optional) Linux App Framework and version for the AppService, e.g. `DOCKER| (golang:latest)`.
- `http2_enabled` - (Optional) Specifies whether or not the http2 protocol should be enabled. Defaults to `false`.

- `cors` - (Optional) A `cors` block as defined below.

A `cors` block supports the following:

- `allowed_origins` - (Optional) A list of origins which should be able to make cross-origin calls. `*` can be used to allow all calls.
- `support_credentials` - (Optional) Are credentials supported?

`identity` supports the following:

- `type` - (Required) Specifies the identity type of the App Service. At this time the only allowed value is `SystemAssigned`.

An `auth_settings` block supports the following:

- `enabled` - (Required) Is Authentication enabled?
- `active_directory` - (Optional) A `active_directory` block as defined below.
- `additional_login_params` - (Optional) Login parameters to send to the OpenID Connect authorization endpoint when a user logs in. Each parameter must be in the form "key=value".
- `allowed_external_redirect_urls` - (Optional) External URLs that can be redirected to as part of logging in or logging out of the app.
- `default_provider` - (Optional) The default provider to use when multiple providers have been set up. Possible values are `AzureActiveDirectory`, `Facebook`, `Google`, `MicrosoftAccount` and `Twitter`.

NOTE: When using multiple providers, the default provider must be set for settings like `unauthenticated_client_action` to work.

- `facebook` - (Optional) A `facebook` block as defined below.
- `google` - (Optional) A `google` block as defined below.
- `issuer` - (Optional) Issuer URI. When using Azure Active Directory, this value is the URI of the directory tenant, e.g. `https://sts.windows.net/{tenant-guid}/` (`https://sts.windows.net/%7Btenant-guid%7D/`).
- `microsoft` - (Optional) A `microsoft` block as defined below.
- `runtime_version` - (Optional) The runtime version of the Authentication/Authorization module.
- `token_refresh_extension_hours` - (Optional) The number of hours after session token expiration that a session token can be used to call the token refresh API. Defaults to 72.
- `token_store_enabled` - (Optional) If enabled the module will durably store platform-specific security tokens that are obtained during login flows. Defaults to false.
- `twitter` - (Optional) A `twitter` block as defined below.
- `unauthenticated_client_action` - (Optional) The action to take when an unauthenticated client attempts to access the app. Possible values are `AllowAnonymous` and `RedirectToLoginPage`.

An `active_directory` block supports the following:

- `client_id` - (Required) The Client ID of this relying party application. Enables OpenIDConnection authentication with Azure Active Directory.
- `client_secret` - (Optional) The Client Secret of this relying party application. If no secret is provided, implicit flow will be used.
- `allowed_audiences` (Optional) Allowed audience values to consider when validating JWTs issued by Azure Active Directory.

A `facebook` block supports the following:

- `app_id` - (Required) The App ID of the Facebook app used for login
- `app_secret` - (Required) The App Secret of the Facebook app used for Facebook Login.
- `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Facebook Login authentication. <https://developers.facebook.com/docs/facebook-login> (<https://developers.facebook.com/docs/facebook-login>)

A `google` block supports the following:

- `client_id` - (Required) The OpenID Connect Client ID for the Google web application.
- `client_secret` - (Required) The client secret associated with the Google web application.
- `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Google Sign-In authentication. <https://developers.google.com/identity/sign-in/web/> (<https://developers.google.com/identity/sign-in/web/>)

A `microsoft` block supports the following:

- `client_id` - (Required) The OAuth 2.0 client ID that was created for the app used for authentication.
- `client_secret` - (Required) The OAuth 2.0 client secret that was created for the app used for authentication.
- `oauth_scopes` (Optional) The OAuth 2.0 scopes that will be requested as part of Microsoft Account authentication. <https://msdn.microsoft.com/en-us/library/dn631845.aspx> (<https://msdn.microsoft.com/en-us/library/dn631845.aspx>)

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Function App
- `default_hostname` - The default hostname associated with the Function App - such as `mysite.azurewebsites.net`
- `outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12`
- `possible_outbound_ip_addresses` - A comma separated list of outbound IP addresses - such as `52.23.25.3,52.143.43.12,52.143.43.17` - not all of which are necessarily in use. Superset of `outbound_ip_addresses`.

- `identity` - An `identity` block as defined below, which contains the Managed Service Identity information for this App Service.
 - `site_credential` - A `site_credential` block as defined below, which contains the site-level credentials used to publish to this App Service.
 - `kind` - The Function App kind - such as `functionapp`, `linux`, `container`
-

`identity` exports the following:

- `principal_id` - The Principal ID for the Service Principal associated with the Managed Service Identity of this App Service.
- `tenant_id` - The Tenant ID for the Service Principal associated with the Managed Service Identity of this App Service.

`site_credential` exports the following:

- `username` - The username which can be used to publish to this App Service
- `password` - The password associated with the username, which can be used to publish to this App Service.

Import

Function Apps can be imported using the `resource_id`, e.g.

```
terraform import azurerm_function_app.functionapp1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Web/sites/functionapp1
```

azurerm_hdinsight_hadoop_cluster

Manages a HDInsight Hadoop Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_hadoop_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"

  component_version {
    hadoop = "2.7"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key   = "${azurerm_storage_account.example.primary_access_key}"
    is_default             = true
  }

  roles {
    head_node {
      vm_size = "Standard_D3_V2"
      username = "acctestusrvn"
      password = "AccTestvdSC4daf986!"
    }
  }
}
```

```

worker_node {
  vm_size          = "Standard_D4_V2"
  username          = "acctestusrvm"
  password          = "AccTestvdSC4daf986!"
  target_instance_count = 3
}

zookeeper_node {
  vm_size = "Standard_D3_V2"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name for this HDInsight Hadoop Cluster. Changing this forces a new resource to be created.
 - **resource_group_name** - (Required) Specifies the name of the Resource Group in which this HDInsight Hadoop Cluster should exist. Changing this forces a new resource to be created.
 - **location** - (Required) Specifies the Azure Region which this HDInsight Hadoop Cluster should exist. Changing this forces a new resource to be created.
 - **cluster_version** - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
 - **component_version** - (Required) A `component_version` block as defined below.
 - **gateway** - (Required) A `gateway` block as defined below.
 - **roles** - (Required) A `roles` block as defined below.
 - **storage_account** - (Required) One or more `storage_account` block as defined below.
 - **tier** - (Required) Specifies the Tier which should be used for this HDInsight Hadoop Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.
-
- **tags** - (Optional) A map of Tags which should be assigned to this HDInsight Hadoop Cluster.
-

A `component_version` block supports the following:

- **hadoop** - (Required) The version of Hadoop which should be used for this HDInsight Hadoop Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight Hadoop Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight Hadoop Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight Hadoop Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight Hadoop Cluster.

Import

HDInsight Hadoop Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_hadoop_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000
0/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_hbase_cluster

Manages a HDInsight HBase Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_hbase_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"

  component_version {
    hbase = "1.1"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_D3_V2"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }
  }
}
```



```

worker_node {
  vm_size          = "Standard_D3_V2"
  username         = "acctestusrvm"
  password         = "AccTestvdSC4daf986!"
  target_instance_count = 3
}

zookeeper_node {
  vm_size = "Standard_D3_V2"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight HBase Cluster. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight HBase Cluster should exist. Changing this forces a new resource to be created.
 - `location` - (Required) Specifies the Azure Region which this HDInsight HBase Cluster should exist. Changing this forces a new resource to be created.
 - `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
 - `component_version` - (Required) A `component_version` block as defined below.
 - `gateway` - (Required) A `gateway` block as defined below.
 - `roles` - (Required) A `roles` block as defined below.
 - `storage_account` - (Required) One or more `storage_account` block as defined below.
 - `tier` - (Required) Specifies the Tier which should be used for this HDInsight HBase Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.
-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight HBase Cluster.
-

A `component_version` block supports the following:

- `hbase` - (Required) The version of HBase which should be used for this HDInsight HBase Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight HBase Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight HBase Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight HBase Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight HBase Cluster.

Import

HDInsight HBase Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_hbase_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_interactive_query_cluster

Manages a HDInsight Interactive Query Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_interactive_query_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"

  component_version {
    interactive_hive = "2.1"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_D13_V2"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }
  }
}
```

```

worker_node {
  vm_size          = "Standard_D14_V2"
  username         = "acctestusrvm"
  password         = "AccTestvdSC4daf986!"
  target_instance_count = 3
}

zookeeper_node {
  vm_size = "Standard_A4_V2"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight Interactive Query Cluster. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight Interactive Query Cluster should exist. Changing this forces a new resource to be created.
 - `location` - (Required) Specifies the Azure Region which this HDInsight Interactive Query Cluster should exist. Changing this forces a new resource to be created.
 - `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
 - `component_version` - (Required) A `component_version` block as defined below.
 - `gateway` - (Required) A `gateway` block as defined below.
 - `roles` - (Required) A `roles` block as defined below.
 - `storage_account` - (Required) One or more `storage_account` block as defined below.
 - `tier` - (Required) Specifies the Tier which should be used for this HDInsight Interactive Query Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.
-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight Interactive Query Cluster.
-

A `component_version` block supports the following:

- `interactive_query` - (Required) The version of Interactive Query which should be used for this HDInsight Interactive Query Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.

NOTE: High memory instances must be specified for the Head Node (Azure suggests a `Standard_D13_V2`).

- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight Interactive Query Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.

NOTE: High memory instances must be specified for the Head Node (Azure suggests a `Standard_D14_V2`).

- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight Interactive Query Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight Interactive Query Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight Interactive Query Cluster.

Import

HDInsight Interactive Query Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_interactive_query_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_kafka_cluster

Manages a HDInsight Kafka Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                     = "hdinsightstor"
  resource_group_name     = "${azurerm_resource_group.example.name}"
  location                 = "${azurerm_resource_group.example.location}"
  account_tier             = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                     = "hdinsight"
  resource_group_name     = "${azurerm_resource_group.example.name}"
  storage_account_name    = "${azurerm_storage_account.example.name}"
  container_access_type   = "private"
}

resource "azurerm_hdinsight_kafka_cluster" "example" {
  name                     = "example-hdicluster"
  resource_group_name     = "${azurerm_resource_group.example.name}"
  location                 = "${azurerm_resource_group.example.location}"
  cluster_version         = "3.6"
  tier                     = "Standard"

  component_version {
    kafka = "2.3"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size   = "Standard_D3_V2"
      username  = "acctestusrvm"
      password  = "AccTestvdSC4daf986!"
    }
  }
}
```

```

worker_node {
  vm_size           = "Standard_D3_V2"
  username          = "acctestusrvm"
  password          = "AccTestvdSC4daf986!"
  number_of_disks_per_node = 3
  target_instance_count  = 3
}

zookeeper_node {
  vm_size = "Standard_D3_V2"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight Kafka Cluster. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight Kafka Cluster should exist. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the Azure Region which this HDInsight Kafka Cluster should exist. Changing this forces a new resource to be created.
- `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
- `component_version` - (Required) A `component_version` block as defined below.
- `gateway` - (Required) A `gateway` block as defined below.
- `roles` - (Required) A `roles` block as defined below.
- `storage_account` - (Required) One or more `storage_account` block as defined below.
- `tier` - (Required) Specifies the Tier which should be used for this HDInsight Kafka Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.

-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight Kafka Cluster.
-

A `component_version` block supports the following:

- `kafka` - (Required) The version of Kafka which should be used for this HDInsight Kafka Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight Kafka Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `number_of_disks_per_node` - (Required) The number of Data Disks which should be assigned to each Worker Node, which can be between 1 and 8. Changing this forces a new resource to be created.
- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight Kafka Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight Kafka Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight Kafka Cluster.

Import

HDInsight Kafka Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_kafka_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_ml_services_cluster

Manages a HDInsight ML Services Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_ml_services_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"
  rstudio                   = true

  gateway {
    enabled = true
    username = "acctestusrgw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_D3_v2"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }

    worker_node {
      vm_size = "Standard_D4_V2"
    }
  }
}
```

```

    -
    username          = "acctestusrvm"
    password          = "AccTestvdSC4daf986!"
    target_instance_count = 3
  }

  zookeeper_node {
    vm_size = "Standard_D3_v2"
    username = "acctestusrvm"
    password = "AccTestvdSC4daf986!"
  }

  edge_node {
    vm_size = "Standard_D3_v2"
    username = "acctestusrvm"
    password = "AccTestvdSC4daf986!"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight ML Services Cluster. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight ML Services Cluster should exist. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the Azure Region which this HDInsight ML Services Cluster should exist. Changing this forces a new resource to be created.
- `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
- `gateway` - (Required) A `gateway` block as defined below.
- `roles` - (Required) A `roles` block as defined below.
- `rstudio` - (Required) Should R Studio community edition for ML Services be installed? Changing this forces a new resource to be created.
- `storage_account` - (Required) One or more `storage_account` block as defined below.
- `tier` - (Required) Specifies the Tier which should be used for this HDInsight ML Services Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.

-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight ML Services Cluster.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.

- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `edge_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Edge Node. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Edge Node. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Edge Node. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Edge Node. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Edge Node should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Edge Node should be provisioned within. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `edge_node` - (Required) A `edge_node` block as defined above.
- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight ML Services Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `' " ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight ML Services Cluster.
- `edge_ssh_endpoint` - The SSH Connectivity Endpoint for the Edge Node of the HDInsight ML Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight ML Services Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight ML Services Cluster.

Import

HDInsight ML Services Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_ml_services_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_rserver_cluster

Manages a HDInsight RServer Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_rserver_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"
  rstudio                   = true

  gateway {
    enabled = true
    username = "acctestusrgw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_D3_v2"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }

    worker_node {
      vm_size = "Standard_D4_V2"
    }
  }
}
```

```

    username      = "acctestusrvm"
    password      = "AccTestvdSC4daf986!"
    target_instance_count = 3
  }

  zookeeper_node {
    vm_size = "Standard_D3_v2"
    username = "acctestusrvm"
    password = "AccTestvdSC4daf986!"
  }

  edge_node {
    vm_size = "Standard_D3_v2"
    username = "acctestusrvm"
    password = "AccTestvdSC4daf986!"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight RServer Cluster. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight RServer Cluster should exist. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the Azure Region which this HDInsight RServer Cluster should exist. Changing this forces a new resource to be created.
- `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
- `gateway` - (Required) A `gateway` block as defined below.
- `roles` - (Required) A `roles` block as defined below.
- `rstudio` - (Required) Should R Studio community edition for RServer be installed? Changing this forces a new resource to be created.
- `storage_account` - (Required) One or more `storage_account` block as defined below.
- `tier` - (Required) Specifies the Tier which should be used for this HDInsight RServer Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.

-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight RServer Cluster.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.

- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `edge_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Edge Node. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Edge Node. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Edge Node. Changing this forces a new resource to be created.
- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Edge Node. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Edge Node should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Edge Node should be provisioned within. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
 - `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
-

A `roles` block supports the following:

- `edge_node` - (Required) A `edge_node` block as defined above.
 - `head_node` - (Required) A `head_node` block as defined above.
 - `worker_node` - (Required) A `worker_node` block as defined below.
 - `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.
-

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight RServer Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `' " ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight RServer Cluster.
- `edge_ssh_endpoint` - The SSH Connectivity Endpoint for the Edge Node of the HDInsight RServer Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight RServer Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight RServer Cluster.

Import

HDInsight RServer Clusters can be imported using the `resource_id`, e.g.

```
terraform import azurerm_hdinsight_rserver_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000  
00/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_spark_cluster

Manages a HDInsight Spark Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                = "hdinsight"
  resource_group_name = "${azurerm_resource_group.example.name}"
  storage_account_name = "${azurerm_storage_account.example.name}"
  container_access_type = "private"
}

resource "azurerm_hdinsight_spark_cluster" "example" {
  name                = "example-hdicluster"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"
  cluster_version     = "3.6"
  tier                 = "Standard"

  component_version {
    spark = "2.3"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_A3"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }
  }
}
```

```

worker_node {
  vm_size          = "Standard_A3"
  username         = "acctestusrvm"
  password         = "AccTestvdSC4daf986!"
  target_instance_count = 3
}

zookeeper_node {
  vm_size = "Medium"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight Spark Cluster. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight Spark Cluster should exist. Changing this forces a new resource to be created.
 - `location` - (Required) Specifies the Azure Region which this HDInsight Spark Cluster should exist. Changing this forces a new resource to be created.
 - `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
 - `component_version` - (Required) A `component_version` block as defined below.
 - `gateway` - (Required) A `gateway` block as defined below.
 - `roles` - (Required) A `roles` block as defined below.
 - `storage_account` - (Required) One or more `storage_account` block as defined below.
 - `tier` - (Required) Specifies the Tier which should be used for this HDInsight Spark Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.
-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight Spark Cluster.
-

A `component_version` block supports the following:

- `spark` - (Required) The version of Spark which should be used for this HDInsight Spark Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'"` ```).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight Spark Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight Spark Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight Spark Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight Spark Cluster.

Import

HDInsight Spark Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_spark_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_hdinsight_storm_cluster

Manages a HDInsight Storm Cluster.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "example" {
  name                        = "hdinsightstor"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "example" {
  name                      = "hdinsight"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
  container_access_type     = "private"
}

resource "azurerm_hdinsight_storm_cluster" "example" {
  name                      = "example-hdicluster"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  location                  = "${azurerm_resource_group.example.location}"
  cluster_version           = "3.6"
  tier                      = "Standard"

  component_version {
    storm = "1.1"
  }

  gateway {
    enabled = true
    username = "acctestusrqw"
    password = "TerrAform123!"
  }

  storage_account {
    storage_container_id = "${azurerm_storage_container.example.id}"
    storage_account_key  = "${azurerm_storage_account.example.primary_access_key}"
    is_default           = true
  }

  roles {
    head_node {
      vm_size = "Standard_A3"
      username = "acctestusrvm"
      password = "AccTestvdSC4daf986!"
    }
  }
}
```



```

worker_node {
  vm_size          = "Standard_D3_V2"
  username         = "acctestusrvm"
  password         = "AccTestvdSC4daf986!"
  target_instance_count = 3
}

zookeeper_node {
  vm_size = "Standard_A4_V2"
  username = "acctestusrvm"
  password = "AccTestvdSC4daf986!"
}
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name for this HDInsight Storm Cluster. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) Specifies the name of the Resource Group in which this HDInsight Storm Cluster should exist. Changing this forces a new resource to be created.
 - `location` - (Required) Specifies the Azure Region which this HDInsight Storm Cluster should exist. Changing this forces a new resource to be created.
 - `cluster_version` - (Required) Specifies the Version of HDInsights which should be used for this Cluster. Changing this forces a new resource to be created.
 - `component_version` - (Required) A `component_version` block as defined below.
 - `gateway` - (Required) A `gateway` block as defined below.
 - `roles` - (Required) A `roles` block as defined below.
 - `storage_account` - (Required) One or more `storage_account` block as defined below.
 - `tier` - (Required) Specifies the Tier which should be used for this HDInsight Storm Cluster. Possible values are `Standard` or `Premium`. Changing this forces a new resource to be created.
-
- `tags` - (Optional) A map of Tags which should be assigned to this HDInsight Storm Cluster.
-

A `component_version` block supports the following:

- `storm` - (Required) The version of Storm which should be used for this HDInsight Storm Cluster. Changing this forces a new resource to be created.
-

A `gateway` block supports the following:

- `enabled` - (Required) Is the Ambari portal enabled? Changing this forces a new resource to be created.
- `password` - (Required) The password used for the Ambari Portal. Changing this forces a new resource to be created.

NOTE: This password must be different from the one used for the `head_node`, `worker_node` and `zookeeper_node` roles.

- `username` - (Required) The username used for the Ambari Portal. Changing this forces a new resource to be created.

A `head_node` block supports the following:

NOTE: This is also known as the `nimbus` node.

- `username` - (Required) The Username of the local administrator for the Head Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Head Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Head Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Head Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Head Nodes should be provisioned within. Changing this forces a new resource to be created.

A `roles` block supports the following:

- `head_node` - (Required) A `head_node` block as defined above.
- `worker_node` - (Required) A `worker_node` block as defined below.
- `zookeeper_node` - (Required) A `zookeeper_node` block as defined below.

A `storage_account` block supports the following:

- `is_default` - (Required) Is this the Default Storage Account for the HDInsight Storm Cluster? Changing this forces a new resource to be created.

NOTE: One of the `storage_account` blocks must be marked as the default.

- `storage_account_key` - (Required) The Access Key which should be used to connect to the Storage Account. Changing this forces a new resource to be created.
- `storage_container_id` - (Required) The ID of the Storage Container. Changing this forces a new resource to be created.

NOTE: This can be obtained from the `id` of the `azurerm_storage_container` resource.

A `worker_node` block supports the following:

NOTE: This is also known as the `supervisor` node.

- `username` - (Required) The Username of the local administrator for the Worker Nodes. Changing this forces a new resource to be created.
- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Worker Nodes. Changing this forces a new resource to be created.
- `min_instance_count` - (Optional) The minimum number of instances which should be run for the Worker Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Worker Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters `'" ``).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Worker Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.
- `target_instance_count` - (Optional) The number of instances which should be run for the Worker Nodes.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Worker Nodes should be provisioned within. Changing this forces a new resource to be created.

A `zookeeper_node` block supports the following:

- `username` - (Required) The Username of the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

- `vm_size` - (Required) The Size of the Virtual Machine which should be used as the Zookeeper Nodes. Changing this forces a new resource to be created.
- `password` - (Optional) The Password associated with the local administrator for the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: If specified, this password must be at least 10 characters in length and must contain at least one digit, one uppercase and one lower case letter, one non-alphanumeric character (except characters ' " `).

- `ssh_keys` - (Optional) A list of SSH Keys which should be used for the local administrator on the Zookeeper Nodes. Changing this forces a new resource to be created.

NOTE: Either a `password` or one or more `ssh_keys` must be specified - but not both.

- `subnet_id` - (Optional) The ID of the Subnet within the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.
- `virtual_network_id` - (Optional) The ID of the Virtual Network where the Zookeeper Nodes should be provisioned within. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HDInsight Storm Cluster.
- `https_endpoint` - The HTTPS Connectivity Endpoint for this HDInsight Storm Cluster.
- `ssh_endpoint` - The SSH Connectivity Endpoint for this HDInsight Storm Cluster.

Import

HDInsight Storm Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_hdinsight_storm_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.HDInsight/clusters/cluster1}
```

azurerm_healthcare_service

Manages a Healthcare Service Resource.

Example Usage

```
resource "azurerm_healthcare_service" "example" {
  name                = "uniquefhirname"
  resource_group_name = "sample-resource-group"
  location            = "westus2"
  kind               = "fhir-R4"
  cosmosdb_throughput = "2000"

  access_policy_object_ids = ["xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx"]

  tags = {
    "environment" = "testenv"
    "purpose"     = "AcceptanceTests"
  }

  authentication_configuration {
    authority          = "https://login.microsoftonline.com/${data.azurerm_client_config.current.tenant_id}%7D"
    audience           = "https://azurehealthcareapis.com/"
    smart_proxy_enabled = "true"
  }

  cors_configuration {
    allowed_origins      = ["http://www.example.com", "http://www.example2.com"]
    allowed_headers      = ["x-tempo-*", "x-tempo2-*"]
    allowed_methods      = ["GET", "PUT"]
    max_age_in_seconds   = "500"
    allow_credentials     = "true"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the service instance. Used for service endpoint, must be unique within the audience.
- `resource_group_name` - (Required) The name of the Resource Group in which to create the Service.
- `location` - (Required) Specifies the supported Azure Region where the Service should be created.

Please Note: Not all locations support this resource. Some are `West US 2`, `North Central US`, and `UK West`.

- `access_policy_ids` - (Optional) A set of Azure object id's that are allowed to access the Service. If not configured, the default value is the object id of the service principal or user that is running Terraform.

- `authentication_configuration` - (Optional) An `authentication_configuration` block as defined below.
- `cosmosdb_throughput` - (Optional) The provisioned throughput for the backing database. Range of 400 - 1000 . Defaults to 400 .
- `cors_configuration` - (Optional) A `cors_configuration` block as defined below.
- `kind` - (Optional) The type of the service. Values at time of publication are: `fhir` , `fhir-Stu3` and `fhir-R4` . Default value is `fhir` .
- `tags` - (Optional) A mapping of tags to assign to the resource.

An `authentication_configuration` supports the following:

- `authority` - (Optional) The Azure Active Directory (tenant) that serves as the authentication authority to access the service. The default authority is the Directory defined in the authentication scheme in use when running Terraform. Authority must be registered to Azure AD and in the following format: `https://{Azure-AD-endpoint}/{tenant-id}`.
- `audience` - (Optional) The intended audience to receive authentication tokens for the service. The default value is `https://azurehealthcareapis.com` (`https://azurehealthcareapis.com`)
- `smart_proxy_enabled` - (Boolean) Enables the 'SMART on FHIR' option for mobile and web implementations.

A `cors_configuration` block supports the following:

- `allowed_origins` - (Required) A set of origins to be allowed via CORS.
- `allowed_headers` - (Required) A set of headers to be allowed via CORS.
- `allowed_methods` - (Required) The methods to be allowed via CORS.
- `max_age_in_seconds` - (Required) The max age to be allowed via CORS.
- `allow_credentials` - (Boolean) If credentials are allowed via CORS.

Attributes Reference

The following attributes are exported:

- `id` - The `id` of the Healthcare Service.

Import

Healthcare Service can be imported using the resource `id` , e.g.

```
terraform import azurerm_healthcare_service.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource_group/providers/Microsoft.HealthcareApis/services/service_name
```

azurerm_image

Manages a custom virtual machine image that can be used to create virtual machines.

Example Usage Creating from VHD

```
resource "azurerm_resource_group" "test" {
  name     = "acctest"
  location = "West US"
}

resource "azurerm_image" "test" {
  name                = "acctest"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  os_disk {
    os_type   = "Linux"
    os_state  = "Generalized"
    blob_uri  = "{blob_uri}"
    size_gb   = 30
  }
}
```

Example Usage Creating from Virtual Machine (VM must be generalized beforehand)

```
resource "azurerm_resource_group" "test" {
  name     = "acctest"
  location = "West US"
}

resource "azurerm_image" "test" {
  name                = "acctest"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  source_virtual_machine_id = "{vm_id}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the image. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the image. Changing this

forces a new resource to be created.

- `location` - (Required) Specified the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `source_virtual_machine_id` - (Optional) The Virtual Machine ID from which to create the image.
- `os_disk` - (Optional) One or more `os_disk` elements as defined below.
- `data_disk` - (Optional) One or more `data_disk` elements as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zone_resilient` - (Optional) Is zone resiliency enabled? Defaults to `false`. Changing this forces a new resource to be created.

Note: `zone_resilient` can only be set to `true` if the image is stored in a region that supports availability zones.

`os_disk` supports the following:

- `os_type` - (Required) Specifies the type of operating system contained in the the virtual machine image. Possible values are: Windows or Linux.
- `os_state` - (Required) Specifies the state of the operating system contained in the blob. Currently, the only value is Generalized.
- `managed_disk_id` - (Optional) Specifies the ID of the managed disk resource that you want to use to create the image.
- `blob_uri` - (Optional) Specifies the URI in Azure storage of the blob that you want to use to create the image.
- `caching` - (Optional) Specifies the caching mode as `ReadWrite`, `ReadOnly`, or `None`. The default is `None`.
- `size_gb` - (Optional) Specifies the size of the image to be created. The target size can't be smaller than the source size.

`data_disk` supports the following:

- `lun` - (Required) Specifies the logical unit number of the data disk.
- `managed_disk_id` - (Optional) Specifies the ID of the managed disk resource that you want to use to create the image.
- `blob_uri` - (Optional) Specifies the URI in Azure storage of the blob that you want to use to create the image.
- `caching` - (Optional) Specifies the caching mode as `ReadWrite`, `ReadOnly`, or `None`. The default is `None`.
- `size_gb` - (Optional) Specifies the size of the image to be created. The target size can't be smaller than the source size.

Attributes Reference

The following attributes are exported:

- `id` - The managed image ID.

Import

Image can be imported using the `resource id`, e.g.

```
terraform import azurerm_image.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.compute/images/image1
```

azurerm_iot_dps_certificate

Manages an IoT Device Provisioning Service Certificate.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "resourceGroup1"
  location = "West US"
}

resource "azurerm_iot_dps" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  sku {
    name     = "S1"
    tier      = "Standard"
    capacity = "1"
  }
}

resource "azurerm_iot_dps_certificate" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  iot_dps_name        = "${azurerm_iot_dps.example.name}"

  certificate_content = "${filebase64("example.cer")}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the IoT Device Provisioning Service Certificate resource. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group under which the IoT Device Provisioning Service Certificate resource has to be created. Changing this forces a new resource to be created.
- **iot_dps_name** - (Required) The name of the IoT Device Provisioning Service that this certificate will be attached to. Changing this forces a new resource to be created.
- **certificate_content** - (Required) The Base-64 representation of the X509 leaf certificate .cer file or just a .pem file content.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the IoT Device Provisioning Service Certificate.

Import

IoT Device Provisioning Service Certificate can be imported using the `resource id`, e.g.

```
terraform import azurerm_iot_dps_certificate.example /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.Devices/provisioningServices/example/certificates/example
```

azurerm_iot_dps

Manages an IoT Device Provisioning Service.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_iot_dps" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"

  sku {
    name      = "S1"
    tier       = "Standard"
    capacity  = "1"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the IoT Device Provisioning Service resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group under which the IoT Device Provisioning Service resource has to be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource has to be created. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `linked_hub` - (Optional) A `linked_hub` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The name of the sku. Possible values are `B1`, `B2`, `B3`, `F1`, `S1`, `S2`, and `S3`.
- `tier` - (Required) The billing tier for the IoT Device Provisioning Service. Possible values are `Basic`, `Free` or `Standard`.
- `capacity` - (Required) The number of provisioned IoT Device Provisioning Service units.

A `linked_hub` block supports the following:

- `connection_string` - (Required) The connection string to connect to the IoT Hub. Changing this forces a new resource.
- `location` - (Required) The location of the IoT hub. Changing this forces a new resource.
- `apply_application_policy` - (Optional) Determines whether to apply application policies to the IoT Hub. Defaults to `false`.
- `allocation_weight` - (Optional) The weight applied to the IoT Hub. Defaults to 0.
- `hostname` - (Computed) The IoT Hub hostname.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the IoT Device Provisioning Service.

Import

IoT Device Provisioning Service can be imported using the `resource id`, e.g.

```
terraform import azurerm_iot_dps.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Devices/provisioningServices/example
```

azurerm_iothub_consumer_group

Manages a Consumer Group within an IoT Hub

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_iothub" "test" {
  name                = "test"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  sku {
    name      = "S1"
    tier       = "Standard"
    capacity  = "1"
  }

  tags = {
    purpose = "testing"
  }
}

resource "azurerm_iothub_consumer_group" "test" {
  name                = "terraform"
  iothub_name         = "${azurerm_iothub.test.name}"
  eventhub_endpoint_name = "events"
  resource_group_name = "${azurerm_resource_group.foo.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of this Consumer Group. Changing this forces a new resource to be created.
- `iothub_name` - (Required) The name of the IoT Hub. Changing this forces a new resource to be created.
- `eventhub_endpoint_name` - (Required) The name of the Event Hub-compatible endpoint in the IoT hub. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group that contains the IoT hub. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the IoT Hub Consumer Group.

Import

IoT Hub Consumer Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_iothub_consumer_group.group1 /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.Devices/IotHubs/hub1/eventHubEndpoints/events/ConsumerGroups  
/group1
```

azurerm_iothub

Manages an IoT Hub

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_storage_account" "test" {
  name                        = "teststa"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "test" {
  name                      = "test"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  storage_account_name      = "${azurerm_storage_account.test.name}"
  container_access_type     = "private"
}

resource "azurerm_iothub" "test" {
  name                      = "test"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  location                  = "${azurerm_resource_group.test.location}"

  sku {
    name      = "S1"
    tier       = "Standard"
    capacity  = "1"
  }

  endpoint {
    type                        = "AzureIoTHub.StorageContainer"
    connection_string          = "${azurerm_storage_account.test.primary_blob_connection_string}"
    name                       = "export"
    batch_frequency_in_seconds = 60
    max_chunk_size_in_bytes    = 10485760
    container_name              = "test"
    encoding                    = "Avro"
    file_name_format            = "${iothub}/{partition}_{YYYY}_{MM}_{DD}_{HH}_{mm}"
  }

  route {
    name          = "export"
    source        = "DeviceMessages"
    condition     = "true"
    endpoint_names = ["export"]
    enabled       = true
  }
}
```



```

fallback_route {
  enabled = true
}

file_upload {
  connection_string = "${azurerm_storage_account.test.primary_blob_connection_string}"
  container_name    = "${azurerm_storage_container.test.name}"
  sas_ttl           = "PT1H"
  notifications     = true
  lock_duration     = "PT1M"
  default_ttl       = "PT1H"
  max_delivery_count = 10
}

tags = {
  purpose = "testing"
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the IoT Hub resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group under which the IoT Hub resource has to be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource has to be created. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `endpoint` - (Optional) An `endpoint` block as defined below.
- `ip_filter_rule` - (Optional) One or more `ip_filter_rule` blocks as defined below.
- `route` - (Optional) A `route` block as defined below.
- `fallback_route` - (Optional) A `fallback_route` block as defined below. If the fallback route is enabled, messages that don't match any of the supplied routes are automatically sent to this route. Defaults to `messages/events`.
- `file_upload` - (Optional) A `file_upload` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The name of the sku. Possible values are `B1`, `B2`, `B3`, `F1`, `S1`, `S2`, and `S3`.
- `tier` - (Required) The billing tier for the IoT Hub. Possible values are `Basic`, `Free` or `Standard`.

NOTE: Only one IoT Hub can be on the `Free` tier per subscription.

-
- `capacity` - (Required) The number of provisioned IoT Hub units.
-

An `endpoint` block supports the following:

- `type` - (Required) The type of the endpoint. Possible values are `AzureIotHub.StorageContainer` , `AzureIotHub.ServiceBusQueue` , `AzureIotHub.ServiceBusTopic` or `AzureIotHub.EventHub` .
 - `connection_string` - (Required) The connection string for the endpoint.
 - `name` - (Required) The name of the endpoint. The name must be unique across endpoint types. The following names are reserved: `events` , `operationsMonitoringEvents` , `fileNotifications` and `$default` .
 - `batch_frequency_in_seconds` - (Optional) Time interval at which blobs are written to storage. Value should be between 60 and 720 seconds. Default value is 300 seconds. This attribute is mandatory for endpoint type `AzureIotHub.StorageContainer` .
 - `max_chunk_size_in_bytes` - (Optional) Maximum number of bytes for each blob written to storage. Value should be between 10485760(10MB) and 524288000(500MB). Default value is 314572800(300MB). This attribute is mandatory for endpoint type `AzureIotHub.StorageContainer` .
 - `container_name` - (Optional) The name of storage container in the storage account. This attribute is mandatory for endpoint type `AzureIotHub.StorageContainer` .
 - `encoding` - (Optional) Encoding that is used to serialize messages to blobs. Supported values are 'avro' and 'avrodeflate'. Default value is 'avro'. This attribute is mandatory for endpoint type `AzureIotHub.StorageContainer` .
 - `file_name_format` - (Optional) File name format for the blob. Default format is `{iothub}/{partition}/{YYYY}/{MM}/{DD}/{HH}/{mm}` . All parameters are mandatory but can be reordered. This attribute is mandatory for endpoint type `AzureIotHub.StorageContainer` .
-

An `ip_filter_rule` block supports the following:

- `name` - (Required) The name of the filter.
 - `ip_mask` - (Required) The IP address range in CIDR notation for the rule.
 - `action` - (Required) The desired action for requests captured by this rule. Possible values are `Accept` , `Reject`
-

A `route` block supports the following:

- `name` - (Required) The name of the route.
- `source` - (Required) The source that the routing rule is to be applied to, such as `DeviceMessages` . Possible values include: `RoutingSourceInvalid` , `RoutingSourceDeviceMessages` , `RoutingSourceTwinChangeEvents` , `RoutingSourceDeviceLifecycleEvents` , `RoutingSourceDeviceJobLifecycleEvents` .
- `condition` - (Optional) The condition that is evaluated to apply the routing rule. If no condition is provided, it evaluates to true by default. For grammar, see: <https://docs.microsoft.com/azure/iot-hub/iot-hub-devguide-query-language> (<https://docs.microsoft.com/azure/iot-hub/iot-hub-devguide-query-language>).
- `endpoint_names` - (Required) The list of endpoints to which messages that satisfy the condition are routed.
- `enabled` - (Required) Used to specify whether a route is enabled.

A `fallback_route` block supports the following:

- `source` - (Optional) The source that the routing rule is to be applied to, such as `DeviceMessages`. Possible values include: `RoutingSourceInvalid`, `RoutingSourceDeviceMessages`, `RoutingSourceTwinChangeEvents`, `RoutingSourceDeviceLifecycleEvents`, `RoutingSourceDeviceJobLifecycleEvents`.
- `condition` - (Optional) The condition that is evaluated to apply the routing rule. If no condition is provided, it evaluates to true by default. For grammar, see: <https://docs.microsoft.com/azure/iot-hub/iot-hub-devguide-query-language> (<https://docs.microsoft.com/azure/iot-hub/iot-hub-devguide-query-language>).
- `endpoint_names` - (Optional) The endpoints to which messages that satisfy the condition are routed. Currently only 1 endpoint is allowed.
- `enabled` - (Optional) Used to specify whether the fallback route is enabled.

A `file_upload` block supports the following:

- `connection_string` - (Required) The connection string for the Azure Storage account to which files are uploaded.
- `container_name` - (Required) The name of the root container where you upload files. The container need not exist but should be creatable using the `connection_string` specified.
- `sas_ttl` - (Optional) The period of time for which the SAS URI generated by IoT Hub for file upload is valid, specified as an ISO 8601 timespan duration (https://en.wikipedia.org/wiki/ISO_8601#Durations). This value must be between 1 minute and 24 hours, and evaluates to 'PT1H' by default.
- `notifications` - (Optional) Used to specify whether file notifications are sent to IoT Hub on upload. It evaluates to false by default.
- `lock_duration` - (Optional) The lock duration for the file upload notifications queue, specified as an ISO 8601 timespan duration (https://en.wikipedia.org/wiki/ISO_8601#Durations). This value must be between 5 and 300 seconds, and evaluates to 'PT1M' by default.
- `default_ttl` - (Optional) The period of time for which a file upload notification message is available to consume before it is expired by the IoT hub, specified as an ISO 8601 timespan duration (https://en.wikipedia.org/wiki/ISO_8601#Durations). This value must be between 1 minute and 48 hours, and evaluates to 'PT1H' by default.
- `max_delivery_count` - (Optional) The number of times the IoT hub attempts to deliver a file upload notification message. It evaluates to 10 by default.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the IoT Hub.
- `event_hub_events_endpoint` - The EventHub compatible endpoint for events data
- `event_hub_events_path` - The EventHub compatible path for events data
- `event_hub_operations_endpoint` - The EventHub compatible endpoint for operational data

- `event_hub_operations_path` - The EventHub compatible path for operational data

NOTE: These fields can be used in conjunction with the `shared_access_policy` block to build a connection string

- `hostname` - The hostname of the IoT Hub Resource.
- `shared_access_policy` - One or more `shared_access_policy` blocks as defined below.

A `shared_access_policy` block contains the following:

- `key_name` - The name of the shared access policy.
- `primary_key` - The primary key.
- `secondary_key` - The secondary key.
- `permissions` - The permissions assigned to the shared access policy.

Import

IoT Hubs can be imported using the `resource_id`, e.g.

```
terraform import azurerm_iothub.hub1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Devices/IotHubs/hub1
```

azurerm_iothub_shared_access_policy

Manages an IoT Hub Shared Access Policy

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_iothub" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location            = "${azurerm_resource_group.example.location}"

  sku {
    name      = "S1"
    tier       = "Standard"
    capacity  = "1"
  }
}

resource "azurerm_iothub_shared_access_policy" "example" {
  name                = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  iothub_name         = "${azurerm_iothub.example.name}"

  registry_read = true
  registry_write = true
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the IoT Hub Shared Access Policy resource. Changing this forces a new resource to be created.
 - **resource_group_name** - (Required) The name of the resource group under which the IoT Hub Shared Access Policy resource has to be created. Changing this forces a new resource to be created.
 - **iothub_name** - (Required) The name of the IoT Hub to which this Shared Access Policy belongs. Changing this forces a new resource to be created.
 - **registry_read** - (Optional) Adds `RegistryRead` permission to this Shared Access Account. It allows read access to the identity registry.
 - **registry_write** - (Optional) Adds `RegistryWrite` permission to this Shared Access Account. It allows write access to the identity registry.
-

NOTE When `registry_write` is set to `true`, `registry_read` must also be set to `true`. This is a limitation of the Azure REST API

- `service_connect` - (Optional) Adds `ServiceConnect` permission to this Shared Access Account. It allows sending and receiving on the cloud-side endpoints.
- `device_connect` - (Optional) Adds `DeviceConnect` permission to this Shared Access Account. It allows sending and receiving on the device-side endpoints.

NOTE At least one of `registry_read`, `registry_write`, `service_connect`, `device_connect` permissions must be set to `true`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the IoT Hub Shared Access Policy.
- `primary_key` - The primary key used to create the authentication token.
- `primary_connection_string` - The primary connection string of the Shared Access Policy.
- `secondary_key` - The secondary key used to create the authentication token.
- `secondary_connection_string` - The secondary connection string of the Shared Access Policy.

Import

IoT Hub Shared Access Policies can be imported using the `resource id`, e.g.

```
terraform import azurerm_iothub_shared_access_policy.shared_access_policy1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Devices/IotHubs/hub1/IotHubKeys/shared_access_policy1
```

azurerm_key_vault_access_policy

Manages a Key Vault Access Policy.

NOTE: It's possible to define Key Vault Access Policies both within the `azurerm_key_vault` resource ([/docs/providers/azurerm/r/key_vault.html](/docs/providers/azurerm/r/key_vault.html)) via the `access_policy` block and by using the `azurerm_key_vault_access_policy` resource ([/docs/providers/azurerm/r/key_vault_access_policy.html](/docs/providers/azurerm/r/key_vault_access_policy.html)). However it's not possible to use both methods to manage Access Policies within a KeyVault, since there'll be conflicts.

NOTE: Azure permits a maximum of 1024 Access Policies per Key Vault - more information can be found in this document (<https://docs.microsoft.com/en-us/azure/key-vault/key-vault-secure-your-key-vault#data-plane-access-control>).

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location = "${azurerm_resource_group.test.location}"
}

resource "azurerm_key_vault" "test" {
  name                = "testvault"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku_name = "standard"

  tenant_id = "22222222-2222-2222-2222-222222222222"

  enabled_for_disk_encryption = true

  tags = {
    environment = "Production"
  }
}

resource "azurerm_key_vault_access_policy" "test" {
  key_vault_id = "${azurerm_key_vault.test.id}"

  tenant_id = "00000000-0000-0000-0000-000000000000"
  object_id = "11111111-1111-1111-1111-111111111111"

  key_permissions = [
    "get",
  ]

  secret_permissions = [
    "get",
  ]
}

```

Argument Reference

The following arguments are supported:

- `key_vault_id` - (Required) Specifies the id of the Key Vault resource. Changing this forces a new resource to be created.

NOTE: At this time the Key Vault <-> Key Vault Access Policy associations need to be configured using the field `key_vault_id` or using both fields `vault_name` and `resource_group_name`. These fields are now deprecated and will be removed in favour of `key_vault_id` in the next major version (2.0) of the AzureRM Provider.

- `vault_name` - (Required / **Deprecated**) Specifies the name of the Key Vault resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required / **Deprecated**) The name of the resource group in which to create the namespace. Changing this forces a new resource to be created.

- `tenant_id` - (Required) The Azure Active Directory tenant ID that should be used for authenticating requests to the key vault. Changing this forces a new resource to be created.
- `object_id` - (Required) The object ID of a user, service principal or security group in the Azure Active Directory tenant for the vault. The object ID must be unique for the list of access policies. Changing this forces a new resource to be created.
- `application_id` - (Optional) The object ID of an Application in Azure Active Directory.
- `certificate_permissions` - (Optional) List of certificate permissions, must be one or more from the following: `backup`, `create`, `delete`, `deleteissuers`, `get`, `getissuers`, `import`, `list`, `listissuers`, `managecontacts`, `manageissuers`, `purge`, `recover`, `restore`, `setissuers` and `update`.
- `key_permissions` - (Required) List of key permissions, must be one or more from the following: `backup`, `create`, `decrypt`, `delete`, `encrypt`, `get`, `import`, `list`, `purge`, `recover`, `restore`, `sign`, `unwrapKey`, `update`, `verify` and `wrapKey`.
- `secret_permissions` - (Required) List of secret permissions, must be one or more from the following: `backup`, `delete`, `get`, `list`, `purge`, `recover`, `restore` and `set`.
- `storage_permissions` - (Optional) List of storage permissions, must be one or more from the following: `backup`, `delete`, `deletesas`, `get`, `getsas`, `list`, `listsas`, `purge`, `recover`, `regeneratekey`, `restore`, `set`, `setsas` and `update`.

Attributes Reference

The following attributes are exported:

- `id` - Key Vault Access Policy ID.

NOTE: This Identifier is unique to Terraform and doesn't map to an existing object within Azure.

Import

Key Vault Access Policies can be imported using the Resource ID of the Key Vault, plus some additional metadata.

If both an `object_id` and `application_id` are specified, then the Access Policy can be imported using the following code:

```
terraform import azurerm_key_vault_access_policy.test /subscriptions/00000000-0000-0000-0000-000000000000
/resourceGroups/mygroup1/providers/Microsoft.KeyVault/vaults/test-vault/objectId/11111111-1111-1111-1111-
111111111111/applicationId/22222222-2222-2222-2222-222222222222
```

where `11111111-1111-1111-1111-111111111111` is the `object_id` and `22222222-2222-2222-2222-222222222222` is the `application_id`.

Access Policies with an `object_id` but no `application_id` can be imported using the following command:

```
terraform import azurerm_key_vault_access_policy.test /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.KeyVault/vaults/test-vault/objectId/11111111-1111-1111-1111-  
111111111111
```

where 11111111-1111-1111-1111-111111111111 is the object_id .

NOTE: Both Identifiers are unique to Terraform and don't map to an existing object within Azure.

azurerm_key_vault_certificate

Manages a Key Vault Certificate.

Example Usage (Importing a PFX)

Note: this example assumed the PFX file is located in the same directory at `certificate-to-import.pfx`.

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "test" {
  name      = "key-vault-certificate-example"
  location  = "West Europe"
}

resource "azurerm_key_vault" "test" {
  name                = "keyvaultcertexample"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  tenant_id           = "${data.azurerm_client_config.current.tenant_id}"

  sku_name = "standard"

  access_policy {
    tenant_id = "${data.azurerm_client_config.current.tenant_id}"
    object_id = "${data.azurerm_client_config.current.service_principal_object_id}"

    certificate_permissions = [
      "create",
      "delete",
      "deleteissuers",
      "get",
      "getissuers",
      "import",
      "list",
      "listissuers",
      "managecontacts",
      "manageissuers",
      "setissuers",
      "update",
    ]

    key_permissions = [
      "backup",
      "create",
      "decrypt",
      "delete",
      "encrypt",
      "get",
      "import",
      "list",
      "purge",
      "recover",
      "restore"
    ]
  }
}
```

```

    "restore",
    "sign",
    "unwrapKey",
    "update",
    "verify",
    "wrapKey",
  ]

  secret_permissions = [
    "backup",
    "delete",
    "get",
    "list",
    "purge",
    "recover",
    "restore",
    "set",
  ]
}

tags = {
  environment = "Production"
}
}

resource "azurerm_key_vault_certificate" "test" {
  name          = "imported-cert"
  key_vault_id = "${azurerm_key_vault.test.id}"

  certificate {
    contents = "${filebase64("certificate-to-import.pfx")}"
    password = ""
  }

  certificate_policy {
    issuer_parameters {
      name = "Self"
    }
  }

  key_properties {
    exportable = true
    key_size   = 2048
    key_type   = "RSA"
    reuse_key  = false
  }

  secret_properties {
    content_type = "application/x-pkcs12"
  }
}
}

```

Example Usage (Generating a new certificate)

```

data "azurerm_client_config" "current" {}

```

```

resource "azurerm_resource_group" "test" {
  name      = "key-vault-certificate-example"
  location  = "West Europe"
}

resource "azurerm_key_vault" "test" {
  name                = "keyvaultcertexample"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  tenant_id           = "${data.azurerm_client_config.current.tenant_id}"

  sku_name = "standard"

  access_policy {
    tenant_id = "${data.azurerm_client_config.current.tenant_id}"
    object_id = "${data.azurerm_client_config.current.service_principal_object_id}"

    certificate_permissions = [
      "create",
      "delete",
      "deleteissuers",
      "get",
      "getissuers",
      "import",
      "list",
      "listissuers",
      "managecontacts",
      "manageissuers",
      "setissuers",
      "update",
    ]

    key_permissions = [
      "backup",
      "create",
      "decrypt",
      "delete",
      "encrypt",
      "get",
      "import",
      "list",
      "purge",
      "recover",
      "restore",
      "sign",
      "unwrapKey",
      "update",
      "verify",
      "wrapKey",
    ]

    secret_permissions = [
      "backup",
      "delete",
      "get",
      "list",
      "purge",
      "recover",
      "restore",
      "set",
    ]
  }
}

```

```
    "set",
  ]
}

tags = {
  environment = "Production"
}
}

resource "azurerm_key_vault_certificate" "test" {
  name          = "generated-cert"
  key_vault_id = "${azurerm_key_vault.test.id}"

  certificate_policy {
    issuer_parameters {
      name = "Self"
    }

    key_properties {
      exportable = true
      key_size   = 2048
      key_type   = "RSA"
      reuse_key  = true
    }

    lifetime_action {
      action {
        action_type = "AutoRenew"
      }

      trigger {
        days_before_expiry = 30
      }
    }
  }

  secret_properties {
    content_type = "application/x-pkcs12"
  }

  x509_certificate_properties {
    # Server Authentication = 1.3.6.1.5.5.7.3.1
    # Client Authentication = 1.3.6.1.5.5.7.3.2
    extended_key_usage = ["1.3.6.1.5.5.7.3.1"]

    key_usage = [
      "cRLSign",
      "dataEncipherment",
      "digitalSignature",
      "keyAgreement",
      "keyCertSign",
      "keyEncipherment",
    ]

    subject_alternative_names {
      dns_names = ["internal.contoso.com", "domain.hello.world"]
    }

    subject          = "CN=hello-world"
    validity_in_months = 12
  }
}
```

```
}  
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault Certificate. Changing this forces a new resource to be created.
- `key_vault_id` - (Required) The ID of the Key Vault where the Certificate should be created.
- `certificate` - (Optional) A `certificate` block as defined below, used to Import an existing certificate.
- `certificate_policy` - (Required) A `certificate_policy` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`certificate` supports the following:

- `contents` - (Required) The base64-encoded certificate contents. Changing this forces a new resource to be created.
- `password` - (Optional) The password associated with the certificate. Changing this forces a new resource to be created.

`certificate_policy` supports the following:

- `issuer_parameters` - (Required) A `issuer_parameters` block as defined below.
- `key_properties` - (Required) A `key_properties` block as defined below.
- `lifetime_action` - (Optional) A `lifetime_action` block as defined below.
- `secret_properties` - (Required) A `secret_properties` block as defined below.
- `x509_certificate_properties` - (Optional) A `x509_certificate_properties` block as defined below.

`issuer_parameters` supports the following:

- `name` - (Required) The name of the Certificate Issuer. Possible values include `Self`, or the name of a certificate issuing authority supported by Azure. Changing this forces a new resource to be created.

`key_properties` supports the following:

- `exportable` - (Required) Is this Certificate Exportable? Changing this forces a new resource to be created.
- `key_size` - (Required) The size of the Key used in the Certificate. Possible values include `2048` and `4096`. Changing this forces a new resource to be created.
- `key_type` - (Required) Specifies the Type of Key, such as `RSA`. Changing this forces a new resource to be created.
- `reuse_key` - (Required) Is the key reusable? Changing this forces a new resource to be created.

`lifetime_action` supports the following:

- `action` - (Required) A `action` block as defined below.

- `trigger` - (Required) A `trigger` block as defined below.

`action` supports the following:

- `action_type` - (Required) The Type of action to be performed when the lifetime trigger is triggered. Possible values include `AutoRenew` and `EmailContacts`. Changing this forces a new resource to be created.

`trigger` supports the following:

- `days_before_expiry` - (Optional) The number of days before the Certificate expires that the action associated with this Trigger should run. Changing this forces a new resource to be created. Conflicts with `lifetime_percentage`.
- `lifetime_percentage` - (Optional) The percentage at which during the Certificate's Lifetime the action associated with this Trigger should run. Changing this forces a new resource to be created. Conflicts with `days_before_expiry`.

`secret_properties` supports the following:

- `content_type` - (Required) The Content-Type of the Certificate, such as `application/x-pkcs12` for a PFX or `application/x-pem-file` for a PEM. Changing this forces a new resource to be created.

`x509_certificate_properties` supports the following:

- `extended_key_usage` - (Optional) A list of Extended/Enhanced Key Usages. Changing this forces a new resource to be created.
- `key_usage` - (Required) A list of uses associated with this Key. Possible values include `cRLSign`, `dataEncipherment`, `decipherOnly`, `digitalSignature`, `encipherOnly`, `keyAgreement`, `keyCertSign`, `keyEncipherment` and `nonRepudiation` and are case-sensitive. Changing this forces a new resource to be created.
- `subject` - (Required) The Certificate's Subject. Changing this forces a new resource to be created.
- `subject_alternative_names` - (Optional) A `subject_alternative_names` block as defined below.
- `validity_in_months` - (Required) The Certificate's Validity Period in Months. Changing this forces a new resource to be created.

`subject_alternative_names` supports the following:

- `dns_names` - (Optional) A list of alternative DNS names (FQDNs) identified by the Certificate. Changing this forces a new resource to be created.
- `emails` - (Optional) A list of email addresses identified by this Certificate. Changing this forces a new resource to be created.
- `upns` - (Optional) A list of User Principal Names identified by the Certificate. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Key Vault Certificate ID.
- `secret_id` - The ID of the associated Key Vault Secret.
- `version` - The current version of the Key Vault Certificate.

- `certificate_data` - The raw Key Vault Certificate data represented as a hexadecimal string.
- `thumbprint` - The X509 Thumbprint of the Key Vault Certificate represented as a hexadecimal string.

Import

Key Vault Certificates can be imported using the `resource_id`, e.g.

```
terraform import azurerm_key_vault_certificate.test https://example-keyvault.vault.azure.net/certificates/example/fdf067c93bbb4b22bffa4d8b7a9a56217
```

azurerm_key_vault

Manages a Key Vault.

NOTE: It's possible to define Key Vault Access Policies both within the `azurerm_key_vault` resource (/docs/providers/azurerm/r/key_vault.html) via the `access_policy` block and by using the `azurerm_key_vault_access_policy` resource (/docs/providers/azurerm/r/key_vault_access_policy.html). However it's not possible to use both methods to manage Access Policies within a KeyVault, since there'll be conflicts.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_key_vault" "test" {
  name                        = "testvault"
  location                   = "${azurerm_resource_group.test.location}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  enabled_for_disk_encryption = true
  tenant_id                  = "d6e396d0-5584-41dc-9fc0-268df99bc610"

  sku_name = "standard"

  access_policy {
    tenant_id = "d6e396d0-5584-41dc-9fc0-268df99bc610"
    object_id = "d746815a-0433-4a21-b95d-fc437d2d475b"

    key_permissions = [
      "get",
    ]

    secret_permissions = [
      "get",
    ]

    storage_permissions = [
      "get",
    ]
  }

  network_acls {
    default_action = "Deny"
    bypass         = "AzureServices"
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Key Vault. Changing this forces a new resource to be created.
- `sku` - (Optional **Deprecated**) A `sku` block as described below.
- `sku_name` - (Optional) The Name of the SKU used for this Key Vault. Possible values are `standard` and `premium`.
- `tenant_id` - (Required) The Azure Active Directory tenant ID that should be used for authenticating requests to the key vault.
- `access_policy` - (Optional) A list (</docs/configuration/attr-as-blocks.html>) of up to 16 objects describing access policies, as described below.

NOTE: It's possible to define Key Vault Access Policies both within the `azurerm_key_vault` resource ([/docs/providers/azurerm/r/key_vault.html](/docs/providers/azurerm/r/key_vault.html)) via the `access_policy` block and by using the `azurerm_key_vault_access_policy` resource ([/docs/providers/azurerm/r/key_vault_access_policy.html](/docs/providers/azurerm/r/key_vault_access_policy.html)). However it's not possible to use both methods to manage Access Policies within a KeyVault, since there'll be conflicts.

- `enabled_for_deployment` - (Optional) Boolean flag to specify whether Azure Virtual Machines are permitted to retrieve certificates stored as secrets from the key vault. Defaults to `false`.
- `enabled_for_disk_encryption` - (Optional) Boolean flag to specify whether Azure Disk Encryption is permitted to retrieve secrets from the vault and unwrap keys. Defaults to `false`.
- `enabled_for_template_deployment` - (Optional) Boolean flag to specify whether Azure Resource Manager is permitted to retrieve secrets from the key vault. Defaults to `false`.
- `network_acls` - (Optional) A `network_acls` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The Name of the SKU used for this Key Vault. Possible values are `standard` and `premium`.

A `access_policy` block supports the following:

Elements of `access_policy` support:

- `tenant_id` - (Required) The Azure Active Directory tenant ID that should be used for authenticating requests to the key vault. Must match the `tenant_id` used above.
- `object_id` - (Required) The object ID of a user, service principal or security group in the Azure Active Directory tenant for the vault. The object ID must be unique for the list of access policies.

- `application_id` - (Optional) The object ID of an Application in Azure Active Directory.
- `certificate_permissions` - (Optional) List of certificate permissions, must be one or more from the following: `backup`, `create`, `delete`, `deleteissuers`, `get`, `getissuers`, `import`, `list`, `listissuers`, `managecontacts`, `manageissuers`, `purge`, `recover`, `restore`, `setissuers` and `update`.
- `key_permissions` - (Optional) List of key permissions, must be one or more from the following: `backup`, `create`, `decrypt`, `delete`, `encrypt`, `get`, `import`, `list`, `purge`, `recover`, `restore`, `sign`, `unwrapKey`, `update`, `verify` and `wrapKey`.
- `secret_permissions` - (Optional) List of secret permissions, must be one or more from the following: `backup`, `delete`, `get`, `list`, `purge`, `recover`, `restore` and `set`.
- `storage_permissions` - (Optional) List of storage permissions, must be one or more from the following: `backup`, `delete`, `deletesas`, `get`, `getsas`, `list`, `listsas`, `purge`, `recover`, `regeneratekey`, `restore`, `set`, `setsas` and `update`.

A `network_acls` block supports the following:

- `bypass` - (Required) Specifies which traffic can bypass the network rules. Possible values are `AzureServices` and `None`.
- `default_action` - (Required) The Default Action to use when no rules match from `ip_rules` / `virtual_network_subnet_ids`. Possible values are `Allow` and `Deny`.
- `ip_rules` - (Optional) One or more IP Addresses, or CIDR Blocks which should be able to access the Key Vault.
- `virtual_network_subnet_ids` - (Optional) One or more Subnet ID's which should be able to access this Key Vault.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Key Vault.
- `vault_uri` - The URI of the Key Vault, used for performing operations on keys and secrets.

Import

Key Vault's can be imported using the `resource id`, e.g.

```
terraform import azurerm_key_vault.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.KeyVault/vaults/vault1
```

azurerm_key_vault_key

Manages a Key Vault Key.

Example Usage

```

data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "test" {
  name      = "my-resource-group"
  location  = "West US"
}

resource "random_id" "server" {
  keepers = {
    ami_id = 1
  }

  byte_length = 8
}

resource "azurerm_key_vault" "test" {
  name                = "keyvaultkeyexample"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  tenant_id           = "${data.azurerm_client_config.current.tenant_id}"

  sku_name = "premium"

  access_policy {
    tenant_id = "${data.azurerm_client_config.current.tenant_id}"
    object_id = "${data.azurerm_client_config.current.service_principal_object_id}"

    key_permissions = [
      "create",
      "get",
    ]

    secret_permissions = [
      "set",
    ]
  }

  tags = {
    environment = "Production"
  }
}

resource "azurerm_key_vault_key" "generated" {
  name          = "generated-certificate"
  key_vault_id = "${azurerm_key_vault.test.id}"
  key_type      = "RSA"
  key_size      = 2048

  key_opts = [
    "decrypt",
    "encrypt",
    "sign",
    "unwrapKey",
    "verify",
    "wrapKey",
  ]
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault Key. Changing this forces a new resource to be created.
- `key_vault_id` - (Required) The ID of the Key Vault where the Key should be created. Changing this forces a new resource to be created.
- `key_type` - (Required) Specifies the Key Type to use for this Key Vault Key. Possible values are `EC` (Elliptic Curve), `EC-HSM`, `Oct` (Octet), `RSA` and `RSA-HSM`. Changing this forces a new resource to be created.
- `key_size` - (Optional) Specifies the Size of the RSA key to create in bytes. For example, 1024 or 2048. *Note:* This field is required if `key_type` is `RSA` or `RSA-HSM`. Changing this forces a new resource to be created.
- `curve` - (Optional) Specifies the curve to use when creating an `EC` key. Possible values are `P-256`, `P-384`, `P-521`, and `SECP256K1`. This field will be required in a future release if `key_type` is `EC` or `EC-HSM`. The API will default to `P-256` if nothing is specified. Changing this forces a new resource to be created.
- `key_opts` - (Required) A list of JSON web key operations. Possible values include: `decrypt`, `encrypt`, `sign`, `unwrapKey`, `verify` and `wrapKey`. Please note these values are case sensitive.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Key Vault Key ID.
- `version` - The current version of the Key Vault Key.
- `n` - The RSA modulus of this Key Vault Key.
- `e` - The RSA public exponent of this Key Vault Key.
- `x` - The EC X component of this Key Vault Key.
- `y` - The EC Y component of this Key Vault Key.

Import

Key Vault Key which is Enabled can be imported using the `resource id`, e.g.

```
terraform import azurerm_key_vault_key.test https://example-keyvault.vault.azure.net/keys/example/fdf067c93bbb4b22bff4d8b7a9a56217
```

azurerm_key_vault_secret

Manages a Key Vault Secret.

Note: All arguments including the secret value will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```

data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "test" {
  name      = "my-resource-group"
  location  = "West US"
}

resource "random_id" "server" {
  keepers = {
    ami_id = 1
  }

  byte_length = 8
}

resource "azurerm_key_vault" "test" {
  name                = "${format("%s%s", "kv", random_id.server.hex)}"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  tenant_id            = "${data.azurerm_client_config.current.tenant_id}"

  sku_name = "premium"

  access_policy {
    tenant_id = "${data.azurerm_client_config.current.tenant_id}"
    object_id = "${data.azurerm_client_config.current.service_principal_object_id}"

    key_permissions = [
      "create",
      "get",
    ]

    secret_permissions = [
      "set",
      "get",
      "delete",
    ]
  }

  tags = {
    environment = "Production"
  }
}

resource "azurerm_key_vault_secret" "test" {
  name      = "secret-sauce"
  value     = "szechuan"
  key_vault_id = "${azurerm_key_vault.test.id}"

  tags = {
    environment = "Production"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Key Vault Secret. Changing this forces a new resource to be created.
- `value` - (Required) Specifies the value of the Key Vault Secret.
- `key_vault_id` - (Required) The ID of the Key Vault where the Secret should be created.
- `content_type` - (Optional) Specifies the content type for the Key Vault Secret.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Key Vault Secret ID.
- `version` - The current version of the Key Vault Secret.

Import

Key Vault Secrets which are Enabled can be imported using the `resource id`, e.g.

```
terraform import azurerm_key_vault_secret.test https://example-keyvault.vault.azure.net/secrets/example/df067c93bbb4b22bffa4d8b7a9a56217
```

azurerm_kubernetes_cluster

Manages a Managed Kubernetes Cluster (also known as AKS / Azure Kubernetes Service)

Note: All arguments including the client secret will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

This example provisions a basic Managed Kubernetes Cluster. Other examples of the `azurerm_kubernetes_cluster` resource can be found in the `./examples/kubernetes` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/kubernetes>)

```

resource "azurerm_resource_group" "test" {
  name      = "acctestRG1"
  location  = "East US"
}

resource "azurerm_kubernetes_cluster" "test" {
  name                  = "acctestaks1"
  location              = azurerm_resource_group.test.location
  resource_group_name   = azurerm_resource_group.test.name
  dns_prefix            = "acctestagent1"

  agent_pool_profile {
    name                = "default"
    count               = 1
    vm_size              = "Standard_D1_v2"
    os_type              = "Linux"
    os_disk_size_gb     = 30
  }

  agent_pool_profile {
    name                = "pool2"
    count               = 1
    vm_size              = "Standard_D2_v2"
    os_type              = "Linux"
    os_disk_size_gb     = 30
  }

  service_principal {
    client_id           = "000000000-0000-0000-0000-000000000000"
    client_secret       = "00000000000000000000000000000000"
  }

  tags = {
    Environment = "Production"
  }
}

output "client_certificate" {
  value = "${azurerm_kubernetes_cluster.test.kube_config.0.client_certificate}"
}

output "kube_config" {
  value = "${azurerm_kubernetes_cluster.test.kube_config_raw}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Managed Kubernetes Cluster to create. Changing this forces a new resource to be created.
- `location` - (Required) The location where the Managed Kubernetes Cluster should be created. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) Specifies the Resource Group where the Managed Kubernetes Cluster should exist. Changing this forces a new resource to be created.
- `agent_pool_profile` - (Required) One or more `agent_pool_profile` blocks as defined below.
- `dns_prefix` - (Required) DNS prefix specified when creating the managed cluster. Changing this forces a new resource to be created.

NOTE: The `dns_prefix` must contain between 3 and 45 characters, and can contain only letters, numbers, and hyphens. It must start with a letter and must end with a letter or a number.

- `service_principal` - (Required) A `service_principal` block as documented below.

A `aci_connector_linux` block supports the following:

- `enabled` - (Required) Is the virtual node addon enabled?
- `subnet_name` - (Optional) The subnet name for the virtual nodes to run. This is required when `aci_connector_linux` `enabled` argument is set to `true`.

Note: AKS will add a delegation to the subnet named here. To prevent further runs from failing you should make sure that the subnet you create for virtual nodes has a delegation, like so.

```
resource "azurerm_subnet" "virtual" {

  ...

  delegation {
    name = "aciDelegation"
    service_delegation {
      name      = "Microsoft.ContainerInstance/containerGroups"
      actions = ["Microsoft.Network/virtualNetworks/subnets/action"]
    }
  }
}
```

- `addon_profile` - (Optional) A `addon_profile` block.
- `api_server_authorized_ip_ranges` - (Optional) The IP ranges to whitelist for incoming traffic to the masters.

NOTE: `api_server_authorized_ip_ranges` Is currently in Preview on an opt-in basis. To use it, enable feature `APIServerSecurityPreview` for namespace `Microsoft.ContainerService`. For an example of how to enable a Preview feature, please visit [How to enable the Azure Firewall Public Preview \(https://docs.microsoft.com/en-us/azure/firewall/public-preview\)](https://docs.microsoft.com/en-us/azure/firewall/public-preview)

- `kubernetes_version` - (Optional) Version of Kubernetes specified when creating the AKS managed cluster. If not specified, the latest recommended version will be used at provisioning time (but won't auto-upgrade).
- `linux_profile` - (Optional) A `linux_profile` block.

- `windows_profile` - (Optional) A `windows_profile` block.
- `network_profile` - (Optional) A `network_profile` block.

NOTE: If `network_profile` is not defined, `kubenet` profile will be used by default.

- `role_based_access_control` - (Optional) A `role_based_access_control` block. Changing this forces a new resource to be created.
- `enable_pod_security_policy` - (Optional) Whether Pod Security Policies are enabled. Note that this also requires role based access control to be enabled.

NOTE: Support for `enable_pod_security_policy` is currently in Preview on an opt-in basis. To use it, enable feature `PodSecurityPolicyPreview` for namespace `Microsoft.ContainerService`. For an example of how to enable a Preview feature, please visit [Register scale set feature provider \(https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler#register-scale-set-feature-provider\)](https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler#register-scale-set-feature-provider).

- `node_resource_group` - (Optional) The name of the Resource Group where the the Kubernetes Nodes should exist. Changing this forces a new resource to be created.

NOTE: Azure requires that a new, non-existent Resource Group is used, as otherwise the provisioning of the Kubernetes Service will fail.

- `tags` - (Optional) A mapping of tags to assign to the resource.

A `addon_profile` block supports the following:

- `aci_connector_linux` - (Optional) A `aci_connector_linux` block. For more details, please visit [Create and configure an AKS cluster to use virtual nodes \(https://docs.microsoft.com/en-us/azure/aks/virtual-nodes-portal\)](https://docs.microsoft.com/en-us/azure/aks/virtual-nodes-portal).
- `http_application_routing` - (Optional) A `http_application_routing` block.
- `oms_agent` - (Optional) A `oms_agent` block. For more details, please visit [How to onboard Azure Monitor for containers \(https://docs.microsoft.com/en-us/azure/monitoring/monitoring-container-insights-onboard\)](https://docs.microsoft.com/en-us/azure/monitoring/monitoring-container-insights-onboard).
- `kube_dashboard` - (Optional) A `kube_dashboard` block.
- `azure_policy` - (Optional) A `azure_policy` block. For more details please visit [Understand Azure Policy for Azure Kubernetes Service \(https://docs.microsoft.com/en-ie/azure/governance/policy/concepts/rego-for-aks\)](https://docs.microsoft.com/en-ie/azure/governance/policy/concepts/rego-for-aks)

NOTE: Azure Policy for Azure Kubernetes Service is currently in preview and not available to subscriptions that have not opted-in (<https://docs.microsoft.com/en-us/azure/governance/policy/concepts/rego-for-aks?toc=/azure/aks/toc.json>) to join Azure Policy preview.

A `agent_pool_profile` block supports the following:

- `name` - (Required) Unique name of the Agent Pool Profile in the context of the Subscription and Resource Group. Changing this forces a new resource to be created.

- `count` - (Optional) Number of Agents (VMs) in the Pool. Possible values must be in the range of 1 to 100 (inclusive). Defaults to 1 .
- `vm_size` - (Required) The size of each VM in the Agent Pool (e.g. `Standard_F1`). Changing this forces a new resource to be created.
- `availability_zones` - (Optional) Availability zones for nodes. The property `type` of the `agent_pool_profile` must be set to `VirtualMachineScaleSets` in order to use availability zones.
- `enable_auto_scaling` - (Optional) Whether to enable auto-scaler (<https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler>). Note that auto scaling feature requires the that the `type` is set to `VirtualMachineScaleSets`
- `min_count` - (Optional) Minimum number of nodes for auto-scaling
- `max_count` - (Optional) Maximum number of nodes for auto-scaling
- `max_pods` - (Optional) The maximum number of pods that can run on each agent. Changing this forces a new resource to be created.
- `os_disk_size_gb` - (Optional) The Agent Operating System disk size in GB. Changing this forces a new resource to be created.
- `os_type` - (Optional) The Operating System used for the Agents. Possible values are `Linux` and `Windows` . Changing this forces a new resource to be created. Defaults to `Linux` .
- `type` - (Optional) Type of the Agent Pool. Possible values are `AvailabilitySet` and `VirtualMachineScaleSets` . Changing this forces a new resource to be created. Defaults to `AvailabilitySet` .
- `vnet_subnet_id` - (Optional) The ID of the Subnet where the Agents in the Pool should be provisioned. Changing this forces a new resource to be created.

NOTE: A route table should be configured on this Subnet.

- `node_taints` - (Optional) A list of Kubernetes taints which should be applied to nodes in the agent pool (e.g `key=value:NoSchedule`)

A `azure_active_directory` block supports the following:

- `client_app_id` - (Required) The Client ID of an Azure Active Directory Application. Changing this forces a new resource to be created.
- `server_app_id` - (Required) The Server ID of an Azure Active Directory Application. Changing this forces a new resource to be created.
- `server_app_secret` - (Required) The Server Secret of an Azure Active Directory Application. Changing this forces a new resource to be created.
- `tenant_id` - (Optional) The Tenant ID used for Azure Active Directory Application. If this isn't specified the Tenant ID of the current Subscription is used. Changing this forces a new resource to be created.

A `http_application_routing` block supports the following:

- `enabled` (Required) Is HTTP Application Routing Enabled? Changing this forces a new resource to be created.

A `linux_profile` block supports the following:

- `admin_username` - (Required) The Admin Username for the Cluster. Changing this forces a new resource to be created.
- `ssh_key` - (Required) An `ssh_key` block. Only one is currently allowed. Changing this forces a new resource to be created.

A `windows_profile` block supports the following:

- `admin_username` - (Required) The Admin Username for Windows VMs.
- `admin_password` - (Required) The Admin Password for Windows VMs.

A `network_profile` block supports the following:

- `network_plugin` - (Required) Network plugin to use for networking. Currently supported values are `azure` and `kubenet`. Changing this forces a new resource to be created.

NOTE: When `network_plugin` is set to `azure` - the `vnet_subnet_id` field in the `agent_pool_profile` block must be set and `pod_cidr` must not be set.

- `network_policy` - (Optional) Sets up network policy to be used with Azure CNI. Network policy allows us to control the traffic flow between pods (<https://docs.microsoft.com/en-us/azure/aks/use-network-policies>). This field can only be set when `network_plugin` is set to `azure`. Currently supported values are `calico` and `azure`. Changing this forces a new resource to be created.
- `dns_service_ip` - (Optional) IP address within the Kubernetes service address range that will be used by cluster service discovery (kube-dns). This is required when `network_plugin` is set to `azure`. Changing this forces a new resource to be created.
- `docker_bridge_cidr` - (Optional) IP address (in CIDR notation) used as the Docker bridge IP address on nodes. This is required when `network_plugin` is set to `azure`. Changing this forces a new resource to be created.
- `pod_cidr` - (Optional) The CIDR to use for pod IP addresses. This field can only be set when `network_plugin` is set to `kubenet`. Changing this forces a new resource to be created.
- `service_cidr` - (Optional) The Network Range used by the Kubernetes service. This is required when `network_plugin` is set to `azure`. Changing this forces a new resource to be created.

NOTE: This range should not be used by any network element on or connected to this VNet. Service address CIDR must be smaller than /12.

Examples of how to use AKS with Advanced Networking (<https://docs.microsoft.com/en-us/azure/aks/networking-overview#advanced-networking>) can be found in the `./examples/kubernetes/` directory in the Github repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/kubernetes>).

- `load_balancer_sku` - (Optional) Specifies the SKU of the Load Balancer used for this Kubernetes Cluster. Possible

values are basic and standard. Defaults to basic.

A `oms_agent` block supports the following:

- `enabled` - (Required) Is the OMS Agent Enabled?
 - `log_analytics_workspace_id` - (Optional) The ID of the Log Analytics Workspace which the OMS Agent should send data to. Must be present if `enabled` is `true`.
-

A `kube_dashboard` block supports the following:

* `enabled` - (Required) Is the Kubernetes Dashboard enabled?

A `azure_policy` block supports the following:

- `enabled` - (Required) Is the Azure Policy for Kubernetes Add On enabled?
-

A `role_based_access_control` block supports the following:

- `azure_active_directory` - (Optional) An `azure_active_directory` block. Changing this forces a new resource to be created.
 - `enabled` - (Required) Is Role Based Access Control Enabled? Changing this forces a new resource to be created.
-

A `service_principal` block supports the following:

- `client_id` - (Required) The Client ID for the Service Principal.
 - `client_secret` - (Required) The Client Secret for the Service Principal.
-

A `ssh_key` block supports the following:

- `key_data` - (Required) The Public SSH Key used to access the cluster. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Kubernetes Managed Cluster ID.
- `fqdn` - The FQDN of the Azure Kubernetes Managed Cluster.
- `kube_admin_config` - A `kube_admin_config` block as defined below. This is only available when Role Based Access Control with Azure Active Directory is enabled.

- kube_admin_config_raw - Raw Kubernetes config for the admin account to be used by kubectl (<https://kubernetes.io/docs/reference/kubectl/overview/>) and other compatible tools. This is only available when Role Based Access Control with Azure Active Directory is enabled.
- kube_config - A kube_config block as defined below.
- kube_config_raw - Raw Kubernetes config to be used by kubectl (<https://kubernetes.io/docs/reference/kubectl/overview/>) and other compatible tools
- http_application_routing - A http_application_routing block as defined below.
- node_resource_group - The auto-generated Resource Group which contains the resources for this Managed Kubernetes Cluster.

A http_application_routing block exports the following:

- http_application_routing_zone_name - The Zone Name of the HTTP Application Routing.

The kube_admin_config and kube_config blocks export the following::

- client_key - Base64 encoded private key used by clients to authenticate to the Kubernetes cluster.
- client_certificate - Base64 encoded public certificate used by clients to authenticate to the Kubernetes cluster.
- cluster_ca_certificate - Base64 encoded public CA certificate used as the root of trust for the Kubernetes cluster.
- host - The Kubernetes cluster server host.
- username - A username used to authenticate to the Kubernetes cluster.
- password - A password or token used to authenticate to the Kubernetes cluster.

NOTE: It's possible to use these credentials with the Kubernetes Provider ([/docs/providers/kubernetes/index.html](https://docs/providers/kubernetes/index.html)) like so:

```
provider "kubernetes" {
  host          = "${azurerm_kubernetes_cluster.main.kube_config.0.host}"
  username      = "${azurerm_kubernetes_cluster.main.kube_config.0.username}"
  password      = "${azurerm_kubernetes_cluster.main.kube_config.0.password}"
  client_certificate = "${base64decode(azurerm_kubernetes_cluster.main.kube_config.0.client_certificate)}"
  client_key     = "${base64decode(azurerm_kubernetes_cluster.main.kube_config.0.client_key)}"
  cluster_ca_certificate = "${base64decode(azurerm_kubernetes_cluster.main.kube_config.0.cluster_ca_certificate)}"
}
```

Import

Managed Kubernetes Clusters can be imported using the `resource_id`, e.g.

```
terraform import azurerm_kubernetes_cluster.cluster1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/group1/providers/Microsoft.ContainerService/managedClusters/cluster1
```

azurerm_kusto_cluster

Manages a Kusto (also known as Azure Data Explorer) Cluster

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "my-kusto-cluster-rg"
  location  = "East US"
}

resource "azurerm_kusto_cluster" "test" {
  name                = "kustocluster"
  location            = "${azurerm_resource_group.rg.location}"
  resource_group_name = "${azurerm_resource_group.rg.name}"

  sku {
    name      = "Standard_D13_v2"
    capacity = 2
  }

  tags = {
    Environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Kusto Cluster to create. Changing this forces a new resource to be created.
- `location` - (Required) The location where the Kusto Cluster should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the Resource Group where the Kusto Cluster should exist. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) The name of the SKU. Valid values are: `Dev(No SLA)_Standard_D11_v2`, `Standard_D11_v2`, `Standard_D12_v2`, `Standard_D13_v2`, `Standard_D14_v2`, `Standard_DS13_v2+1TB_PS`, `Standard_DS13_v2+2TB_PS`, `Standard_DS14_v2+3TB_PS`, `Standard_DS14_v2+4TB_PS`, `Standard_L16s`, `Standard_L4s` and `Standard_L8s`
- `capacity` - (Required) Specifies the node count for the cluster. Boundaries depend on the sku name.

Attributes Reference

The following attributes are exported:

- `id` - The Kusto Cluster ID.
 - `uri` - The FQDN of the Azure Kusto Cluster.
 - `data_ingestion_uri` - The Kusto Cluster URI to be used for data ingestion.
-

Import

Kusto Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_kusto_cluster.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Kusto/Clusters/cluster1
```

azurerm_kusto_database

Manages a Kusto (also known as Azure Data Explorer) Database

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "my-kusto-rg"
  location  = "East US"
}

resource "azurerm_kusto_cluster" "cluster" {
  name            = "kustocluster"
  location        = "${azurerm_resource_group.rg.location}"
  resource_group_name = "${azurerm_resource_group.rg.name}"

  sku {
    name      = "Standard_D13_v2"
    capacity = 2
  }
}

resource "azurerm_kusto_database" "database" {
  name                = "my-kusto-database"
  resource_group_name = azurerm_resource_group.rg.name
  location            = azurerm_resource_group.rg.location
  cluster_name        = azurerm_kusto_cluster.cluster.name

  hot_cache_period    = "P7D"
  soft_delete_period = "P31D"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Kusto Database to create. Changing this forces a new resource to be created.
- `location` - (Required) The location where the Kusto Database should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the Resource Group where the Kusto Database should exist. Changing this forces a new resource to be created.
- `cluster_name` - (Required) Specifies the name of the Kusto Cluster this database will be added to. Changing this forces a new resource to be created.
- `hot_cache_period` - (Optional) The time the data that should be kept in cache for fast queries as ISO 8601 timespan. Default is unlimited. For more information see: ISO 8601 Timespan (https://en.wikipedia.org/wiki/ISO_8601#Durations)

- `soft_delete_period` - (Optional) The time the data should be kept before it stops being accessible to queries as ISO 8601 timespan. Default is unlimited. For more information see: ISO 8601 Timespan (https://en.wikipedia.org/wiki/ISO_8601#Durations)

Attributes Reference

The following attributes are exported:

- `id` - The Kusto Cluster ID.
- `size` - The size of the database in bytes.

Import

Kusto Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_kusto_database.test /subscriptions/00000000-0000-0000-0000-000000000000/resource
Groups/group1/providers/Microsoft.Kusto/Clusters/cluster1/Databases/database1
```

azurerm_kusto_eventhub_data_connection

Manages a Kusto (also known as Azure Data Explorer) EventHub Data Connection

Example Usage

```

resource "azurerm_resource_group" "rg" {
  name      = "my-kusto-rg"
  location  = "East US"
}

resource "azurerm_kusto_cluster" "cluster" {
  name            = "kustocluster"
  location        = "${azurerm_resource_group.rg.location}"
  resource_group_name = "${azurerm_resource_group.rg.name}"

  sku {
    name      = "Standard_D13_v2"
    capacity = 2
  }
}

resource "azurerm_kusto_database" "database" {
  name            = "my-kusto-database"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  location        = "${azurerm_resource_group.rg.location}"
  cluster_name    = "${azurerm_kusto_cluster.cluster.name}"
  hot_cache_period = "P7D"
  soft_delete_period = "P31D"
}

resource "azurerm_eventhub_namespace" "eventhub_ns" {
  name            = "my-eventhub-ns"
  location        = "${azurerm_resource_group.rg.location}"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  sku             = "Standard"
}

resource "azurerm_eventhub" "eventhub" {
  name            = "my-eventhub"
  namespace_name  = "${azurerm_eventhub_namespace.eventhub_ns.name}"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  partition_count = 1
  message_retention = 1
}

resource "azurerm_eventhub_consumer_group" "consumer_group" {
  name            = "my-eventhub-consumergroup"
  namespace_name  = "${azurerm_eventhub_namespace.eventhub_ns.name}"
  eventhub_name    = "${azurerm_eventhub.eventhub.name}"
  resource_group_name = "${azurerm_resource_group.rg.name}"
}

resource "azurerm_kusto_eventhub_data_connection" "eventhub_connection" {
  name            = "my-kusto-eventhub-data-connection"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  location        = "${azurerm_resource_group.rg.location}"
  cluster_name    = "${azurerm_kusto_cluster.cluster.name}"
  database_name    = "${azurerm_kusto_database.database.name}"

  eventhub_id      = "${azurerm_eventhub.eventhub.id}"
  consumer_group    = "${azurerm_eventhub_consumer_group.consumer_group.name}"

  table_name        = "my-table"           #(Optional)
  mapping_rule_name = "my-table-mapping" #(Optional)
  data_format        = "JSON"              #(Optional)
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Kusto EventHub Data Connection to create. Changing this forces a new resource to be created.
- `location` - (Required) The location where the Kusto Database should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) Specifies the Resource Group where the Kusto Database should exist. Changing this forces a new resource to be created.
- `cluster_name` - (Required) Specifies the name of the Kusto Cluster this data connection will be added to. Changing this forces a new resource to be created.
- `database_name` - (Required) Specifies the name of the Kusto Database this data connection will be added to. Changing this forces a new resource to be created.
- `eventhub_id` - (Required) Specifies the resource id of the EventHub this data connection will use for ingestion. Changing this forces a new resource to be created.
- `consumer_group` - (Required) Specifies the EventHub consumer group this data connection will use for ingestion. Changing this forces a new resource to be created.
- `table_name` - (Optional) Specifies the target table name used for the message ingestion. Table must exist before resource is created.
- `mapping_rule_name` - (Optional) Specifies the mapping rule used for the message ingestion. Mapping rule must exist before resource is created.
- `data_format` - (Optional) Specifies the data format of the EventHub messages. Allowed values: AVRO , CSV , JSON , MULTIJSON , PSV , RAW , SCSV , SINGLEJSON , SOHSV , TSV and TXT

Attributes Reference

The following attributes are exported:

- `id` - The EventHub Data Connection ID.

Import

Kusto EventHub Data Connections can be imported using the `resource id` , e.g.

```
terraform import azurerm_kusto_eventhub_data_connection.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Kusto/Clusters/cluster1/Databases/database1/DataConnections/eventHubConnection1
```

azurerm_linux_virtual_machine_scale_set

NOTE: This resource is in Beta and as such the Schema can change in Minor versions of the Provider.

NOTE: All arguments including the administrator login and password will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Manages a Linux Virtual Machine Scale Set.

NOTE: This resource does not support Unmanaged Disks. If you need to use Unmanaged Disks you can continue to use the `azurerm_virtual_machine_scale_set` resource ([/docs/providers/azurerm/r/virtual_machine_scale_set.html](/docs/providers/azurerm/r/virtual_machine_scale_set.html)) instead

NOTE: Terraform will automatically update & reimage the nodes in the Scale Set if Required during an Update (for example, when changing Sku) - you can opt out of this by setting the `terraform_should_roll_instances_when_required` field to `false`.

Example Usage

This example provisions a basic Linux Virtual Machine Scale Set on an internal network. Additional examples of how to use the `azurerm_linux_virtual_machine_scale_set` resource can be found in the `./examples/vm-scale-set/linux`` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/vm-scale-set/linux>).

```

resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name                = "example-network"
  resource_group_name = azurerm_resource_group.example.name
  location            = azurerm_resource_group.example.location
  address_space       = ["10.0.0.0/16"]
}

resource "azurerm_subnet" "internal" {
  name                = "internal"
  resource_group_name = azurerm_resource_group.example.name
  virtual_network_name = azurerm_virtual_network.example.name
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_linux_virtual_machine_scale_set" "example" {
  name                = "example-vmss"
  resource_group_name = azurerm_resource_group.example.name
  location            = azurerm_resource_group.example.location
  sku                 = "Standard_F2"
  instances            = 1
  admin_username      = "adminuser"

  admin_ssh_key {
    username   = "adminuser"
    public_key = file("~/ssh/id_rsa.pub")
  }

  source_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }

  os_disk {
    storage_account_type = "Standard_LRS"
    caching               = "ReadWrite"
  }

  network_interface {
    name      = "example"
    primary   = true

    ip_configuration {
      name       = "internal"
      primary    = true
      subnet_id  = azurerm_subnet.internal.id
    }
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Linux Virtual Machine Scale Set. Changing this forces a new resource to be created.
- `location` - (Required) The Azure location where the Linux Virtual Machine Scale Set should exist. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Linux Virtual Machine Scale Set should be exist. Changing this forces a new resource to be created.
- `admin_username` - (Required) The username of the local administrator on each Virtual Machine Scale Set instance. Changing this forces a new resource to be created.
- `instances` - (Required) The number of Virtual Machines in the Scale Set.

NOTE: If you're using AutoScaling, you may wish to use Terraform's `ignore_changes` functionality (https://www.terraform.io/docs/configuration/resources.html#ignore_changes) to ignore changes to this field.

- `sku` - (Required) The Virtual Machine SKU for the Scale Set, such as `Standard_F2`.
- `network_interface` - (Required) One or more `network_interface` blocks as defined below.
- `os_disk` - (Required) An `os_disk` block as defined below.

- `additional_capabilities` - (Optional) A `additional_capabilities` block as defined below.
- `admin_password` - (Optional) The Password which should be used for the local-administrator on this Virtual Machine. Changing this forces a new resource to be created.

NOTE: When an `admin_password` is specified `disable_password_authentication` must be set to `false`.

NOTE: One of either `admin_password` or `admin_ssh_key` must be specified.

- `admin_ssh_key` - (Optional) One or more `admin_ssh_key` blocks as defined below.

NOTE: One of either `admin_password` or `admin_ssh_key` must be specified.

- `automatic_os_upgrade_policy` - (Optional) A `automatic_os_upgrade_policy` block as defined below. This is Required and can only be specified when `upgrade_mode` is set to `Automatic`.
- `boot_diagnostics` - (Optional) A `boot_diagnostics` block as defined below.
- `computer_name_prefix` - (Optional) The prefix which should be used for the name of the Virtual Machines in this Scale Set. If unspecified this defaults to the value for the `name` field.
- `custom_data` - (Optional) The Base64-Encoded Custom Data which should be used for this Virtual Machine Scale Set.

NOTE: When Custom Data has been configured, it's not possible to remove it without tainting the Virtual Machine Scale Set, due to a limitation of the Azure API.

- `data_disk` - (Optional) One or more `data_disk` blocks as defined below.
- `disable_password_authentication` - Should Password Authentication be disabled on this Virtual Machine Scale Set? Defaults to `true`.

In general we'd recommend using SSH Keys for authentication rather than Passwords - but there's tradeoffs to each - please see this thread for more information (<https://security.stackexchange.com/questions/69407/why-is-using-an-ssh-key-more-secure-than-using-passwords>).

NOTE: When an `admin_password` is specified `disable_password_authentication` must be set to `false`.

- `do_not_run_extensions_on_overprovisioned_machines` - (Optional) Should Virtual Machine Extensions be run on Overprovisioned Virtual Machines in the Scale Set? Defaults to `false`.
- `eviction_policy` - (Optional) The Policy which should be used Virtual Machines are Evicted from the Scale Set. Changing this forces a new resource to be created.

NOTE: This can only be configured when `priority` is set to `Low`.

- `health_probe_id` - (Optional) The ID of a Load Balancer Probe which should be used to determine the health of an instance. Changing this forces a new resource to be created. This is Required and can only be specified when `upgrade_mode` is set to `Automatic` or `Rolling`.
- `identity` - (Optional) A `identity` block as defined below.
- `max_bid_price` - (Optional) The maximum price you're willing to pay for a low-priority VM Scale Set, in US Dollars; which must be greater than the current low-priority price. If this bid price falls below the current low-priority price the Virtual Machines in the Scale Set will be evicted using the `eviction_policy`. Defaults to `-1`, which means that this VM Scale Set should not be evicted for price reasons.

NOTE: This can only be configured when `priority` is set to `Low`.

- `overprovision` - (Optional) Should Azure over-provision Virtual Machines in this Scale Set? This means that multiple Virtual Machines will be provisioned and Azure will keep the instances which become available first - which improves provisioning success rates and improves deployment time. You're not billed for these over-provisioned VM's and they don't count towards the Subscription Quota. Defaults to `false`.
- `priority` - (Optional) The Priority of this Virtual Machine Scale Set. Possible values are `Regular` and `Low`. Defaults to `Regular`. Changing this value forces a new resource.

NOTE: When `priority` is set to `Low` an `eviction_policy` must be specified.

- `provision_vm_agent` - (Optional) Should the Azure VM Agent be provisioned on each Virtual Machine in the Scale Set? Defaults to `true`. Changing this value forces a new resource to be created.

- `proximity_placement_group_id` - (Optional) The ID of the Proximity Placement Group in which the Virtual Machine Scale Set should be assigned to. Changing this forces a new resource to be created.
- `rolling_upgrade_policy` - (Optional) A `rolling_upgrade_policy` block as defined below. This is Required and can only be specified when `upgrade_mode` is set to `Automatic` or `Rolling`.
- `secret` - (Optional) One or more `secret` blocks as defined below.
- `single_placement_group` - (Optional) Should this Virtual Machine Scale Set be limited to a Single Placement Group, which means the number of instances will be capped at 100 Virtual Machines. Defaults to `true`.
- `source_image_id` - (Optional) The ID of an Image which each Virtual Machine in this Scale Set should be based on.

NOTE: One of either `source_image_id` or `source_image_reference` must be set.

- `source_image_reference` - (Optional) A `source_image_reference` block as defined below.

NOTE: One of either `source_image_id` or `source_image_reference` must be set.

- `tags` - (Optional) A mapping of tags which should be assigned to this Virtual Machine Scale Set.
- `terraform_should_roll_instances_when_required` - (Optional) Should Terraform automatically roll instances within the Virtual Machine Scale Set when required? This happens when the `data_disk`, `os_disk`, `sku`, `source_image_id`, or `source_image_reference` fields change. This field defaults to `true`.

NOTE: This field is specific to Terraform, when required Terraform will automatically roll the instances in a Scale Set one at a time.

- `upgrade_mode` - (Optional) Specifies how Upgrades (e.g. changing the Image/SKU) should be performed to Virtual Machine Instances. Possible values are `Automatic`, `Manual` and `Rolling`. Defaults to `Manual`.
- `zone_balance` - (Optional) Should the Virtual Machines in this Scale Set be strictly evenly distributed across Availability Zones? Defaults to `false`. Changing this forces a new resource to be created.

NOTE: This can only be set to `true` when one or more `zones` are configured.

- `zones` - (Optional) A list of Availability Zones in which the Virtual Machines in this Scale Set should be created in. Changing this forces a new resource to be created.

A `additional_capabilities` block supports the following:

- `ultra_ssd_enabled` - (Optional) Should the capacity to enable Data Disks of the `UltraSSD_LRS` storage account type be supported on this Virtual Machine Scale Set? Defaults to `false`. Changing this forces a new resource to be created.

A `admin_ssh_key` block supports the following:

- `public_key` - (Required) The Public Key which should be used for authentication, which needs to be at least 2048-bit and in `ssh-rsa` format.

- `username` - (Required) The Username for which this Public SSH Key should be configured.

NOTE: The Azure VM Agent only allows creating SSH Keys at the path `/home/{username}/.ssh/authorized_keys` - as such this public key will be added/appended to the authorized keys file.

A `automatic_os_upgrade_policy` block supports the following:

- `disable_automatic_rollback` - (Required) Should automatic rollbacks be disabled? Changing this forces a new resource to be created.
- `enable_automatic_os_upgrade` - (Required) Should OS Upgrades automatically be applied to Scale Set instances in a rolling fashion when a newer version of the OS Image becomes available? Changing this forces a new resource to be created.

A `boot_diagnostics` block supports the following:

- `storage_account_uri` - (Required) The Primary/Secondary Endpoint for the Azure Storage Account which should be used to store Boot Diagnostics, including Console Output and Screenshots from the Hypervisor.

A `certificate` block supports the following:

- `url` - (Required) The Secret URL of a Key Vault Certificate.

NOTE: This can be sourced from the `secret_id` field within the `azurerm_key_vault_certificate` Resource.

A `data_disk` block supports the following:

- `caching` - (Required) The type of Caching which should be used for this Data Disk. Possible values are `None` , `ReadOnly` and `ReadWrite` .
- `disk_size_gb` - (Required) The size of the Data Disk which should be created.
- `lun` - (Required) The Logical Unit Number of the Data Disk, which must be unique within the Virtual Machine.
- `storage_account_type` - (Required) The Type of Storage Account which should back this Data Disk. Possible values include `Standard_LRS` , `StandardSSD_LRS` , `Premium_LRS` and `UltraSSD_LRS` .

NOTE: `UltraSSD_LRS` is only supported when `ultra_ssd_enabled` within the `additional_capabilities` block is enabled.

- `write_accelerator_enabled` - (Optional) Should Write Accelerator be enabled for this Data Disk? Defaults to `false` .

NOTE: This requires that the `storage_account_type` is set to `Premium_LRS` and that `caching` is set to `None` .

A `diff_disk_settings` block supports the following:

`option` - (Required) Specifies the Ephemeral Disk Settings for the OS Disk. At this time the only possible value is `Local` . Changing this forces a new resource to be created.

A `identity` block supports the following:

- `type` - (Required) The type of Managed Identity which should be assigned to the Linux Virtual Machine Scale Set. Possible values are `SystemAssigned`, `UserAssigned` and `SystemAssigned, UserAssigned` .
- `identity_ids` - (Optional) A list of User Managed Identity ID's which should be assigned to the Linux Virtual Machine Scale Set.

NOTE: This is required when `type` is set to `UserAssigned` .

A `ip_configuration` block supports the following:

- `name` - (Required) The Name which should be used for this IP Configuration.
- `application_gateway_backend_address_pool_ids` - (Optional) A list of Backend Address Pools ID's from a Application Gateway which this Virtual Machine Scale Set should be connected to.
- `application_security_group_ids` - (Optional) A list of Application Security Group ID's which this Virtual Machine Scale Set should be connected to.
- `load_balancer_backend_address_pool_ids` - (Optional) A list of Backend Address Pools ID's from a Load Balancer which this Virtual Machine Scale Set should be connected to.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `load_balancer_inbound_nat_rules_ids` - (Optional) A list of NAT Rule ID's from a Load Balancer which this Virtual Machine Scale Set should be connected to.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `primary` - (Optional) Is this the Primary IP Configuration for this Network Interface? Defaults to `false` .

NOTE: One `ip_configuration` block must be marked as Primary for each Network Interface.

- `public_ip_address` - (Optional) A `public_ip_address` block as defined below.
- `subnet_id` - (Optional) The ID of the Subnet which this IP Configuration should be connected to.

`subnet_id` is required if `version` is set to `IPv4` .

- `version` - (Optional) The Internet Protocol Version which should be used for this IP Configuration. Possible values are `IPv4` and `IPv6` . Defaults to `IPv4` .
-

A `ip_tag` block supports the following:

- `tag` - The IP Tag associated with the Public IP, such as `SQL` or `Storage` .
- `type` - The Type of IP Tag, such as `FirstPartyUsage` .

A `network_interface` block supports the following:

- `name` - (Required) The Name which should be used for this Network Interface. Changing this forces a new resource to be created.
- `ip_configuration` - (Required) One or more `ip_configuration` blocks as defined above.
- `dns_servers` - (Optional) A list of IP Addresses of DNS Servers which should be assigned to the Network Interface.
- `enable_accelerated_networking` - (Optional) Does this Network Interface support Accelerated Networking? Defaults to `false` .
- `enable_ip_forwarding` - (Optional) Does this Network Interface support IP Forwarding? Defaults to `false` .
- `network_security_group_id` - (Optional) The ID of a Network Security Group which should be assigned to this Network Interface.
- `primary` - (Optional) Is this the Primary IP Configuration?

NOTE: If multiple `network_interface` blocks are specified, one must be set to `primary` .

A `os_disk` block supports the following:

- `caching` - (Required) The Type of Caching which should be used for the Internal OS Disk. Possible values are `None` , `ReadOnly` and `ReadWrite` .
- `storage_account_type` - (Required) The Type of Storage Account which should back this the Internal OS Disk. Possible values include `Standard_LRS` , `StandardSSD_LRS` and `Premium_LRS` .
- `diff_disk_settings` - (Optional) A `diff_disk_settings` block as defined above. Changing this forces a new resource to be created.
- `disk_size_gb` - (Optional) The Size of the Internal OS Disk in GB, if you wish to vary from the size used in the image this Virtual Machine Scale Set is sourced from.

NOTE: If specified this must be equal to or larger than the size of the Image the VM Scale Set is based on. When creating a larger disk than exists in the image you'll need to repartition the disk to use the remaining space.

- `write_accelerator_enabled` - (Optional) Should Write Accelerator be Enabled for this OS Disk? Defaults to `false` .

NOTE: This requires that the `storage_account_type` is set to `Premium_LRS` and that `caching` is set to `None` .

A `public_ip_address` block supports the following:

- `name` - (Required) The Name of the Public IP Address Configuration.

- `domain_name_label` - (Optional) The Prefix which should be used for the Domain Name Label for each Virtual Machine Instance. Azure concatenates the Domain Name Label and Virtual Machine Index to create a unique Domain Name Label for each Virtual Machine.
- `idle_timeout_in_minutes` - (Optional) The Idle Timeout in Minutes for the Public IP Address. Possible values are in the range 4 to 32 .
- `ip_tag` - (Optional) One or more `ip_tag` blocks as defined above.
- `public_ip_prefix_id` - (Optional) The ID of the Public IP Address Prefix from where Public IP Addresses should be allocated. Changing this forces a new resource to be created.

NOTE: This functionality is in Preview and must be opted into via `az feature register --namespace Microsoft.Network --name AllowBringYourOwnPublicIpAddress` and then `az provider register -n Microsoft.Network` .

A `rolling_upgrade_policy` block supports the following:

- `max_batch_instance_percent` - (Required) The maximum percent of total virtual machine instances that will be upgraded simultaneously by the rolling upgrade in one batch. As this is a maximum, unhealthy instances in previous or future batches can cause the percentage of instances in a batch to decrease to ensure higher reliability. Changing this forces a new resource to be created.
- `max_unhealthy_instance_percent` - (Required) The maximum percentage of the total virtual machine instances in the scale set that can be simultaneously unhealthy, either as a result of being upgraded, or by being found in an unhealthy state by the virtual machine health checks before the rolling upgrade aborts. This constraint will be checked prior to starting any batch. Changing this forces a new resource to be created.
- `max_unhealthy_upgraded_instance_percent` - (Required) The maximum percentage of upgraded virtual machine instances that can be found to be in an unhealthy state. This check will happen after each batch is upgraded. If this percentage is ever exceeded, the rolling update aborts. Changing this forces a new resource to be created.
- `pause_time_between_batches` - (Required) The wait time between completing the update for all virtual machines in one batch and starting the next batch. The time duration should be specified in ISO 8601 format. Changing this forces a new resource to be created.

A `secret` block supports the following:

- `certificate` - (Required) One or more `certificate` blocks as defined above.
- `key_vault_id` - (Required) The ID of the Key Vault from which all Secrets should be sourced.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the Linux Virtual Machine Scale Set.
- `identity` - An `identity` block as defined below.

- `unique_id` - The Unique ID for this Linux Virtual Machine Scale Set.
-

An `identity` block exports the following:

- `principal_id` - The ID of the System Managed Service Principal.

Timeouts

The `timeouts` block allows you to specify timeouts

(<https://www.terraform.io/docs/configuration/resources.html#timeouts>) for certain actions:

- `create` - (Defaults to 30 minutes) Used when creating the Linux Virtual Machine Scale Set.
- `update` - (Defaults to 60 minutes) Used when updating (and rolling the instances of) the Linux Virtual Machine Scale Set (e.g. when changing SKU).
- `delete` - (Defaults to 30 minutes) Used when deleting the Linux Virtual Machine Scale Set.

Import

Linux Virtual Machine Scale Sets can be imported using the `resource_id`, e.g.

```
terraform import azurerm_linux_virtual_machine_scale_set.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/Microsoft.Compute/virtualMachineScaleSets/scaleset1
```

azurerm_lb_backend_address_pool

Manages a Load Balancer Backend Address Pool.

NOTE: When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration Attached

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name                = "PublicIPForLB"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
}

resource "azurerm_lb" "test" {
  name                = "TestLoadBalancer"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name                = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_backend_address_pool" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id      = "${azurerm_lb.test.id}"
  name                 = "BackEndAddressPool"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Backend Address Pool.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which to create the Backend Address Pool.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Backend Address Pool.
- `backend_ip_configurations` - The Backend IP Configurations associated with this Backend Address Pool.
- `load_balancing_rules` - The Load Balancing Rules associated with this Backend Address Pool.

Import

Load Balancer Backend Address Pools can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb_backend_address_pool.test /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/backendAddressPools/pool1
```

azurerm_lb

Manages a Load Balancer Resource.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name            = "PublicIPForLB"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
}

resource "azurerm_lb" "test" {
  name            = "TestLoadBalancer"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name            = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Load Balancer.
- `resource_group_name` - (Required) The name of the Resource Group in which to create the Load Balancer.
- `location` - (Required) Specifies the supported Azure Region where the Load Balancer should be created.
- `frontend_ip_configuration` - (Optional) One or multiple `frontend_ip_configuration` blocks as documented below.
- `sku` - (Optional) The SKU of the Azure Load Balancer. Accepted values are `Basic` and `Standard`. Defaults to `Basic`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`frontend_ip_configuration` supports the following:

- `name` - (Required) Specifies the name of the frontend ip configuration.
- `subnet_id` - The ID of the Subnet which should be associated with the IP Configuration.

- `private_ip_address` - (Optional) Private IP Address to assign to the Load Balancer. The last one and first four IPs in any range are reserved and cannot be manually assigned.
- `private_ip_address_allocation` - (Optional) The allocation method for the Private IP Address used by this Load Balancer. Possible values as `Dynamic` and `Static`.
- `public_ip_address_id` - (Optional) The ID of a Public IP Address which should be associated with the Load Balancer.
- `public_ip_prefix_id` - (Optional) The ID of a Public IP Prefix which should be associated with the Load Balancer. Public IP Prefix can only be used with outbound rules.
- `zones` - (Optional) A list of Availability Zones which the Load Balancer's IP Addresses should be created in.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

Attributes Reference

The following attributes are exported:

- `id` - The Load Balancer ID.
- `private_ip_address` - The first private IP address assigned to the load balancer in `frontend_ip_configuration` blocks, if any.
- `private_ip_addresses` - The list of private IP address assigned to the load balancer in `frontend_ip_configuration` blocks, if any.

Import

Load Balancers can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1
```


azurerm_lb_nat_pool

Manages a Load Balancer NAT pool.

NOTE: This resource cannot be used with with virtual machines, instead use the `azurerm_lb_nat_rule` resource.

NOTE When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration Attached

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name            = "PublicIPForLB"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
}

resource "azurerm_lb" "test" {
  name            = "TestLoadBalancer"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name            = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_nat_pool" "test" {
  resource_group_name      = "${azurerm_resource_group.test.name}"
  loadbalancer_id          = "${azurerm_lb.test.id}"
  name                     = "SampleApplicationPool"
  protocol                 = "Tcp"
  frontend_port_start      = 80
  frontend_port_end        = 81
  backend_port              = 8080
  frontend_ip_configuration_name = "PublicIPAddress"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the NAT pool.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which to create the NAT pool.
- `frontend_ip_configuration_name` - (Required) The name of the frontend IP configuration exposing this rule.
- `protocol` - (Required) The transport protocol for the external endpoint. Possible values are `Udp` or `Tcp`.
- `frontend_port_start` - (Required) The first port number in the range of external ports that will be used to provide Inbound Nat to NICs associated with this Load Balancer. Possible values range between 1 and 65534, inclusive.
- `frontend_port_end` - (Required) The last port number in the range of external ports that will be used to provide Inbound Nat to NICs associated with this Load Balancer. Possible values range between 1 and 65534, inclusive.
- `backend_port` - (Required) The port used for the internal endpoint. Possible values range between 1 and 65535, inclusive.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer NAT pool.

Import

Load Balancer NAT Pools can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb_nat_pool.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/inboundNatPools/pool1
```

azurerm_lb_nat_rule

Manages a Load Balancer NAT Rule.

NOTE: This resource cannot be used with with virtual machine scale sets, instead use the `azurerm_lb_nat_pool` resource.

NOTE When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration Attached

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name                = "PublicIPForLB"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
}

resource "azurerm_lb" "test" {
  name                = "TestLoadBalancer"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name                = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_nat_rule" "test" {
  resource_group_name      = "${azurerm_resource_group.test.name}"
  loadbalancer_id          = "${azurerm_lb.test.id}"
  name                    = "RDPAccess"
  protocol                 = "Tcp"
  frontend_port            = 3389
  backend_port             = 3389
  frontend_ip_configuration_name = "PublicIPAddress"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the NAT Rule.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which to create the NAT Rule.
- `frontend_ip_configuration_name` - (Required) The name of the frontend IP configuration exposing this rule.
- `protocol` - (Required) The transport protocol for the external endpoint. Possible values are `Udp` , `Tcp` or `All` .
- `frontend_port` - (Required) The port for the external endpoint. Port numbers for each Rule must be unique within the Load Balancer. Possible values range between 1 and 65534, inclusive.
- `backend_port` - (Required) The port used for internal connections on the endpoint. Possible values range between 1 and 65535, inclusive.
- `enable_floating_ip` - (Optional) Enables the Floating IP Capacity, required to configure a SQL AlwaysOn Availability Group.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer NAT Rule.

Import

Load Balancer NAT Rules can be imported using the `resource id` , e.g.

```
terraform import azurerm_lb_nat_rule.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/inboundNatRules/rule1
```

azurerm_lb_outbound_rule

Manages a Load Balancer Outbound Rule.

NOTE When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration and a Backend Address Pool Attached.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name            = "PublicIPForLB"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
}

resource "azurerm_lb" "test" {
  name            = "TestLoadBalancer"
  location        = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name            = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_backend_address_pool" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id     = "${azurerm_lb.test.id}"
  name                = "be-%d"
}

resource "azurerm_lb_outbound_rule" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id     = "${azurerm_lb.test.id}"
  name                = "OutboundRule"
  protocol             = "Tcp"
  backend_address_pool_id = "${azurerm_lb_backend_address_pool.test.id}"

  frontend_ip_configuration {
    name = "PublicIPAddress"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Outbound Rule. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource. Changing this forces a new resource to be created.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which to create the Outbound Rule. Changing this forces a new resource to be created.
- `frontend_ip_configuration` - (Required) One or more `frontend_ip_configuration` blocks as defined below.
- `backend_address_pool_id` - (Required) The ID of the Backend Address Pool. Outbound traffic is randomly load balanced across IPs in the backend IPs.
- `protocol` - (Required) The transport protocol for the external endpoint. Possible values are `Udp`, `Tcp` or `All`.
- `enable_tcp_reset` - (Optional) Receive bidirectional TCP Reset on TCP flow idle timeout or unexpected connection termination. This element is only used when the protocol is set to TCP.
- `allocated_outbound_ports` - (Optional) The number of outbound ports to be used for NAT.
- `idle_timeout_in_minutes` - (Optional) The timeout for the TCP idle connection

A `frontend_ip_configuration` block supports the following:

- `name` - (Required) The name of the Frontend IP Configuration.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer Outbound Rule.

Import

Load Balancer Outbound Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb_outbound_rule.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/outboundRules/rule1
```

azurerm_lb_probe

Manages a LoadBalancer Probe Resource.

NOTE When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration Attached

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name                = "PublicIPForLB"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
}

resource "azurerm_lb" "test" {
  name                = "TestLoadBalancer"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name                = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_probe" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id      = "${azurerm_lb.test.id}"
  name                 = "ssh-running-probe"
  port                 = 22
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Probe.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource.
- `loadbalancer_id` - (Required) The ID of the LoadBalancer in which to create the NAT Rule.
- `protocol` - (Optional) Specifies the protocol of the end point. Possible values are `Http`, `Https` or `Tcp`. If `Tcp` is specified, a received ACK is required for the probe to be successful. If `Http` is specified, a 200 OK response from the

specified URI is required for the probe to be successful.

- `port` - (Required) Port on which the Probe queries the backend endpoint. Possible values range from 1 to 65535, inclusive.
- `request_path` - (Optional) The URI used for requesting health status from the backend endpoint. Required if protocol is set to Http. Otherwise, it is not allowed.
- `interval_in_seconds` - (Optional) The interval, in seconds between probes to the backend endpoint for health status. The default value is 15, the minimum value is 5.
- `number_of_probes` - (Optional) The number of failed probe attempts after which the backend endpoint is removed from rotation. The default value is 2. `NumberOfProbes` multiplied by `intervalInSeconds` value must be greater or equal to 10. Endpoints are returned to rotation when at least one probe is successful.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer Probe.

Import

Load Balancer Probes can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb_probe.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/probes/probe1
```


azurerm_lb_rule

Manages a Load Balancer Rule.

NOTE When using this resource, the Load Balancer needs to have a FrontEnd IP Configuration Attached

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "LoadBalancerRG"
  location  = "West US"
}

resource "azurerm_public_ip" "test" {
  name                = "PublicIPForLB"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
}

resource "azurerm_lb" "test" {
  name                = "TestLoadBalancer"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name                = "PublicIPAddress"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_rule" "test" {
  resource_group_name      = "${azurerm_resource_group.test.name}"
  loadbalancer_id          = "${azurerm_lb.test.id}"
  name                     = "LBRule"
  protocol                 = "Tcp"
  frontend_port            = 3389
  backend_port             = 3389
  frontend_ip_configuration_name = "PublicIPAddress"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the LB Rule.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource.
- `loadbalancer_id` - (Required) The ID of the Load Balancer in which to create the Rule.

- `frontend_ip_configuration_name` - (Required) The name of the frontend IP configuration to which the rule is associated.
- `protocol` - (Required) The transport protocol for the external endpoint. Possible values are `Tcp`, `Udp` or `All`.
- `frontend_port` - (Required) The port for the external endpoint. Port numbers for each Rule must be unique within the Load Balancer. Possible values range between 0 and 65534, inclusive.
- `backend_port` - (Required) The port used for internal connections on the endpoint. Possible values range between 0 and 65535, inclusive.
- `backend_address_pool_id` - (Optional) A reference to a Backend Address Pool over which this Load Balancing Rule operates.
- `probe_id` - (Optional) A reference to a Probe used by this Load Balancing Rule.
- `enable_floating_ip` - (Optional) Floating IP is pertinent to failover scenarios: a "floating" IP is reassigned to a secondary server in case the primary server fails. Floating IP is required for SQL AlwaysOn.
- `idle_timeout_in_minutes` - (Optional) Specifies the timeout for the `Tcp` idle connection. The value can be set between 4 and 30 minutes. The default value is 4 minutes. This element is only used when the protocol is set to `Tcp`.
- `load_distribution` - (Optional) Specifies the load balancing distribution type to be used by the Load Balancer. Possible values are: `Default` – The load balancer is configured to use a 5 tuple hash to map traffic to available servers. `SourceIP` – The load balancer is configured to use a 2 tuple hash to map traffic to available servers. `SourceIPProtocol` – The load balancer is configured to use a 3 tuple hash to map traffic to available servers. Also known as Session Persistence, where the options are called `None`, `Client IP` and `Client IP and Protocol` respectively.
- `disable_outbound_snat` - (Optional) Indicates whether outbound snat is disabled or enabled. Default false. ## Attributes Reference

The following attributes are exported:

- `id` - The ID of the Load Balancer Rule.

Import

Load Balancer Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_lb_rule.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/loadBalancingRules/rule1
```

azurerm_local_network_gateway

Manages a local network gateway connection over which specific connections can be configured.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "localNetworkGWTest"
  location  = "West US"
}

resource "azurerm_local_network_gateway" "home" {
  name                = "backHome"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  gateway_address     = "12.13.14.15"
  address_space       = ["10.0.0.0/16"]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the local network gateway. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the local network gateway.
- `location` - (Required) The location/region where the local network gateway is created. Changing this forces a new resource to be created.
- `gateway_address` - (Required) The IP address of the gateway to which to connect.
- `address_space` - (Required) The list of string CIDRs representing the address spaces the gateway exposes.
- `bgp_settings` - (Optional) A `bgp_settings` block as defined below containing the Local Network Gateway's BGP speaker settings.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`bgp_settings` supports the following:

- `asn` - (Required) The BGP speaker's ASN.
- `bgp_peering_address` - (Required) The BGP peering address and BGP identifier of this BGP speaker.
- `peer_weight` - (Optional) The weight added to routes learned from this BGP speaker.

Attributes Reference

The following attributes are exported:

- `id` - The local network gateway unique ID within Azure.

Import

Local Network Gateways can be imported using the `resource id`, e.g.

```
terraform import azurerm_local_network_gateway.lng1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/localNetworkGateways/lng1
```

azurerm_log_analytics_linked_service

Links a Log Analytics (formally Operational Insights) Workspace to another resource. The (currently) only linkable service is an Azure Automation Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourcegroup-01"
  location  = "West Europe"
}

resource "azurerm_automation_account" "test" {
  name                = "automation-01"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name = "Basic"
  }

  tags = {
    environment = "development"
  }
}

resource "azurerm_log_analytics_workspace" "test" {
  name                = "workspace-01"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "PerGB2018"
  retention_in_days   = 30
}

resource "azurerm_log_analytics_linked_service" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  workspace_name      = "${azurerm_log_analytics_workspace.test.name}"
  resource_id         = "${azurerm_automation_account.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which the Log Analytics Linked Service is created. Changing this forces a new resource to be created.
- `workspace_name` - (Required) Name of the Log Analytics Workspace that will contain the linkedServices resource. Changing this forces a new resource to be created.
- `linked_service_name` - (Optional) Name of the type of linkedServices resource to connect to the Log Analytics

Workspace specified in `workspace_name` . Currently it defaults to and only supports `automation` as a value. Changing this forces a new resource to be created.

- `resource_id` - (Optional) The ID of the Resource that will be linked to the workspace.
 - `linked_service_properties` - (Optional **Deprecated**) A `linked_service_properties` block as defined below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

`linked_service_properties` supports the following:

- `resource_id` - (Optional **Deprecated**) The resource id of the resource that will be linked to the workspace. This field has been deprecated in favour of the top-level `resource_id` field and will be removed in v2.0 of the AzureRM Provider.

Attributes Reference

The following attributes are exported:

- `id` - The Log Analytics Linked Service ID.
- `name` - The automatically generated name of the Linked Service. This cannot be specified. The format is always `<workspace_name>/<linked_service_name>` e.g. `workspace1/Automation`

Import

Log Analytics Workspaces can be imported using the `resource id` , e.g.

```
terraform import azurerm_log_analytics_linked_service.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.OperationalInsights/workspaces/workspace1/linkedservices/automation
```

azurerm_log_analytics_solution

Manages a Log Analytics (formally Operational Insights) Solution.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "k8s-log-analytics-test"
  location  = "westeurope"
}

resource "random_id" "workspace" {
  keepers = {
    # Generate a new id each time we switch to a new resource group
    group_name = "${azurerm_resource_group.test.name}"
  }

  byte_length = 8
}

resource "azurerm_log_analytics_workspace" "test" {
  name                = "k8s-workspace-${random_id.workspace.hex}"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                 = "PerGB2018"
}

resource "azurerm_log_analytics_solution" "test" {
  solution_name      = "ContainerInsights"
  location           = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  workspace_resource_id = "${azurerm_log_analytics_workspace.test.id}"
  workspace_name     = "${azurerm_log_analytics_workspace.test.name}"

  plan {
    publisher = "Microsoft"
    product   = "OMSGallery/ContainerInsights"
  }
}
```

Argument Reference

The following arguments are supported:

- `solution_name` - (Required) Specifies the name of the solution to be deployed. See [here](https://docs.microsoft.com/en-us/azure/log-analytics/log-analytics-add-solutions) for options (https://docs.microsoft.com/en-us/azure/log-analytics/log-analytics-add-solutions). Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which the Log Analytics solution is created. Changing this forces a new resource to be created. Note: The solution and it's related workspace can only exist in the same resource group.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `workspace_resource_id` - (Required) The full resource ID of the Log Analytics workspace with which the solution will be linked. Changing this forces a new resource to be created.
- `workspace_name` - (Required) The full name of the Log Analytics workspace with which the solution will be linked. Changing this forces a new resource to be created.
- `plan` - (Required) A `plan` block as documented below.

A `plan` block includes:

- `publisher` - (Required) The publisher of the solution. For example `Microsoft` . Changing this forces a new resource to be created.
- `product` - (Required) The product name of the solution. For example `OMSGallery/Containers` . Changing this forces a new resource to be created.
- `promotion_code` - (Optional) A promotion code to be used with the solution.

Import

Log Analytics Solutions can be imported using the `resource_id` , e.g.

```
terraform import azurerm_log_analytics_solution.solution1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.OperationsManagement/solutions/solution1
```


azurerm_log_analytics_workspace

Manages a Log Analytics (formally Operational Insights) Workspace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acctestRG-01"
  location = "East US"
}

resource "azurerm_log_analytics_workspace" "test" {
  name                = "acctest-01"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                  = "PerGB2018"
  retention_in_days    = 30
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Log Analytics Workspace. Workspace name should include 4-63 letters, digits or '-'. The '-' shouldn't be the first or the last symbol. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the Log Analytics workspace is created. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **sku** - (Required) Specifies the Sku of the Log Analytics Workspace. Possible values are `Free`, `PerNode`, `Premium`, `Standard`, `Standalone`, `Unlimited`, and `PerGB2018` (new Sku as of 2018-04-03).

NOTE: A new pricing model took effect on 2018-04-03, which requires the SKU `PerGB2018`. If you're provisioned resources before this date you have the option of remaining with the previous Pricing SKU and using the other SKU's defined above. More information about the Pricing SKU's is available at the following URI (<http://aka.ms/PricingTierWarning>).

- **retention_in_days** - (Optional) The workspace data retention in days. Possible values range between 30 and 730.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Log Analytics Workspace ID.
- `primary_shared_key` - The Primary shared key for the Log Analytics Workspace.
- `secondary_shared_key` - The Secondary shared key for the Log Analytics Workspace.
- `workspace_id` - The Workspace (or Customer) ID for the Log Analytics Workspace.
- `portal_url` - The Portal URL for the Log Analytics Workspace.

Import

Log Analytics Workspaces can be imported using the `resource id`, e.g.

```
terraform import azurerm_log_analytics_workspace.workspace1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.OperationalInsights/workspaces/workspace1
```

azurerm_log_analytics_workspace_linked_service

Links a Log Analytics (formally Operational Insights) Workspace to another resource. The (currently) only linkable service is an Azure Automation Account.

NOTE: This resource has been deprecated in favour of the `azurerm_log_analytics_linked_service` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourcegroup-01"
  location  = "West Europe"
}

resource "azurerm_automation_account" "test" {
  name                = "automation-01"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name = "Basic"
  }

  tags = {
    environment = "development"
  }
}

resource "azurerm_log_analytics_workspace" "test" {
  name                = "workspace-01"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku                = "PerGB2018"
  retention_in_days   = 30
}

resource "azurerm_log_analytics_workspace_linked_service" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  workspace_name      = "${azurerm_log_analytics_workspace.test.name}"
  resource_id         = "${azurerm_automation_account.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which the Log Analytics Linked Service is created. Changing this forces a new resource to be created.
- `workspace_name` - (Required) Name of the Log Analytics Workspace that will contain the linkedServices resource. Changing this forces a new resource to be created.
- `linked_service_name` - (Optional) Name of the type of linkedServices resource to connect to the Log Analytics Workspace specified in `workspace_name` . Currently it defaults to and only supports `automation` as a value. Changing this forces a new resource to be created.
- `resource_id` - (Optional) The ID of the Resource that will be linked to the workspace.
- `linked_service_properties` - (Optional **Deprecated**) A `linked_service_properties` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

`linked_service_properties` supports the following:

- `resource_id` - (Optional **Deprecated**) The resource id of the resource that will be linked to the workspace. This field has been deprecated in favour of the top-level `resource_id` field and will be removed in v2.0 of the AzureRM Provider.

Attributes Reference

The following attributes are exported:

- `id` - The Log Analytics Linked Service ID.
- `name` - The automatically generated name of the Linked Service. This cannot be specified. The format is always `<workspace_name>/<linked_service_name>` e.g. `workspace1/Automation`

Import

Log Analytics Workspaces can be imported using the `resource_id`, e.g.

```
terraform import azurerm_log_analytics_workspace_linked_service.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.OperationalInsights/workspaces/workspace1/linkedservices/automation
```

azurerm_logic_app_action_custom

Manages a Custom Action within a Logic App Workflow

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name            = "workflow1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_logic_app_action_custom" "test" {
  name          = "example-action"
  logic_app_id = "${azurerm_logic_app_workflow.test.id}"

  body = <<BODY
{
  "description": "A variable to configure the auto expiration age in days. Configured in negative number. Default is -30 (30 days old).",
  "inputs": {
    "variables": [
      {
        "name": "ExpirationAgeInDays",
        "type": "Integer",
        "value": -30
      }
    ]
  },
  "runAfter": {},
  "type": "InitializeVariable"
}
BODY
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the HTTP Action to be created within the Logic App Workflow. Changing this forces a new resource to be created.

NOTE: This name must be unique across all Actions within the Logic App Workflow.

- **logic_app_id** - (Required) Specifies the ID of the Logic App Workflow. Changing this forces a new resource to be

created.

- `body` - (Required) Specifies the JSON Blob defining the Body of this Custom Action.

NOTE: To make the Action more readable, you may wish to consider using HEREDOC syntax (as shown above) or the `local_file` resource (<https://www.terraform.io/docs/providers/local/d/file.html>) to load the schema from a file on disk.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Action within the Logic App Workflow.

Import

Logic App Custom Actions can be imported using the `resource id`, e.g.

```
terraform import azurerm_logic_app_action_custom.custom1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1/actions/custom1
```

NOTE: This ID is unique to Terraform and doesn't directly match to any other resource. To compose this ID, you can take the ID Logic App Workflow and append `/actions/{name of the action}`.

azurerm_logic_app_action_http

Manages an HTTP Action within a Logic App Workflow

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name            = "workflow1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_logic_app_action_http" "test" {
  name          = "webhook"
  logic_app_id = "${azurerm_logic_app_workflow.test.id}"
  method        = "GET"
  uri           = "http://example.com/some-webhook"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the HTTP Action to be created within the Logic App Workflow. Changing this forces a new resource to be created.

NOTE: This name must be unique across all Actions within the Logic App Workflow.

- `logic_app_id` - (Required) Specifies the ID of the Logic App Workflow. Changing this forces a new resource to be created.
- `method` - (Required) Specifies the HTTP Method which should be used for this HTTP Action. Possible values include DELETE, GET, PATCH, POST and PUT.
- `uri` - (Required) Specifies the URI which will be called when this HTTP Action is triggered.
- `body` - (Optional) Specifies the HTTP Body that should be sent to the `uri` when this HTTP Action is triggered.
- `headers` - (Optional) Specifies a Map of Key-Value Pairs that should be sent to the `uri` when this HTTP Action is triggered.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HTTP Action within the Logic App Workflow.

Import

Logic App HTTP Actions can be imported using the `resource id`, e.g.

```
terraform import azurerm_logic_app_action_http.webhook1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1/actions/webhook1
```

NOTE: This ID is unique to Terraform and doesn't directly match to any other resource. To compose this ID, you can take the ID Logic App Workflow and append `/actions/{name of the action}`.

azurerm_logic_app_trigger_custom

Manages a Custom Trigger within a Logic App Workflow

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name            = "workflow1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_logic_app_trigger_custom" "test" {
  name          = "example-trigger"
  logic_app_id = "${azurerm_logic_app_workflow.test.id}"

  body = <<BODY
{
  "recurrence": {
    "frequency": "Day",
    "interval": 1
  },
  "type": "Recurrence"
}
BODY
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the HTTP Trigger to be created within the Logic App Workflow. Changing this forces a new resource to be created.

NOTE: This name must be unique across all Triggers within the Logic App Workflow.

- `logic_app_id` - (Required) Specifies the ID of the Logic App Workflow. Changing this forces a new resource to be created.
- `body` - (Required) Specifies the JSON Blob defining the Body of this Custom Trigger.

NOTE: To make the Trigger more readable, you may wish to consider using HEREDOC syntax (as shown above) or the `local_file` resource (<https://www.terraform.io/docs/providers/local/d/file.html>) to load the schema from a file on disk.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Trigger within the Logic App Workflow.

Import

Logic App Custom Triggers can be imported using the `resource id`, e.g.

```
terraform import azurerm_logic_app_trigger_custom.custom1 /subscriptions/00000000-0000-0000-0000-000000000000/0000/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1/triggers/custom1
```

NOTE: This ID is unique to Terraform and doesn't directly match to any other resource. To compose this ID, you can take the ID Logic App Workflow and append `/triggers/{name of the trigger}`.

azurerm_logic_app_trigger_http_request

Manages a HTTP Request Trigger within a Logic App Workflow

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name            = "workflow1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_logic_app_trigger_http_request" "test" {
  name          = "some-http-trigger"
  logic_app_id = "${azurerm_logic_app_workflow.test.id}"

  schema = <<SCHEMA
{
  "type": "object",
  "properties": {
    "hello": {
      "type": "string"
    }
  }
}
SCHEMA
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the HTTP Request Trigger to be created within the Logic App Workflow. Changing this forces a new resource to be created.

NOTE: This name must be unique across all Triggers within the Logic App Workflow.

- `logic_app_id` - (Required) Specifies the ID of the Logic App Workflow. Changing this forces a new resource to be created.
- `schema` - (Required) A JSON Blob defining the Schema of the incoming request. This needs to be valid JSON.

NOTE: To make the Trigger more readable, you may wish to consider using HEREDOC syntax (as shown above) or the `local_file` resource (<https://www.terraform.io/docs/providers/local/d/file.html>) to load the schema from a file on

disk.

- `method` - (Optional) Specifies the HTTP Method which the request be using. Possible values include `DELETE`, `GET`, `PATCH`, `POST` or `PUT`.
- `relative_path` - (Optional) Specifies the Relative Path used for this Request.

NOTE: When `relative_path` is set a `method` must also be set.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the HTTP Request Trigger within the Logic App Workflow.

Import

Logic App HTTP Request Triggers can be imported using the `resource id`, e.g.

```
terraform import azurerm_logic_app_trigger_http_request.request1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1/triggers/request1
```

NOTE: This ID is unique to Terraform and doesn't directly match to any other resource. To compose this ID, you can take the ID Logic App Workflow and append `/triggers/{name of the trigger}`.

azurerm_logic_app_trigger_recurrence

Manages a Recurrence Trigger within a Logic App Workflow

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name            = "workflow1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_logic_app_trigger_recurrence" "test" {
  name          = "run-every-day"
  logic_app_id  = "${azurerm_logic_app_workflow.test.id}"
  frequency     = "Day"
  interval      = 1
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Recurrence Triggers to be created within the Logic App Workflow. Changing this forces a new resource to be created.

NOTE: This name must be unique across all Triggers within the Logic App Workflow.

- **logic_app_id** - (Required) Specifies the ID of the Logic App Workflow. Changing this forces a new resource to be created.
- **frequency** - (Required) Specifies the Frequency at which this Trigger should be run. Possible values include `Month`, `Week`, `Day`, `Hour`, `Minute` and `Second`.
- **interval** - (Required) Specifies interval used for the Frequency, for example a value of `4` for `interval` and `hour` for `frequency` would run the Trigger every 4 hours.

Attributes Reference

The following attributes are exported:

- **id** - The ID of the Recurrence Trigger within the Logic App Workflow.

Import

Logic App Recurrence Triggers can be imported using the `resource_id`, e.g.

```
terraform import azurerm_logic_app_trigger_recurrence.daily /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1/triggers/daily
```

NOTE: This ID is unique to Terraform and doesn't directly match to any other resource. To compose this ID, you can take the ID Logic App Workflow and append `/triggers/{name of the trigger}`.

azurerm_logic_app_workflow

Manages a Logic App Workflow.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "workflow-resources"
  location  = "East US"
}

resource "azurerm_logic_app_workflow" "test" {
  name                = "workflow1"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Logic App Workflow. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Logic App Workflow should be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Logic App Workflow exists. Changing this forces a new resource to be created.
- `workflow_schema` - (Optional) Specifies the Schema to use for this Logic App Workflow. Defaults to `https://schema.management.azure.com/providers/Microsoft.Logic/schemas/2016-06-01/workflowdefinition.json#`. Changing this forces a new resource to be created.
- `workflow_version` - (Optional) Specifies the version of the Schema used for this Logic App Workflow. Defaults to `1.0.0.0`. Changing this forces a new resource to be created.
- `parameters` - (Optional) A map of Key-Value pairs.

NOTE: Any parameters specified must exist in the Schema defined in `workflow_schema`.

- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Logic App Workflow ID.

- `access_endpoint` - The Access Endpoint for the Logic App Workflow

Import

Logic App Workflows can be imported using the `resource_id`, e.g.

```
terraform import azurerm_logic_app_workflow.workflow1 /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.Logic/workflows/workflow1
```


azurerm_managed_disk

Manages a managed disk.

Example Usage with Create Empty

```
resource "azurerm_resource_group" "test" {
  name     = "acctestRG"
  location = "West US 2"
}

resource "azurerm_managed_disk" "test" {
  name                = "acctestmd"
  location             = "West US 2"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_type = "Standard_LRS"
  create_option        = "Empty"
  disk_size_gb         = "1"

  tags = {
    environment = "staging"
  }
}
```

Example Usage with Create Copy

```

resource "azurerm_resource_group" "test" {
  name      = "acctestRG"
  location  = "West US 2"
}

resource "azurerm_managed_disk" "source" {
  name                = "acctestmd1"
  location             = "West US 2"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_type = "Standard_LRS"
  create_option        = "Empty"
  disk_size_gb         = "1"

  tags = {
    environment = "staging"
  }
}

resource "azurerm_managed_disk" "copy" {
  name                = "acctestmd2"
  location             = "West US 2"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_type = "Standard_LRS"
  create_option        = "Copy"
  source_resource_id   = "${azurerm_managed_disk.source.id}"
  disk_size_gb         = "1"

  tags = {
    environment = "staging"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the managed disk. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the managed disk.
- `location` - (Required) Specified the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `storage_account_type` - (Required) The type of storage to use for the managed disk. Allowable values are `Standard_LRS`, `Premium_LRS`, `StandardSSD_LRS` or `UltraSSD_LRS`.

Note: A `storage_account_type` of type `UltraSSD_LRS` and the arguments `disk_iops_read_write` and `disk_mbps_read_write` are currently in private preview and are not available to subscriptions that have not requested onboarding to Azure Ultra Disk Storage private preview. Azure Ultra Disk Storage is only available in East US 2, North Europe, and Southeast Asia regions. For more information see the [Azure Ultra Disk Storage product documentation](https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-enable-ultra-ssd) (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-enable-ultra-ssd>), [product blog](https://azure.microsoft.com/en-us/blog/announcing-the-general-availability-of-azure-ultra-disk-storage/) (<https://azure.microsoft.com/en-us/blog/announcing-the-general-availability-of-azure-ultra-disk-storage/>) and [FAQ](https://docs.microsoft.com/en-us/azure/virtual-machines/windows/faq-for-disks#ultra-disks) (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/faq-for-disks#ultra-disks>).

- `create_option` - (Required) The method to use when creating the managed disk. Possible values include:
 - `Import` - Import a VHD file in to the managed disk (VHD specified with `source_uri`).
 - `Empty` - Create an empty managed disk.
 - `Copy` - Copy an existing managed disk or snapshot (specified with `source_resource_id`).
 - `FromImage` - Copy a Platform Image (specified with `image_reference_id`)
 - `Restore` - Set by Azure Backup or Site Recovery on a restored disk (specified with `source_resource_id`).
- `source_uri` - (Optional) URI to a valid VHD file to be used when `create_option` is `Import` .
- `source_resource_id` - (Optional) ID of an existing managed disk to copy `create_option` is `Copy` or the recovery point to restore when `create_option` is `Restore`
- `image_reference_id` - (Optional) ID of an existing platform/marketplace disk image to copy when `create_option` is `FromImage` .
- `os_type` - (Optional) Specify a value when the source of an `Import` or `Copy` operation targets a source that contains an operating system. Valid values are `Linux` or `Windows`
- `disk_size_gb` - (Optional, Required for a new managed disk) Specifies the size of the managed disk to create in gigabytes. If `create_option` is `Copy` or `FromImage` , then the value must be equal to or greater than the source's size.
- `disk_iops_read_write` - (Optional) The number of IOPS allowed for this disk; only settable for UltraSSD disks. One operation can transfer between 4k and 256k bytes.
- `disk_mbps_read_write` - (Optional) The bandwidth allowed for this disk; only settable for UltraSSD disks. MBps means millions of bytes per second.
- `encryption_settings` - (Optional) an `encryption_settings` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zones` - (Optional) A collection containing the availability zone to allocate the Managed Disk in.

Note: Availability Zones are only supported in select regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

For more information on managed disks, such as sizing options and pricing, please check out the [azure documentation](https://docs.microsoft.com/en-us/azure/storage/storage-managed-disks-overview) (<https://docs.microsoft.com/en-us/azure/storage/storage-managed-disks-overview>).

`encryption_settings` supports:

- `enabled` - (Required) Is Encryption enabled on this Managed Disk? Changing this forces a new resource to be created.
- `disk_encryption_key` - (Optional) A `disk_encryption_key` block as defined below.
- `key_encryption_key` - (Optional) A `key_encryption_key` block as defined below.

`disk_encryption_key` supports:

- `secret_url` - (Required) The URL to the Key Vault Secret used as the Disk Encryption Key. This can be found as `id` on the `azurerm_key_vault_secret` resource.
- `source_vault_id` - (Required) The URL of the Key Vault. This can be found as `vault_uri` on the `azurerm_key_vault` resource.

`key_encryption_key` supports:

- `key_url` - (Required) The URL to the Key Vault Key used as the Key Encryption Key. This can be found as `id` on the `azurerm_key_vault_key` resource.
- `source_vault_id` - (Required) The URL of the Key Vault. This can be found as `vault_uri` on the `azurerm_key_vault` resource.

Attributes Reference

The following attributes are exported:

- `id` - The managed disk ID.

Import

Managed Disks can be imported using the `resource id`, e.g.

```
terraform import azurerm_managed_disk.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.compute/disks/manageddisk1
```

azurerm_management_group

Manages a Management Group.

Example Usage

```
data "azurerm_subscription" "current" {}

resource "azurerm_management_group" "example_parent" {
  display_name = "ParentGroup"

  subscription_ids = [
    "${data.azurerm_subscription.current.subscription_id}",
  ]
}

resource "azurerm_management_group" "example_child" {
  display_name           = "ChildGroup"
  parent_management_group_id = "${azurerm_management_group.example_parent.id}"

  subscription_ids = [
    "${data.azurerm_subscription.current.subscription_id}",
  ]

  # other subscription IDs can go here
}
```

Argument Reference

The following arguments are supported:

- `group_id` - (Optional) The UUID for this Management Group, which needs to be unique across your tenant - which will be generated if not provided. Changing this forces a new resource to be created.
- `display_name` - (Optional) A friendly name for this Management Group. If not specified, this'll be the same as the `group_id`.
- `parent_management_group_id` - (Optional) The ID of the Parent Management Group. Changing this forces a new resource to be created.
- `subscription_ids` - (Optional) A list of Subscription GUIDs which should be assigned to the Management Group.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Management Group.

Import

Management Groups can be imported using the `management_group_resource_id`, e.g.

```
terraform import azurerm_management_group.test /providers/Microsoft.Management/ManagementGroups/group1
```

azurerm_management_lock

Manages a Management Lock which is scoped to a Subscription, Resource Group or Resource.

Example Usage (Subscription Level Lock)

```
data "azurerm_subscription" "current" {}

resource "azurerm_management_lock" "subscription-level" {
  name      = "subscription-level"
  scope     = "${data.azurerm_subscription.current.id}"
  lock_level = "CanNotDelete"
  notes     = "Items can't be deleted in this subscription!"
}
```

Example Usage (Resource Group Level Lock)

```
resource "azurerm_resource_group" "test" {
  name      = "locked-resource-group"
  location = "West Europe"
}

resource "azurerm_management_lock" "resource-group-level" {
  name      = "resource-group-level"
  scope     = "${azurerm_resource_group.test.id}"
  lock_level = "ReadOnly"
  notes     = "This Resource Group is Read-Only"
}
```

Example Usage (Resource Level Lock)

```
resource "azurerm_resource_group" "test" {
  name      = "locked-resource-group"
  location  = "West Europe"
}

resource "azurerm_public_ip" "test" {
  name                = "locked-publicip"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
  idle_timeout_in_minutes = 30
}

resource "azurerm_management_lock" "public-ip" {
  name      = "resource-ip"
  scope     = "${azurerm_public_ip.test.id}"
  lock_level = "CanNotDelete"
  notes     = "Locked because it's needed by a third-party"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Management Lock. Changing this forces a new resource to be created.
- `scope` - (Required) Specifies the scope at which the Management Lock should be created. Changing this forces a new resource to be created.
- `lock_level` - (Required) Specifies the Level to be used for this Lock. Possible values are `CanNotDelete` and `ReadOnly`. Changing this forces a new resource to be created.

Note: `CanNotDelete` means authorized users are able to read and modify the resources, but not delete. `ReadOnly` means authorized users can only read from a resource, but they can't modify or delete it.

- `notes` - (Optional) Specifies some notes about the lock. Maximum of 512 characters. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Management Lock

Import

Management Locks can be imported using the `resource id`, e.g.


```
terraform import azurerm_management_lock.lock1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Authorization/locks/lock1
```

azurerm_maps_account

Manages an Azure Maps Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_maps_account" "test" {
  name                  = "example-maps-account"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku_name              = "s1"

  tags = {
    environment = "Test"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Azure Maps Account. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Azure Maps Account should exist. Changing this forces a new resource to be created.
- `sku_name` - (Required) The sku of the Azure Maps Account. Possible values are `s0` and `s1`.
- `tags` - (Optional) A mapping of tags to assign to the Azure Maps Account.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the Azure Maps Account.
- `primary_access_key` - The primary key used to authenticate and authorize access to the Maps REST APIs.
- `secondary_access_key` - The secondary key used to authenticate and authorize access to the Maps REST APIs.
- `x_ms_client_id` - A unique identifier for the Maps Account.

Import

A Maps Account can be imported using the `resource_id`, e.g.

```
terraform import azurerm_maps_account.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Maps/accounts/my-maps-account
```

azurerm_mariadb_configuration

Sets a MariaDB Configuration value on a MariaDB Server.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mariadb_server" "example" {
  name                = "mariadb-server-1"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name      = "B_Gen5_2"
    capacity = 2
    tier       = "Basic"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb            = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "mariadbadmin"
  administrator_login_password = "HqSh1CoR3!"
  version                     = "10.2"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_mariadb_configuration" "example" {
  name                = "interactive_timeout"
  resource_group_name = "${azurerm_resource_group.example.name}"
  server_name         = "${azurerm_mariadb_server.example.name}"
  value               = "600"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the MariaDB Configuration, which needs to be a valid MariaDB configuration name (<https://mariadb.com/kb/en/library/server-system-variables/>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the MariaDB Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the MariaDB Server exists. Changing this forces a new resource to be created.
- `value` - (Required) Specifies the value of the MariaDB Configuration. See the MariaDB documentation for valid values.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MariaDB Configuration.

Import

MariaDB Configurations can be imported using the `resource id`, e.g.

```
terraform import azurerm_mariadb_configuration.interactive_timeout /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMariaDB/servers/server1/configurations/interactive_timeout
```

azurerm_mariadb_database

Manages a MariaDB Database within a MariaDB Server

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-mariadb-database-RG"
  location  = "westeurope"
}

resource "azurerm_mariadb_server" "example" {
  name                = "mariadb-svr"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name      = "B_Gen5_2"
    capacity  = 2
    tier       = "Basic"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb          = 51200
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "acctestun"
  administrator_login_password = "HqSh1CoR3!"
  version                     = "10.2"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_mariadb_database" "example" {
  name                = "mariadb_database"
  resource_group_name = "${azurerm_resource_group.example.name}"
  server_name         = "${azurerm_mariadb_server.example.name}"
  charset             = "utf8"
  collation            = "utf8_general_ci"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the MariaDB Database, which needs to be a valid MariaDB identifier (<https://mariadb.com/kb/en/library/identifier-names/>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the MariaDB Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the MariaDB Server exists. Changing this forces a new resource to be created.
- `charset` - (Required) Specifies the Charset for the MariaDB Database, which needs to be a valid MariaDB Charset (<https://mariadb.com/kb/en/library/setting-character-sets-and-collations>). Changing this forces a new resource to be created.
- `collation` - (Required) Specifies the Collation for the MariaDB Database, which needs to be a valid MariaDB Collation (<https://mariadb.com/kb/en/library/setting-character-sets-and-collations>). Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MariaDB Database.

Import

MariaDB Database's can be imported using the `resource id`, e.g.

```
terraform import azurerm_mariadb_database.database1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMariaDB/servers/server1/databases/database1
```

azurerm_mariadb_firewall_rule

Manages a Firewall Rule for a MariaDB Server

Example Usage (Single IP Address)

```
resource "azurerm_mariadb_firewall_rule" "test" {
  name                = "test-rule"
  resource_group_name = "test-rg"
  server_name         = "test-server"
  start_ip_address    = "40.112.8.12"
  end_ip_address      = "40.112.8.12"
}
```

Example Usage (IP Range)

```
resource "azurerm_mariadb_firewall_rule" "test" {
  name                = "test-rule"
  resource_group_name = "test-rg"
  server_name         = "test-server"
  start_ip_address    = "40.112.0.0"
  end_ip_address      = "40.112.255.255"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the MariaDB Firewall Rule. Changing this forces a new resource to be created.
- **server_name** - (Required) Specifies the name of the MariaDB Server. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the MariaDB Server exists. Changing this forces a new resource to be created.
- **start_ip_address** - (Required) Specifies the Start IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.
- **end_ip_address** - (Required) Specifies the End IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MariaDB Firewall Rule.

Import

MariaDB Firewall rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_mariadb_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.DBforMariaDB/servers/server1/firewallRules/rule1
```

azurerm_mariadb_server

Manages a MariaDB Server.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "api-rg-pro"
  location = "West Europe"
}

resource "azurerm_mariadb_server" "test" {
  name                = "mariadb-server-1"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name       = "B_Gen5_2"
    capacity   = 2
    tier        = "Basic"
    family     = "Gen5"
  }

  storage_profile {
    storage_mb           = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
    storage_autogrow     = "Disabled"
  }

  administrator_login          = "mariadbadmin"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "10.2"
  ssl_enforcement              = "Enabled"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the MariaDB Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the MariaDB Server. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `storage_profile` - (Required) A `storage_profile` block as defined below.

- `administrator_login` - (Required) The Administrator Login for the MariaDB Server. Changing this forces a new resource to be created.
- `administrator_login_password` - (Required) The Password associated with the `administrator_login` for the MariaDB Server.
- `version` - (Required) Specifies the version of MariaDB to use. Possible values are `10.2` and `10.3`. Changing this forces a new resource to be created.
- `ssl_enforcement` - (Required) Specifies if SSL should be enforced on connections. Possible values are `Enabled` and `Disabled`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block supports the following:

- `name` - (Required) Specifies the SKU Name for this MariaDB Server. The name of the SKU, follows the `tier + family + cores` pattern (e.g. `B_Gen5_1`, `GP_Gen5_8`). For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/mariadb/servers/create#sku>).
- `capacity` - (Required) The scale up/out capacity, representing server's compute units.
- `tier` - (Required) The tier of the particular SKU. Possible values are `Basic`, `GeneralPurpose`, and `MemoryOptimized`. For more information see the product documentation (<https://docs.microsoft.com/en-us/azure/mariadb/concepts-pricing-tiers>).
- `family` - (Required) The `family` of the hardware (e.g. `Gen5`), before selecting your `family` check the product documentation (<https://docs.microsoft.com/en-us/azure/mariadb/concepts-pricing-tiers#compute-generations-vccores-and-memory>) for availability in your region.

A `storage_profile` block supports the following:

- `storage_mb` - (Required) Max storage allowed for a server. Possible values are between `5120 MB` (5GB) and `1024000 MB` (1TB) for the Basic SKU and between `5120 MB` (5GB) and `4096000 MB` (4TB) for General Purpose/Memory Optimized SKUs. For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/mariadb/servers/create#storageprofile>).
- `backup_retention_days` - (Optional) Backup retention days for the server, supported values are between `7` and `35` days.
- `geo_redundant_backup` - (Optional) Enable Geo-redundant or not for server backup. Valid values for this property are `Enabled` or `Disabled`.
- `auto_grow` - (Optional) Defines whether autogrow is enabled or disabled for the storage. Valid values are `Enabled` or `Disabled`.

NOTE: Geo Redundant Backups cannot be configured when using the `Basic` tier.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MariaDB Server.
- `fqdn` - The FQDN of the MariaDB Server.

Import

MariaDB Server's can be imported using the `resource id`, e.g.

```
terraform import azurerm_mariadb_server.server1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMariaDB/servers/server1
```

azurerm_mariadb_virtual_network_rule

Manages a MariaDB Virtual Network Rule.

NOTE: MariaDB Virtual Network Rules can only be used with SKU Tiers of `GeneralPurpose` or `MemoryOptimized` (<https://docs.microsoft.com/en-us/azure/mariadb/concepts-data-access-security-vnet>)

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-vnet"
  address_space   = ["10.7.29.0/29"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "internal" {
  name                 = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.7.29.0/29"
  service_endpoints    = ["Microsoft.Sql"]
}

resource "azurerm_mariadb_server" "test" {
  name                        = "mariadb-server-1"
  location                  = "${azurerm_resource_group.test.location}"
  resource_group_name       = "${azurerm_resource_group.test.name}"
  administrator_login       = "mariadbadminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                   = "5.7"
  ssl_enforcement           = "Enabled"

  sku {
    name      = "GP_Gen5_2"
    capacity = 2
    tier      = "GeneralPurpose"
    family   = "Gen5"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }
}

resource "azurerm_mariadb_virtual_network_rule" "test" {
  name            = "mariadb-vnet-rule"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name      = "${azurerm_mariadb_server.test.name}"
  subnet_id        = "${azurerm_subnet.internal.id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the MariaDB Virtual Network Rule. Cannot be empty and must only contain

alphanumeric characters and hyphens. Cannot start with a number, and cannot start or end with a hyphen. Changing this forces a new resource to be created.

NOTE: `name` must be between 1-128 characters long and must satisfy all of the requirements below: 1. Contains only alphanumeric and hyphen characters 2. Cannot start with a number or hyphen 3. Cannot end with a hyphen

- `resource_group_name` - (Required) The name of the resource group where the MariaDB server resides. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server to which this MariaDB virtual network rule will be applied to. Changing this forces a new resource to be created.
- `subnet_id` - (Required) The ID of the subnet that the MariaDB server will be connected to.

NOTE: Due to a bug in the Azure API (<https://github.com/Azure/azure-rest-api-specs/issues/3719>) this resource currently doesn't expose the `ignore_missing_vnet_service_endpoint` field and defaults this to `false`. Terraform will check during the provisioning of the Virtual Network Rule that the Subnet contains the Service Rule to verify that the Virtual Network Rule can be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MariaDB Virtual Network Rule.

Import

MariaDB Virtual Network Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_mariadb_virtual_network_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/MariaDB/servers/myserver/virtualNetworkRules/vnetrulename
```

azurerm_marketplace_agreement

Allows accepting the Legal Terms for a Marketplace Image.

Example Usage

```
resource "azurerm_marketplace_agreement" "barracuda" {  
  publisher = "barracudanetworks"  
  offer     = "waf"  
  plan      = "hourly"  
}
```

Argument Reference

The following arguments are supported:

- `offer` - (Required) The Offer of the Marketplace Image. Changing this forces a new resource to be created.
- `plan` - (Required) The Plan of the Marketplace Image. Changing this forces a new resource to be created.
- `publisher` - (Required) The Publisher of the Marketplace Image. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Marketplace Agreement.

Import

Marketplace Agreement can be imported using the `resource id`, e.g.

```
terraform import azurerm_marketplace_agreement.example /subscriptions/000000000-0000-0000-0000-000000000000  
0/providers/Microsoft.MarketplaceOrdering/offerTypes/virtualmachine/publishers/publisher1/offers/offer1/p  
lans/plan1/agreements/current
```


azurerm_media_services_account

Manages a Media Services Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "media-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "examplestoracc"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "GRS"
}

resource "azurerm_media_services_account" "test" {
  name                = "examplemediaacc"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  storage_account {
    id          = "${azurerm_storage_account.test.id}"
    is_primary = true
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Media Services Account. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Media Services Account. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **storage_account** - (Required) One or more `storage_account` blocks as defined below.

A `storage_account` block supports the following:

- **id** - (Required) Specifies the ID of the Storage Account that will be associated with the Media Services instance.

- `is_primary` - (Required) Specifies whether the storage account should be the primary account or not. Defaults to `false`.

NOTE: Whilst multiple `storage_account` blocks can be specified - one of them must be set to the primary

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Media Services Account.

Import

Media Services Accounts can be imported using the `resource id`, e.g.

```
terraform import azurerm_media_services_account.account /subscriptions/00000000-0000-0000-0000-000000000000/00/resourceGroups/mygroup1/providers/Microsoft.Media/mediaservices/account1
```

azurerm_metric_alertrule

Manages a metric-based alert rule (<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitor-quick-resource-metric-alert-portal>) in Azure Monitor.

NOTE: This resource has been deprecated (<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/monitoring-classic-retirement>) in favour of the `azurerm_monitor_metric_alert` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage (CPU Percentage of a virtual machine)

```
resource "azurerm_metric_alertrule" "test" {
  name                = "${azurerm_virtual_machine.test.name}-cpu"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  description = "An alert rule to watch the metric Percentage CPU"

  enabled = true

  resource_id = "${azurerm_virtual_machine.test.id}"
  metric_name = "Percentage CPU"
  operator    = "GreaterThan"
  threshold   = 75
  aggregation = "Average"
  period      = "PT5M"

  email_action {
    send_to_service_owners = false

    custom_emails = [
      "some.user@example.com",
    ]
  }

  webhook_action {
    service_uri = "https://example.com/some-url"

    properties = {
      severity          = "incredible"
      acceptance_test = "true"
    }
  }
}
```

Example Usage (Storage usage of a SQL Database)

```

resource "azurerm_metric_alertrule" "test" {
  name                = "${azurerm_sql_database.test.name}-storage"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  description = "An alert rule to watch the metric Storage"

  enabled = true

  resource_id = "${azurerm_sql_database.test.id}"
  metric_name = "storage"
  operator    = "GreaterThan"
  threshold   = 1073741824
  aggregation = "Maximum"
  period      = "PT10M"

  email_action {
    send_to_service_owners = false

    custom_emails = [
      "some.user@example.com",
    ]
  }

  webhook_action {
    service_uri = "https://example.com/some-url"

    properties = {
      severity      = "incredible"
      acceptance_test = "true"
    }
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the alert rule. Changing this forces a new resource to be created.
 - **resource_group_name** - (Required) The name of the resource group in which to create the alert rule. Changing this forces a new resource to be created.
 - **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
 - **description** - (Optional) A verbose description of the alert rule that will be included in the alert email.
 - **enabled** - (Optional) If `true`, the alert rule is enabled. Defaults to `true`.
-
- **resource_id** - (Required) The ID of the resource monitored by the alert rule.
 - **metric_name** - (Required) The metric that defines what the rule monitors.

For a comprehensive reference of supported `metric_name` values for types of `resource` refer to Supported metrics with Azure Monitor (<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-supported-metrics>) in the Azure documentation. In the referred table, the column "Metric" corresponds to supported values for `metric_name`.

- `operator` - (Required) The operator used to compare the metric data and the threshold. Possible values are `GreaterThan`, `GreaterThanOrEqual`, `LessThan`, and `LessThanOrEqual`.
- `threshold` - (Required) The threshold value that activates the alert.
- `period` - (Required) The period of time formatted in ISO 8601 duration format (https://en.wikipedia.org/wiki/ISO_8601#Durations) that is used to monitor the alert activity based on the threshold. The period must be between 5 minutes and 1 day.
- `aggregation` - (Required) Defines how the metric data is combined over time. Possible values are `Average`, `Minimum`, `Maximum`, `Total`, and `Last`.

- `email_action` - (Optional) A `email_action` block as defined below.
- `webhook_action` - (Optional) A `webhook_action` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource. Changing this forces a new resource to be created.

`email_action` supports the following:

- `send_to_service_owners` - (Optional) If `true`, the administrators (service and co-administrators) of the subscription are notified when the alert is triggered. Defaults to `false`.
- `custom_emails` - (Optional) A list of email addresses to be notified when the alert is triggered.

`webhook_action` supports the following:

- `service_uri` - (Required) The service uri of the webhook to POST the notification when the alert is triggered.
- `properties` - (Optional) A dictionary of custom properties to include with the webhook POST operation payload.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the alert rule.

Import

Metric Alert Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_metric_alertrule.alertrule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/alertrules/alertrule1
```


azurerm_monitor_action_group

Manages an Action Group within Azure Monitor.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "monitoring-resources"
  location  = "West US"
}

resource "azurerm_monitor_action_group" "test" {
  name                        = "CriticalAlertsAction"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  short_name                 = "p0action"

  email_receiver {
    name          = "sendtoadmin"
    email_address = "admin@contoso.com"
  }

  email_receiver {
    name          = "sendtodevops"
    email_address = "devops@contoso.com"
  }

  sms_receiver {
    name          = "oncallmsg"
    country_code = "1"
    phone_number = "1231231234"
  }

  webhook_receiver {
    name                = "callmyapiaswell"
    service_uri         = "http://example.com/alert"
    use_common_alert_schema = true
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Action Group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Action Group instance.
- `short_name` - (Required) The short name of the action group. This will be used in SMS messages.
- `enabled` - (Optional) Whether this action group is enabled. If an action group is not enabled, then none of its receivers will receive communications. Defaults to `true`.

- `email_receiver` - (Optional) One or more `email_receiver` blocks as defined below.
 - `sms_receiver` - (Optional) One or more `sms_receiver` blocks as defined below.
 - `webhook_receiver` - (Optional) One or more `webhook_receiver` blocks as defined below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

`email_receiver` supports the following:

- `name` - (Required) The name of the email receiver. Names must be unique (case-insensitive) across all receivers within an action group.
 - `email_address` - (Required) The email address of this receiver.
-

`sms_receiver` supports the following:

- `name` - (Required) The name of the SMS receiver. Names must be unique (case-insensitive) across all receivers within an action group.
 - `country_code` - (Required) The country code of the SMS receiver.
 - `phone_number` - (Required) The phone number of the SMS receiver.
-

`webhook_receiver` supports the following:

- `name` - (Required) The name of the webhook receiver. Names must be unique (case-insensitive) across all receivers within an action group.
- `service_uri` - (Required) The URI where webhooks should be sent.
- `use_common_alert_schema` - (Optional) Enables or disables the common alert schema.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Action Group.

Import

Action Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_action_group.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Insights/actionGroups/myagname
```


azurerm_monitor_activity_log_alert

Manages an Activity Log Alert within Azure Monitor.

Example Usage

```
resource "azurerm_resource_group" "main" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_monitor_action_group" "main" {
  name                = "example-actiongroup"
  resource_group_name = "${azurerm_resource_group.main.name}"
  short_name          = "p0action"

  webhook_receiver {
    name      = "callmyapi"
    service_uri = "http://example.com/alert"
  }
}

resource "azurerm_storage_account" "to_monitor" {
  name                = "examplesa"
  resource_group_name = "${azurerm_resource_group.main.name}"
  location             = "${azurerm_resource_group.main.location}"
  account_tier         = "Standard"
  account_replication_type = "GRS"
}

resource "azurerm_monitor_activity_log_alert" "main" {
  name                = "example-activitylogalert"
  resource_group_name = "${azurerm_resource_group.main.name}"
  scopes              = ["${azurerm_resource_group.main.id}"]
  description          = "This alert will monitor a specific storage account updates."

  criteria {
    resource_id      = "${azurerm_storage_account.to_monitor.id}"
    operation_name    = "Microsoft.Storage/storageAccounts/write"
    category          = "Recommendation"
  }

  action {
    action_group_id = "${azurerm_monitor_action_group.main.id}"

    webhook_properties = {
      from = "terraform"
    }
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the activity log alert. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the activity log alert instance.
- `scopes` - (Required) The Scope at which the Activity Log should be applied, for example a the Resource ID of a Subscription or a Resource (such as a Storage Account).
- `criteria` - (Required) A `criteria` block as defined below.
- `action` - (Optional) One or more `action` blocks as defined below.
- `enabled` - (Optional) Should this Activity Log Alert be enabled? Defaults to `true`.
- `description` - (Optional) The description of this activity log alert.
- `tags` - (Optional) A mapping of tags to assign to the resource.

An `action` block supports the following:

- `action_group_id` - (Required) The ID of the Action Group can be sourced from the `azurerm_monitor_action_group` resource ([/docs/providers/azurerm/r/monitor_action_group.html](/docs/providers/azurerm/r/monitor_action_group.html)).
- `webhook_properties` - (Optional) The map of custom string properties to include with the post operation. These data are appended to the webhook payload.

A `criteria` block supports the following:

- `category` - (Required) The category of the operation. Possible values are `Administrative`, `Autoscale`, `Policy`, `Recommendation`, `Security` and `ServiceHealth`.
- `operation_name` - (Optional) The Resource Manager Role-Based Access Control operation name. Supported operation should be of the form: `<resourceProvider>/<resourceType>/<operation>`.
- `resource_provider` - (Optional) The name of the resource provider monitored by the activity log alert.
- `resource_type` - (Optional) The resource type monitored by the activity log alert.
- `resource_group` - (Optional) The name of resource group monitored by the activity log alert.
- `resource_id` - (Optional) The specific resource monitored by the activity log alert. It should be within one of the `scopes`.
- `caller` - (Optional) The email address or Azure Active Directory identifier of the user who performed the operation.
- `level` - (Optional) The severity level of the event. Possible values are `Verbose`, `Informational`, `Warning`, `Error`, and `Critical`.
- `status` - (Optional) The status of the event. For example, `Started`, `Failed`, or `Succeeded`.
- `sub_status` - (Optional) The sub status of the event.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the activity log alert.

Import

Activity log alerts can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_activity_log_alert.test /subscriptions/00000000-0000-0000-0000-000000000000/000/resourceGroups/group1/providers/microsoft.insights/activityLogAlerts/myalertname
```

azurerm_monitor_autoscale_setting

Manages a AutoScale Setting which can be applied to Virtual Machine Scale Sets, App Services and other scalable resources.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "autoscalingTest"
  location  = "West US"
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_monitor_autoscale_setting" "test" {
  name                        = "myAutoscaleSetting"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  target_resource_id         = "${azurerm_virtual_machine_scale_set.test.id}"

  profile {
    name = "defaultProfile"

    capacity {
      default = 1
      minimum = 1
      maximum = 10
    }

    rule {
      metric_trigger {
        metric_name      = "Percentage CPU"
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
        time_grain       = "PT1M"
        statistic         = "Average"
        time_window       = "PT5M"
        time_aggregation  = "Average"
        operator          = "GreaterThan"
        threshold         = 75
      }

      scale_action {
        direction = "Increase"
        type      = "ChangeCount"
        value     = "1"
        cooldown  = "PT1M"
      }
    }
  }

  rule {
    metric_trigger {
      metric_name      = "Percentage CPU"
      metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
      time_grain       = "PT1M"
    }
  }
}
```

```

    ~
    statistic          = "Average"
    time_window        = "PT5M"
    time_aggregation   = "Average"
    operator           = "LessThan"
    threshold          = 25
  }

  scale_action {
    direction = "Decrease"
    type      = "ChangeCount"
    value     = "1"
    cooldown  = "PT1M"
  }
}

notification {
  email {
    send_to_subscription_administrator = true
    send_to_subscription_co_administrator = true
    custom_emails                     = ["admin@contoso.com"]
  }
}
}

```

Example Usage (repeating on weekends)

```

resource "azurerm_resource_group" "test" {
  name     = "autoscalingTest"
  location = "West US"
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_monitor_autoscale_setting" "test" {
  name                = "myAutoscaleSetting"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"

  profile {
    name = "Weekends"

    capacity {
      default = 1
      minimum = 1
      maximum = 10
    }

    rule {
      metric_trigger {
        metric_name      = "Percentage CPU"
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
      }
    }
  }
}

```

```

        time_grain          = "PT1M"
        statistic           = "Average"
        time_window         = "PT5M"
        time_aggregation    = "Average"
        operator            = "GreaterThan"
        threshold           = 90
    }

    scale_action {
        direction = "Increase"
        type      = "ChangeCount"
        value     = "2"
        cooldown  = "PT1M"
    }
}

rule {
    metric_trigger {
        metric_name          = "Percentage CPU"
        metric_resource_id   = "${azurerem_virtual_machine_scale_set.test.id}"
        time_grain           = "PT1M"
        statistic            = "Average"
        time_window          = "PT5M"
        time_aggregation     = "Average"
        operator             = "LessThan"
        threshold            = 10
    }

    scale_action {
        direction = "Decrease"
        type      = "ChangeCount"
        value     = "2"
        cooldown  = "PT1M"
    }
}

recurrence {
    frequency = "Week"
    timezone  = "Pacific Standard Time"
    days      = ["Saturday", "Sunday"]
    hours     = [12]
    minutes   = [0]
}

notification {
    email {
        send_to_subscription_administrator = true
        send_to_subscription_co_administrator = true
        custom_emails                      = ["admin@contoso.com"]
    }
}
}

```

Example Usage (for fixed dates)

```

resource "azurerm_resource_group" "test" {
  name      = "autoscalingTest"
  location  = "West US"
}

resource "azurerm_virtual_machine_scale_set" "test" {
  # ...
}

resource "azurerm_monitor_autoscale_setting" "test" {
  name                = "myAutoscaleSetting"
  enabled              = true
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  target_resource_id  = "${azurerm_virtual_machine_scale_set.test.id}"

  profile {
    name = "forJuly"

    capacity {
      default = 1
      minimum = 1
      maximum = 10
    }

    rule {
      metric_trigger {
        metric_name      = "Percentage CPU"
        metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
        time_grain       = "PT1M"
        statistic         = "Average"
        time_window       = "PT5M"
        time_aggregation = "Average"
        operator          = "GreaterThan"
        threshold         = 90
      }

      scale_action {
        direction = "Increase"
        type      = "ChangeCount"
        value     = "2"
        cooldown  = "PT1M"
      }
    }
  }

  rule {
    metric_trigger {
      metric_name      = "Percentage CPU"
      metric_resource_id = "${azurerm_virtual_machine_scale_set.test.id}"
      time_grain       = "PT1M"
      statistic         = "Average"
      time_window       = "PT5M"
      time_aggregation = "Average"
      operator          = "LessThan"
      threshold         = 10
    }

    scale_action {
      direction = "Decrease"
    }
  }
}

```

```

    type      = "ChangeCount"
    value     = "2"
    cooldown  = "PT1M"
  }
}

fixed_date {
  timezone = "Pacific Standard Time"
  start    = "2020-07-01T00:00:00Z"
  end      = "2020-07-31T23:59:59Z"
}

notification {
  email {
    send_to_subscription_administrator = true
    send_to_subscription_co_administrator = true
    custom_emails                      = ["admin@contoso.com"]
  }
}
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the AutoScale Setting. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in the AutoScale Setting should be created. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the AutoScale Setting should exist. Changing this forces a new resource to be created.
- `profile` - (Required) Specifies one or more (up to 20) `profile` blocks as defined below.
- `target_resource_id` - (Required) Specifies the resource ID of the resource that the autoscale setting should be added to.
- `enabled` - (Optional) Specifies whether automatic scaling is enabled for the target resource. Defaults to `true`.
- `notification` - (Optional) Specifies a `notification` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `profile` block supports the following:

- `name` - (Required) Specifies the name of the profile.
- `capacity` - (Required) A `capacity` block as defined below.
- `rule` - (Required) One or more (up to 10) `rule` blocks as defined below.
- `fixed_date` - (Optional) A `fixed_date` block as defined below. This cannot be specified if a `recurrence` block is specified.

- `recurrence` - (Optional) A `recurrence` block as defined below. This cannot be specified if a `fixed_date` block is specified.

A `capacity` block supports the following:

- `default` - (Required) The number of instances that are available for scaling if metrics are not available for evaluation. The default is only used if the current instance count is lower than the default. Valid values are between 0 and 1000 .
- `maximum` - (Required) The maximum number of instances for this resource. Valid values are between 0 and 1000 .

NOTE: The maximum number of instances is also limited by the amount of Cores available in the subscription.

- `minimum` - (Required) The minimum number of instances for this resource. Valid values are between 0 and 1000 .

A `rule` block supports the following:

- `metric_trigger` - (Required) A `metric_trigger` block as defined below.
- `scale_action` - (Required) A `scale_action` block as defined below.

A `metric_trigger` block supports the following:

- `metric_name` - (Required) The name of the metric that defines what the rule monitors, such as `Percentage CPU` for `Virtual Machine Scale Sets` and `CpuPercentage` for `App Service Plan` .

NOTE: The allowed value of `metric_name` highly depends on the targeting resource type, please visit [Supported metrics with Azure Monitor \(https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported\)](https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported) for more details.

- `metric_resource_id` - (Required) The ID of the Resource which the Rule monitors.
 - `operator` - (Required) Specifies the operator used to compare the metric data and threshold. Possible values are: `Equals` , `NotEquals` , `GreaterThan` , `GreaterThanOrEqual` , `LessThan` , `LessThanOrEqual` .
 - `statistic` - (Required) Specifies how the metrics from multiple instances are combined. Possible values are `Average` , `Min` and `Max` .
 - `time_aggregation` - (Required) Specifies how the data that's collected should be combined over time. Possible values include `Average` , `Count` , `Maximum` , `Minimum` , `Last` and `Total` . Defaults to `Average` .
 - `time_grain` - (Required) Specifies the granularity of metrics that the rule monitors, which must be one of the pre-defined values returned from the metric definitions for the metric. This value must be between 1 minute and 12 hours and can be formatted as an ISO 8601 string.
 - `time_window` - (Required) Specifies the time range for which data is collected, which must be greater than the delay in metric collection (which varies from resource to resource). This value must be between 5 minutes and 12 hours and be formatted as an ISO 8601 string.
 - `threshold` - (Required) Specifies the threshold of the metric that triggers the scale action.
-

A `scale_action` block supports the following:

- `cooldown` - (Required) The amount of time to wait since the last scaling action before this action occurs. Must be between 1 minute and 1 week and formatted as a ISO 8601 string.
 - `direction` - (Required) The scale direction. Possible values are `Increase` and `Decrease`.
 - `type` - (Required) The type of action that should occur. Possible values are `ChangeCount`, `ExactCount` and `PercentChangeCount`.
 - `value` - (Required) The number of instances involved in the scaling action. Defaults to `1`.
-

A `fixed_date` block supports the following:

- `end` - (Required) Specifies the end date for the profile, formatted as an RFC3339 date string.
 - `start` - (Required) Specifies the start date for the profile, formatted as an RFC3339 date string.
 - `timezone` (Optional) The Time Zone of the `start` and `end` times. A list of possible values can be found here (<https://msdn.microsoft.com/en-us/library/azure/dn931928.aspx>). Defaults to `UTC`.
-

A `recurrence` block supports the following:

- `timezone` - (Required) The Time Zone used for the `hours` field. A list of possible values can be found here (<https://msdn.microsoft.com/en-us/library/azure/dn931928.aspx>). Defaults to `UTC`.
 - `days` - (Required) A list of days that this profile takes effect on. Possible values include `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday` and `Sunday`.
 - `hours` - (Required) A list containing a single item, which specifies the Hour interval at which this recurrence should be triggered (in 24-hour time). Possible values are from `0` to `23`.
 - `minutes` - (Required) A list containing a single item which specifies the Minute interval at which this recurrence should be triggered.
-

A `notification` block supports the following:

- `email` - (Required) A `email` block as defined below.
 - `webhook` - (Optional) One or more `webhook` blocks as defined below.
-

A `email` block supports the following:

- `send_to_subscription_administrator` - (Optional) Should email notifications be sent to the subscription administrator? Defaults to `false`.
 - `send_to_subscription_co_administrator` - (Optional) Should email notifications be sent to the subscription co-administrator? Defaults to `false`.
 - `custom_emails` - (Optional) Specifies a list of custom email addresses to which the email notifications will be sent.
-

A `webhook` block supports the following:

- `service_uri` - (Required) The HTTPS URI which should receive scale notifications.
- `properties` - (Optional) A map of settings.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the AutoScale Setting.

Import

AutoScale Setting can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_autoscale_setting.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/microsoft.insights/autoscalesettings/setting1
```

azurerm_monitor_diagnostic_setting

Manages a Diagnostic Setting for an existing Resource.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

data "azurerm_storage_account" "test" {
  name                = "examplestoracc"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

data "azurerm_key_vault" "test" {
  name                = "example-vault"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_monitor_diagnostic_setting" "test" {
  name                        = "example"
  target_resource_id         = "${data.azurerm_key_vault.test.id}"
  storage_account_id         = "${data.azurerm_storage_account.test.id}"

  log {
    category = "AuditEvent"
    enabled  = false

    retention_policy {
      enabled = false
    }
  }

  metric {
    category = "AllMetrics"

    retention_policy {
      enabled = false
    }
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Diagnostic Setting. Changing this forces a new resource to be created.

- `target_resource_id` - (Required) The ID of an existing Resource on which to configure Diagnostic Settings. Changing this forces a new resource to be created.
- `eventhub_name` - (Optional) Specifies the name of the Event Hub where Diagnostics Data should be sent. Changing this forces a new resource to be created.

NOTE: If this isn't specified then the default Event Hub will be used.

- `eventhub_authorization_rule_id` - (Optional) Specifies the ID of an Event Hub Namespace Authorization Rule used to send Diagnostics Data. Changing this forces a new resource to be created.

NOTE: This can be sourced from the `azurerm_eventhub_namespace_authorization_rule` resource ([/docs/providers/azurerm/r/eventhub_namespace_authorization_rule.html](/docs/providers/azurerm/r/eventhub_namespace_authorization_rule.html)) and is different from a `azurerm_eventhub_authorization_rule` resource ([/docs/providers/azurerm/r/eventhub_authorization_rule.html](/docs/providers/azurerm/r/eventhub_authorization_rule.html)).

NOTE: One of `eventhub_authorization_rule_id`, `log_analytics_workspace_id` and `storage_account_id` must be specified.

- `log` - (Optional) One or more `log` blocks as defined below.

NOTE: At least one `log` or `metric` block must be specified.

- `log_analytics_workspace_id` - (Optional) Specifies the ID of a Log Analytics Workspace where Diagnostics Data should be sent. Changing this forces a new resource to be created.

NOTE: One of `eventhub_authorization_rule_id`, `log_analytics_workspace_id` and `storage_account_id` must be specified.

- `metric` - (Optional) One or more `metric` blocks as defined below.

NOTE: At least one `log` or `metric` block must be specified.

- `storage_account_id` - (Optional) With this parameter you can specify a storage account which should be used to send the logs to. Parameter must be a valid Azure Resource ID. Changing this forces a new resource to be created.

NOTE: One of `eventhub_authorization_rule_id`, `log_analytics_workspace_id` and `storage_account_id` must be specified.

- `log_analytics_destination_type` - (Optional) When set to 'Dedicated' logs sent to a Log Analytics workspace will go into resource specific tables, instead of the legacy AzureDiagnostics table.

NOTE: This setting will only have an effect if a `log_analytics_workspace_id` is provided, and the resource is available for resource-specific logs. As of July 2019, this only includes Azure Data Factory. Please see the documentation (<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/diagnostic-logs-stream-log-store#azure-diagnostics-vs-resource-specific>) for more information.

A `log` block supports the following:

- `category` - (Required) The name of a Diagnostic Log Category for this Resource.

NOTE: The Log Categories available vary depending on the Resource being used. You may wish to use the `azurerm_monitor_diagnostic_categories` Data Source ([/docs/providers/azurerm/d/monitor_diagnostic_categories.html](/docs/providers/azurerm/d/monitor_diagnostic_categories.html)) to identify which categories are available for a given Resource.

- `retention_policy` - (Required) A `retention_policy` block as defined below.
- `enabled` - (Optional) Is this Diagnostic Log enabled? Defaults to `true`.

A `metric` block supports the following:

- `category` - (Required) The name of a Diagnostic Metric Category for this Resource.

NOTE: The Metric Categories available vary depending on the Resource being used. You may wish to use the `azurerm_monitor_diagnostic_categories` Data Source ([/docs/providers/azurerm/d/monitor_diagnostic_categories.html](/docs/providers/azurerm/d/monitor_diagnostic_categories.html)) to identify which categories are available for a given Resource.

- `retention_policy` - (Required) A `retention_policy` block as defined below.
- `enabled` - (Optional) Is this Diagnostic Metric enabled? Defaults to `true`.

A `retention_policy` block supports the following:

- `enabled` - (Required) Is this Retention Policy enabled?
- `days` - (Optional) The number of days for which this Retention Policy should apply.

NOTE: Setting this to `0` will retain the events indefinitely.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Diagnostic Setting.

Import

Diagnostic Settings can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_diagnostics.test /subscriptions/XXX/resourcegroups/resource_group/providers/microsoft.keyvault/vaults/vault|logMonitoring
```

NOTE: This is a Terraform specific Resource ID which uses the format `{resourceId}|{diagnosticSettingName}`

azurerm_monitor_log_profile

Manages a Log Profile (<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-overview-activity-logs#export-the-activity-log-with-a-log-profile>). A Log Profile configures how Activity Logs are exported.

NOTE: It's only possible to configure one Log Profile per Subscription. If you are trying to create more than one Log Profile, an error with `StatusCode=409` will occur.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "logprofiletest-rg"
  location  = "eastus"
}

resource "azurerm_storage_account" "test" {
  name                        = "afscsdfytw"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier                = "Standard"
  account_replication_type    = "GRS"
}

resource "azurerm_eventhub_namespace" "test" {
  name          = "logprofileeventhub"
  location      = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  sku           = "Standard"
  capacity      = 2
}

resource "azurerm_monitor_log_profile" "test" {
  name = "default"

  categories = [
    "Action",
    "Delete",
    "Write",
  ]

  locations = [
    "westus",
    "global",
  ]

  # RootManageSharedAccessKey is created by default with listen, send, manage permissions
  servicebus_rule_id = "${azurerm_eventhub_namespace.test.id}/authorizationrules/RootManageSharedAccessKey"
  storage_account_id = "${azurerm_storage_account.test.id}"

  retention_policy {
    enabled = true
    days    = 7
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Log Profile. Changing this forces a new resource to be created.
- **categories** - (Required) List of categories of the logs.
- **locations** - (Required) List of regions for which Activity Log events are stored or streamed.

- `storage_account_id` - (Optional) The resource ID of the storage account in which the Activity Log is stored. At least one of `storage_account_id` or `servicebus_rule_id` must be set.
 - `servicebus_rule_id` - (Optional) The service bus (or event hub) rule ID of the service bus (or event hub) namespace in which the Activity Log is streamed to. At least one of `storage_account_id` or `servicebus_rule_id` must be set.
 - `retention_policy` - (Required) A `retention_policy` block as documented below. A retention policy for how long Activity Logs are retained in the storage account.
-

The `retention_policy` block supports:

- `enabled` - (Required) A boolean value to indicate whether the retention policy is enabled.
- `days` - (Optional) The number of days for the retention policy. Defaults to 0.

Attributes Reference

The following attributes are exported:

- `id` - The Log Profile resource ID.

Import

A Log Profile can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_log_profile.test /subscriptions/00000000-0000-0000-0000-000000000000/providers/microsoft.insights/logprofiles/test
```

azurerm_monitor_metric_alert

Manages a Metric Alert within Azure Monitor.

Example Usage

```

resource "azurerm_resource_group" "main" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_storage_account" "to_monitor" {
  name                        = "examplestorageaccount"
  resource_group_name        = "${azurerm_resource_group.main.name}"
  location                   = "${azurerm_resource_group.main.location}"
  account_tier                = "Standard"
  account_replication_type    = "LRS"
}

resource "azurerm_monitor_action_group" "main" {
  name                      = "example-actiongroup"
  resource_group_name       = "${azurerm_resource_group.main.name}"
  short_name                = "exampleact"

  webhook_receiver {
    name      = "callmyapi"
    service_uri = "http://example.com/alert"
  }
}

resource "azurerm_monitor_metric_alert" "test" {
  name                = "example-metricalert"
  resource_group_name = "${azurerm_resource_group.main.name}"
  scopes              = ["${azurerm_storage_account.to_monitor.id}"]
  description          = "Action will be triggered when Transactions count is greater than 50."

  criteria {
    metric_namespace = "Microsoft.Storage/storageAccounts"
    metric_name       = "Transactions"
    aggregation       = "Total"
    operator           = "GreaterThan"
    threshold          = 50

    dimension {
      name      = "ApiName"
      operator  = "Include"
      values    = ["*"]
    }
  }

  action {
    action_group_id = "${azurerm_monitor_action_group.main.id}"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Metric Alert. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the Metric Alert instance.
 - `scopes` - (Required) A set of strings of resource IDs at which the metric criteria should be applied.
 - `criteria` - (Required) One or more `criteria` blocks as defined below.
 - `action` - (Optional) One or more `action` blocks as defined below.
 - `enabled` - (Optional) Should this Metric Alert be enabled? Defaults to `true`.
 - `auto_mitigate` - (Optional) Should the alerts in this Metric Alert be auto resolved? Defaults to `false`.
 - `description` - (Optional) The description of this Metric Alert.
 - `frequency` - (Optional) The evaluation frequency of this Metric Alert, represented in ISO 8601 duration format. Possible values are `PT1M`, `PT5M`, `PT15M`, `PT30M` and `PT1H`. Defaults to `PT1M`.
 - `severity` - (Optional) The severity of this Metric Alert. Possible values are `0`, `1`, `2`, `3` and `4`. Defaults to `3`.
 - `window_size` - (Optional) The period of time that is used to monitor alert activity, represented in ISO 8601 duration format. This value must be greater than `frequency`. Possible values are `PT1M`, `PT5M`, `PT15M`, `PT30M`, `PT1H`, `PT6H`, `PT12H` and `P1D`. Defaults to `PT5M`.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

An `action` block supports the following:

- `action_group_id` - (Required) The ID of the Action Group can be sourced from the `azurerm_monitor_action_group` resource ([/docs/providers/azurerm/r/monitor_action_group.html](/docs/providers/azurerm/r/monitor_action_group.html))
 - `webhook_properties` - (Optional) The map of custom string properties to include with the post operation. These data are appended to the webhook payload.
-

A `criteria` block supports the following:

- `metric_namespace` - (Required) One of the metric namespaces to be monitored.
 - `metric_name` - (Required) One of the metric names to be monitored.
 - `aggregation` - (Required) The statistic that runs over the metric values. Possible values are `Average`, `Count`, `Minimum`, `Maximum` and `Total`.
 - `operator` - (Required) The criteria operator. Possible values are `Equals`, `NotEquals`, `GreaterThan`, `GreaterThanOrEqual`, `LessThan` and `LessThanOrEqual`.
 - `threshold` - (Required) The criteria threshold value that activates the alert.
 - `dimension` - (Optional) One or more `dimension` blocks as defined below.
-

A `dimension` block supports the following:

- `name` - (Required) One of the dimension names.
- `operator` - (Required) The dimension operator. Possible values are `Include` and `Exclude`.
- `values` - (Required) The list of dimension values.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the metric alert.

Import

Metric Alerts can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_metric_alert.main /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-resources/providers/microsoft.insights/metricalerts/example-metricalert
```

azurerm_monitor_metric_alertrule

Manages a metric-based alert rule (<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitor-quick-resource-metric-alert-portal>) in Azure Monitor.

Example Usage (CPU Percentage of a virtual machine)

```
resource "azurerm_monitor_metric_alertrule" "test" {
  name = "${azurerm_virtual_machine.test.name}-cpu"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location = "${azurerm_resource_group.test.location}"

  description = "An alert rule to watch the metric Percentage CPU"

  enabled = true

  resource_id = "${azurerm_virtual_machine.test.id}"
  metric_name = "Percentage CPU"
  operator = "GreaterThan"
  threshold = 75
  aggregation = "Average"
  period = "PT5M"

  email_action {
    send_to_service_owners = false

    custom_emails = [
      "some.user@example.com",
    ]
  }

  webhook_action {
    service_uri = "https://example.com/some-url"

    properties = {
      severity = "incredible"
      acceptance_test = "true"
    }
  }
}
```

Example Usage (Storage usage of a SQL Database)

```

resource "azurerm_monitor_metric_alertrule" "test" {
  name                = "${azurerm_sql_database.test.name}-storage"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  description = "An alert rule to watch the metric Storage"

  enabled = true

  resource_id = "${azurerm_sql_database.test.id}"
  metric_name = "storage"
  operator    = "GreaterThan"
  threshold   = 1073741824
  aggregation = "Maximum"
  period      = "PT10M"

  email_action {
    send_to_service_owners = false

    custom_emails = [
      "some.user@example.com",
    ]
  }

  webhook_action {
    service_uri = "https://example.com/some-url"

    properties = {
      severity      = "incredible"
      acceptance_test = "true"
    }
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the alert rule. Changing this forces a new resource to be created.
 - **resource_group_name** - (Required) The name of the resource group in which to create the alert rule. Changing this forces a new resource to be created.
 - **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
 - **description** - (Optional) A verbose description of the alert rule that will be included in the alert email.
 - **enabled** - (Optional) If `true`, the alert rule is enabled. Defaults to `true`.
-
- **resource_id** - (Required) The ID of the resource monitored by the alert rule.
 - **metric_name** - (Required) The metric that defines what the rule monitors.

For a comprehensive reference of supported `metric_name` values for types of `resource` refer to Supported metrics with Azure Monitor (<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-supported-metrics>) in the Azure documentation. In the referred table, the column "Metric" corresponds to supported values for `metric_name`.

- `operator` - (Required) The operator used to compare the metric data and the threshold. Possible values are `GreaterThan`, `GreaterThanOrEqual`, `LessThan`, and `LessThanOrEqual`.
- `threshold` - (Required) The threshold value that activates the alert.
- `period` - (Required) The period of time formatted in ISO 8601 duration format (https://en.wikipedia.org/wiki/ISO_8601#Durations) that is used to monitor the alert activity based on the threshold. The period must be between 5 minutes and 1 day.
- `aggregation` - (Required) Defines how the metric data is combined over time. Possible values are `Average`, `Minimum`, `Maximum`, `Total`, and `Last`.

- `email_action` - (Optional) A `email_action` block as defined below.
- `webhook_action` - (Optional) A `webhook_action` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource. Changing this forces a new resource to be created.

`email_action` supports the following:

- `send_to_service_owners` - (Optional) If `true`, the administrators (service and co-administrators) of the subscription are notified when the alert is triggered. Defaults to `false`.
- `custom_emails` - (Optional) A list of email addresses to be notified when the alert is triggered.

`webhook_action` supports the following:

- `service_uri` - (Required) The service uri of the webhook to POST the notification when the alert is triggered.
- `properties` - (Optional) A dictionary of custom properties to include with the webhook POST operation payload.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the alert rule.

Import

Metric Alert Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_monitor_metric_alertrule.alertrule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.insights/alertrules/alertrule1
```


azurerm_mssql_elasticpool

Allows you to manage an Azure SQL Elastic Pool via the 2017-10-01-preview API which allows for vCore and DTU based configurations.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "my-resource-group"
  location  = "westeurope"
}

resource "azurerm_sql_server" "test" {
  name                        = "my-sql-server"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  version                    = "12.0"
  administrator_login        = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_mssql_elasticpool" "test" {
  name                = "test-epool"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  server_name          = "${azurerm_sql_server.test.name}"
  max_size_gb          = 756

  sku {
    name      = "GP_Gen5"
    tier       = "GeneralPurpose"
    family    = "Gen5"
    capacity  = 4
  }

  per_database_settings {
    min_capacity = 0.25
    max_capacity = 4
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the elastic pool. This needs to be globally unique. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the elastic pool. This must be the same as the resource group of the underlying SQL server.

- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server on which to create the elastic pool. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `per_database_settings` - (Required) A `per_database_settings` block as defined below.
- `max_size_gb` - (Optional) The max data size of the elastic pool in gigabytes. Conflicts with `max_size_bytes`.
- `max_size_bytes` - (Optional) The max data size of the elastic pool in bytes. Conflicts with `max_size_gb`.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zone_redundant` - (Optional) Whether or not this elastic pool is zone redundant. `tier` needs to be `Premium` for DTU based or `BusinessCritical` for vCore based `sku`. Defaults to `false`.

`sku` supports the following:

- `name` - (Required) Specifies the SKU Name for this Elasticpool. The name of the SKU, will be either vCore based `tier + family` pattern (e.g. `GP_Gen4`, `BC_Gen5`) or the DTU based `BasicPool`, `StandardPool`, or `PremiumPool` pattern.
- `capacity` - (Required) The scale up/out capacity, representing server's compute units. For more information see the documentation for your Elasticpool configuration: vCore-based (<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-vcore-resource-limits-elastic-pools>) or DTU-based (<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dtu-resource-limits-elastic-pools>).
- `tier` - (Required) The tier of the particular SKU. Possible values are `GeneralPurpose`, `BusinessCritical`, `Basic`, `Standard`, or `Premium`. For more information see the documentation for your Elasticpool configuration: vCore-based (<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-vcore-resource-limits-elastic-pools>) or DTU-based (<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dtu-resource-limits-elastic-pools>).
- `family` - (Optional) The family of hardware `Gen4` or `Gen5`.

`per_database_settings` supports the following:

- `min_capacity` - (Required) The minimum capacity all databases are guaranteed.
 - `max_capacity` - (Required) The maximum capacity any one database can consume.
-

Attributes Reference

The following attributes are exported:

- `id` - The MsSQL Elastic Pool ID.

Import

SQL Elastic Pool can be imported using the `resource id`, e.g.

```
terraform import azurerm_mssql_elasticpool.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/elasticPools/myelasticpoolname
```

azurerm_mysql_configuration

Sets a MySQL Configuration value on a MySQL Server.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mysql_server" "test" {
  name                = "mysql-server-1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name      = "B_Gen5_2"
    capacity = 2
    tier       = "Basic"
    family    = "Gen4"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "psqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "5.7"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_mysql_configuration" "test" {
  name                = "interactive_timeout"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name         = "${azurerm_mysql_server.test.name}"
  value               = "600"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the MySQL Configuration, which needs to be a valid MySQL configuration name (<https://dev.mysql.com/doc/refman/5.7/en/server-configuration.html>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the MySQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the MySQL Server exists. Changing this forces a new resource to be created.
- `value` - (Required) Specifies the value of the MySQL Configuration. See the MySQL documentation for valid values.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MySQL Configuration.

Import

MySQL Configurations can be imported using the `resource id`, e.g.

```
terraform import azurerm_mysql_configuration.interactive_timeout /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMySQL/servers/server1/configurations/interactive_timeout
```

azurerm_mysql_database

Manages a MySQL Database within a MySQL Server

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mysql_server" "test" {
  name                = "mysql-server-1"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name      = "B_Gen5_2"
    capacity  = 2
    tier       = "Basic"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "mysqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "5.7"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_mysql_database" "test" {
  name                = "exampleldb"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name         = "${azurerm_mysql_server.test.name}"
  charset              = "utf8"
  collation            = "utf8_unicode_ci"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the MySQL Database, which needs to be a valid MySQL identifier (<https://dev.mysql.com/doc/refman/5.7/en/identifiers.html>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the MySQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the MySQL Server exists. Changing this forces a new resource to be created.
- `charset` - (Required) Specifies the Charset for the MySQL Database, which needs to be a valid MySQL Charset (<https://dev.mysql.com/doc/refman/5.7/en/charset-charsets.html>). Changing this forces a new resource to be created.
- `collation` - (Required) Specifies the Collation for the MySQL Database, which needs to be a valid MySQL Collation (<https://dev.mysql.com/doc/refman/5.7/en/charset-mysql.html>). Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MySQL Database.

Import

MySQL Database's can be imported using the `resource id`, e.g.

```
terraform import azurerm_mysql_database.database1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMySQL/servers/server1/databases/database1
```

azurerm_mysql_firewall_rule

Manages a Firewall Rule for a MySQL Server

Example Usage (Single IP Address)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mysql_server" "test" {
  # ...
}

resource "azurerm_mysql_firewall_rule" "test" {
  name                        = "office"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  server_name                 = "${azurerm_mysql_server.test.name}"
  start_ip_address            = "40.112.8.12"
  end_ip_address              = "40.112.8.12"
}
```

Example Usage (IP Range)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mysql_server" "test" {
  # ...
}

resource "azurerm_mysql_firewall_rule" "test" {
  name                        = "office"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  server_name                 = "${azurerm_mysql_server.test.name}"
  start_ip_address            = "40.112.0.0"
  end_ip_address              = "40.112.255.255"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the MySQL Firewall Rule. Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the MySQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the MySQL Server exists. Changing this forces a new resource to be created.
- `start_ip_address` - (Required) Specifies the Start IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.
- `end_ip_address` - (Required) Specifies the End IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MySQL Firewall Rule.

Import

MySQL Firewall Rule's can be imported using the `resource id`, e.g.

```
terraform import azurerm_mysql_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMySQL/servers/server1/firewallRules/rule1
```

azurerm_mysql_server

Manages a MySQL Server.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_mysql_server" "test" {
  name                  = "mysql-server-1"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name  = "${azurerm_resource_group.test.name}"

  sku {
    name      = "B_Gen5_2"
    capacity = 2
    tier       = "Basic"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "mysqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "5.7"
  ssl_enforcement              = "Enabled"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the MySQL Server. Changing this forces a new resource to be created. This needs to be globally unique within Azure.
- **resource_group_name** - (Required) The name of the resource group in which to create the MySQL Server. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **sku** - (Required) A `sku` block as defined below.
- **storage_profile** - (Required) A `storage_profile` block as defined below.

- `administrator_login` - (Required) The Administrator Login for the MySQL Server. Changing this forces a new resource to be created.
 - `administrator_login_password` - (Required) The Password associated with the `administrator_login` for the MySQL Server.
 - `version` - (Required) Specifies the version of MySQL to use. Valid values are `5.6` and `5.7`. Changing this forces a new resource to be created.
 - `ssl_enforcement` - (Required) Specifies if SSL should be enforced on connections. Possible values are `Enabled` and `Disabled`.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

`sku` supports the following:

- `name` - (Required) Specifies the SKU Name for this MySQL Server. The name of the SKU, follows the `tier + family + cores` pattern (e.g. `B_Gen4_1`, `GP_Gen5_8`). For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/mysql/servers/create#sku>).
 - `capacity` - (Required) The scale up/out capacity, representing server's compute units.
 - `tier` - (Required) The tier of the particular SKU. Possible values are `Basic`, `GeneralPurpose`, and `MemoryOptimized`. For more information see the product documentation (<https://docs.microsoft.com/en-us/azure/mysql/concepts-pricing-tiers>).
 - `family` - (Required) The family of hardware `Gen4` or `Gen5`, before selecting your family check the product documentation (<https://docs.microsoft.com/en-us/azure/mysql/concepts-pricing-tiers#compute-generations-vcors-and-memory>) for availability in your region.
-

`storage_profile` supports the following:

- `storage_mb` - (Required) Max storage allowed for a server. Possible values are between `5120 MB(5GB)` and `1048576 MB(1TB)` for the Basic SKU and between `5120 MB(5GB)` and `4194304 MB(4TB)` for General Purpose/Memory Optimized SKUs. For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/mysql/servers/create#StorageProfile>).
- `backup_retention_days` - (Optional) Backup retention days for the server, supported values are between `7` and `35` days.
- `geo_redundant_backup` - (Optional) Enable Geo-redundant or not for server backup. Valid values for this property are `Enabled` or `Disabled`, not supported for the `basic` tier.
- `auto_grow` - (Optional) Defines whether autogrow is enabled or disabled for the storage. Valid values are `Enabled` or `Disabled`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MySQL Server.

- `fqdn` - The FQDN of the MySQL Server.

Import

MySQL Server's can be imported using the `resource id`, e.g.

```
terraform import azurerm_mysql_server.server1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforMySQL/servers/server1
```

azurerm_mysql_virtual_network_rule

Manages a MySQL Virtual Network Rule.

NOTE: MySQL Virtual Network Rules can only be used with SKU Tiers of `GeneralPurpose` or `MemoryOptimized` (<https://docs.microsoft.com/en-us/azure/mysql/concepts-data-access-and-security-vnet>)

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-vnet"
  address_space   = ["10.7.29.0/29"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "internal" {
  name                 = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.7.29.0/29"
  service_endpoints    = ["Microsoft.Sql"]
}

resource "azurerm_mysql_server" "test" {
  name                  = "mysql-server-1"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  administrator_login   = "mysqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version              = "5.7"
  ssl_enforcement      = "Enabled"

  sku {
    name      = "GP_Gen5_2"
    capacity = 2
    tier      = "GeneralPurpose"
    family   = "Gen5"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }
}

resource "azurerm_mysql_virtual_network_rule" "test" {
  name            = "mysql-vnet-rule"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name      = "${azurerm_mysql_server.test.name}"
  subnet_id        = "${azurerm_subnet.internal.id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the MySQL Virtual Network Rule. Cannot be empty and must only contain

alphanumeric characters and hyphens. Cannot start with a number, and cannot start or end with a hyphen. Changing this forces a new resource to be created.

NOTE: `name` must be between 1-128 characters long and must satisfy all of the requirements below: 1. Contains only alphanumeric and hyphen characters 2. Cannot start with a number or hyphen 3. Cannot end with a hyphen

- `resource_group_name` - (Required) The name of the resource group where the MySQL server resides. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server to which this MySQL virtual network rule will be applied to. Changing this forces a new resource to be created.
- `subnet_id` - (Required) The ID of the subnet that the MySQL server will be connected to.

NOTE: Due to a bug in the Azure API (<https://github.com/Azure/azure-rest-api-specs/issues/3719>) this resource currently doesn't expose the `ignore_missing_vnet_service_endpoint` field and defaults this to `false`. Terraform will check during the provisioning of the Virtual Network Rule that the Subnet contains the Service Rule to verify that the Virtual Network Rule can be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the MySQL Virtual Network Rule.

Import

MySQL Virtual Network Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_mysql_virtual_network_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.DBforMySQL/servers/myserver/virtualNetworkRules/vnetrulename
```

azurerm_connection_monitor

Configures a Network Connection Monitor to monitor communication between a Virtual Machine and an endpoint using a Network Watcher.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "connection-monitor-rg"
  location  = "West US"
}

resource "azurerm_network_watcher" "test" {
  name                = "network-watcher"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_virtual_network" "test" {
  name                = "production-network"
  address_space       = ["10.0.0.0/16"]
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "cmtest-nic"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                        = "testconfiguration1"
    subnet_id                  = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "cmtest-vm"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_interface_ids = ["${azurerm_network_interface.test.id}"]
  vm_size             = "Standard_F2"

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
```

```

sku          = "10.04-LTS"
version      = "latest"
}

storage_os_disk {
  name          = "osdisk"
  caching       = "ReadWrite"
  create_option = "FromImage"
  managed_disk_type = "Standard_LRS"
}

os_profile {
  computer_name  = "cmtest-vm"
  admin_username = "testadmin"
  admin_password = "Password1234!"
}

os_profile_linux_config {
  disable_password_authentication = false
}
}

resource "azurerm_virtual_machine_extension" "test" {
  name                  = "cmtest-vm-network-watcher"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  virtual_machine_name  = "${azurerm_virtual_machine.test.name}"
  publisher             = "Microsoft.Azure.NetworkWatcher"
  type                 = "NetworkWatcherAgentLinux"
  type_handler_version  = "1.4"
  auto_upgrade_minor_version = true
}

resource "azurerm_network_connection_monitor" "test" {
  name                  = "cmtest-connectionmonitor"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  network_watcher_name  = "${azurerm_network_watcher.test.name}"

  source {
    virtual_machine_id = "${azurerm_virtual_machine.test.id}"
  }

  destination {
    address = "terraform.io"
    port    = 80
  }

  depends_on = ["azurerm_virtual_machine_extension.test"]
}

```

NOTE: This Resource requires that the Network Watcher Agent Virtual Machine Extension (<https://docs.microsoft.com/en-us/azure/network-watcher/connection-monitor>) is installed on the Virtual Machine before monitoring can be started. The extension can be installed via the `azurerm_virtual_machine_extension` resource ([/docs/providers/azurerm/r/virtual_machine_extension.html](https://docs/providers/azurerm/r/virtual_machine_extension.html)).

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Network Connection Monitor. Changing this forces a new resource to be created.
- `network_watcher_name` - (Required) The name of the Network Watcher. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Connection Monitor. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `auto_start` - (Optional) Specifies whether the connection monitor will start automatically once created. Defaults to `true` . Changing this forces a new resource to be created.
- `interval_in_seconds` - (Optional) Monitoring interval in seconds. Defaults to `60` .
- `source` - (Required) A `source` block as defined below.
- `destination` - (Required) A `destination` block as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `source` block contains:

- `virtual_machine_id` - (Required) The ID of the Virtual Machine to monitor connectivity from.
- `port` - (Optional) The port on the Virtual Machine to monitor connectivity from. Defaults to `0` (Dynamic Port Assignment).

A `destination` block contains:

- `virtual_machine_id` - (Optional) The ID of the Virtual Machine to monitor connectivity to.
- `address` - (Optional) IP address or domain name to monitor connectivity to.
- `port` - (Required) The port on the destination to monitor connectivity to.

NOTE: One of `virtual_machine_id` or `address` must be specified.

Attributes Reference

The following attributes are exported:

- `id` - The Connection Monitor ID.

Import

Connection Monitors can be imported using the `resource_id`, e.g.

```
terraform import azurerm_network_connection_monitor.monitor1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/networkWatchers/watcher1/connectionMonitors/monitor1
```

azurerm_network_ddos_protection_plan

Manages an AzureNetwork DDoS Protection Plan.

NOTE Azure only allows one DDoS Protection Plan per region.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_network_ddos_protection_plan" "test" {
  name                = "example-protection-plan"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Network DDoS Protection Plan. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the resource. Changing this forces a new resource to be created.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- **id** - The Resource ID of the DDoS Protection Plan
- **virtual_network_ids** - The Resource ID list of the Virtual Networks associated with DDoS Protection Plan.

Import

Azure DDoS Protection Plan can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_ddos_protection_plan.test /subscriptions/00000000-0000-0000-0000-00000000  
0000/resourceGroups/group1/providers/Microsoft.Network/ddosProtectionPlans/testddosplan
```

azurerm_network_interface_application_gateway_backend_address_pool_association

Manages the association between a Network Interface and a Application Gateway's Backend Address Pool.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "frontend" {
  name                 = "frontend"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.254.0.0/24"
}

resource "azurerm_subnet" "backend" {
  name                 = "backend"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.254.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name             = "example-pip"
  location         = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Dynamic"
}

# since these variables are re-used - a locals block makes this more maintainable
locals {
  backend_address_pool_name = "${azurerm_virtual_network.test.name}-beap"
  frontend_port_name        = "${azurerm_virtual_network.test.name}-feport"
  frontend_ip_configuration_name = "${azurerm_virtual_network.test.name}-feip"
  http_setting_name         = "${azurerm_virtual_network.test.name}-be-htst"
  listener_name             = "${azurerm_virtual_network.test.name}-httplstn"
  request_routing_rule_name = "${azurerm_virtual_network.test.name}-rqrt"
}

resource "azurerm_application_gateway" "network" {
  name                 = "example-appgateway"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"

  sku {
    name     = "Standard_Small"
    tier      = "Standard"
    capacity = 2
  }

  gateway_ip_configuration {
    name        = "my-gateway-ip-configuration"
    subnet_id   = "${azurerm_subnet.frontend.id}"
  }

  frontend_port {
    name = "${local.frontend_port_name}"
    port = 80
  }

  frontend_ip_configuration {
    name                 = "${local.frontend_ip_configuration_name}"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }

  backend_address_pool {
    name = "${local.backend_address_pool_name}"
  }

  backend_http_settings {
    name                  = "${local.backend_http_settings_name}"
    port                 = 80
    protocol              = "Http"
    cookie_based_routing = false
  }
}
```



```

name                = "${local.http_setting_name}"
cookie_based_affinity = "Disabled"
port                = 80
protocol            = "Http"
request_timeout     = 1
}

http_listener {
  name                        = "${local.listener_name}"
  frontend_ip_configuration_name = "${local.frontend_ip_configuration_name}"
  frontend_port_name         = "${local.frontend_port_name}"
  protocol                   = "Http"
}

request_routing_rule {
  name                = "${local.request_routing_rule_name}"
  rule_type           = "Basic"
  http_listener_name  = "${local.listener_name}"
  backend_address_pool_name = "${local.backend_address_pool_name}"
  backend_http_settings_name = "${local.http_setting_name}"
}
}

resource "azurerm_network_interface" "test" {
  name                = "example-nic"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id          = "${azurerm_subnet.frontend.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_network_interface_application_gateway_backend_address_pool_association" "test" {
  network_interface_id = "${azurerm_network_interface.test.id}"
  ip_configuration_name = "testconfiguration1"
  backend_address_pool_id = "${azurerm_application_gateway.network.backend_address_pool.0.id}"
}

```

Argument Reference

The following arguments are supported:

- `network_interface_id` - (Required) The ID of the Network Interface. Changing this forces a new resource to be created.
- `ip_configuration_name` - (Required) The Name of the IP Configuration within the Network Interface which should be connected to the Backend Address Pool. Changing this forces a new resource to be created.
- `backend_address_pool_id` - (Required) The ID of the Application Gateway's Backend Address Pool which this Network Interface which should be connected to. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The (Terraform specific) ID of the Association between the Network Interface and the Application Gateway Backend Address Pool.

Import

Associations between Network Interfaces and Application Gateway Backend Address Pools can be imported using the `resource id`, e.g.

```

terraform import azurerm_network_interface_application_gateway_backend_address_pool_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.network/networkInterfaces/nic1/ipConfigurations/example1/subscriptions/00000000-0000-0000-0000-000000000000/00/resourceGroups/group1/providers/Microsoft.Network/applicationGateways/gateway1/backendAddressPools/pool1

```

NOTE: This ID is specific to Terraform - and is of the format `{networkInterfaceId}/ipConfigurations/{ipConfigurationName}|{backendAddressPoolId}`.

azurerm_network_interface_application_security_group_association

Manages the association between a Network Interface and a Application Security Group.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_application_security_group" "test" {
  name            = "example-asg"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_network_interface" "test" {
  name            = "example-nic"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                       = "testconfiguration1"
    subnet_id                 = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
    application_security_group_ids = ["${azurerm_application_security_group.test.id}"]
  }
}

resource "azurerm_network_interface_application_security_group_association" "test" {
  network_interface_id      = "${azurerm_network_interface.test.id}"
  ip_configuration_name     = "testconfiguration1"
  application_security_group_id = "${azurerm_application_security_group.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `network_interface_id` - (Required) The ID of the Network Interface. Changing this forces a new resource to be created.
- `ip_configuration_name` - (Required) The Name of the IP Configuration within the Network Interface which should be connected to the Application Security Group. Changing this forces a new resource to be created.
- `application_security_group_id` - (Required) The ID of the Application Security Group which this Network Interface which should be

connected to. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The (Terraform specific) ID of the Association between the Network Interface and the Application Security Group.

Import

Associations between Network Interfaces and Application Security Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_interface_application_security_group_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.network/networkInterfaces/nic1/ipConfigurations/example/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/applicationSecurityGroups/securityGroup1
```

NOTE: This ID is specific to Terraform - and is of the format `{networkInterfaceId}/ipConfigurations/{ipConfigurationName}|{applicationSecurityGroupId}`.

azurerm_network_interface_backend_address_pool_association

Manages the association between a Network Interface and a Load Balancer's Backend Address Pool.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name            = "example-pip"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Static"
}

resource "azurerm_lb" "test" {
  name            = "example-lb"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name            = "primary"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_backend_address_pool" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id      = "${azurerm_lb.test.id}"
  name                 = "acctestpool"
}

resource "azurerm_network_interface" "test" {
  name            = "example-nic"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id           = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_network_interface_backend_address_pool_association" "test" {
  network_interface_id = "${azurerm_network_interface.test.id}"
  ip_configuration_name = "testconfiguration1"
  backend_address_pool_id = "${azurerm_lb_backend_address_pool.test.id}"
}

```

Argument Reference

The following arguments are supported:

- `network_interface_id` - (Required) The ID of the Network Interface. Changing this forces a new resource to be created.
- `ip_configuration_name` - (Required) The Name of the IP Configuration within the Network Interface which should be connected to the Backend Address Pool. Changing this forces a new resource to be created.
- `backend_address_pool_id` - (Required) The ID of the Load Balancer Backend Address Pool which this Network Interface which should be connected to. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The (Terraform specific) ID of the Association between the Network Interface and the Load Balancers Backend Address Pool.

Import

Associations between Network Interfaces and Load Balancer Backend Address Pools can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_interface_backend_address_pool_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.network/networkInterfaces/nic1/ipConfigurations/example1/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/backendAddressPools/pool1
```

NOTE: This ID is specific to Terraform - and is of the format

`{networkInterfaceId}/ipConfigurations/{ipConfigurationName}|{backendAddressPoolId}`.

azurerm_network_interface

Manages a Network Interface located in a Virtual Network, usually attached to a Virtual Machine.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_virtual_network" "test" {
  name            = "acceptanceTestVirtualNetwork1"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "testsubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name            = "acceptanceTestNetworkInterface1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                 = "testconfiguration1"
    subnet_id            = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }

  tags = {
    environment = "staging"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the network interface. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the network interface. Changing this forces a new resource to be created.

- `location` - (Required) The location/region where the network interface is created. Changing this forces a new resource to be created.
- `network_security_group_id` - (Optional) The ID of the Network Security Group to associate with the network interface.
- `internal_dns_name_label` - (Optional) Relative DNS name for this NIC used for internal communications between VMs in the same VNet
- `enable_ip_forwarding` - (Optional) Enables IP Forwarding on the NIC. Defaults to `false`.
- `enable_accelerated_networking` - (Optional) Enables Azure Accelerated Networking using SR-IOV. Only certain VM instance sizes are supported. Refer to [Create a Virtual Machine with Accelerated Networking](https://docs.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-cli) (<https://docs.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-cli>). Defaults to `false`.

NOTE: when using Accelerated Networking in an Availability Set - the Availability Set must be deployed on an Accelerated Networking enabled cluster.

- `dns_servers` - (Optional) List of DNS servers IP addresses to use for this NIC, overrides the VNet-level server list
- `ip_configuration` - (Required) One or more `ip_configuration` associated with this NIC as documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `ip_configuration` block supports:

- `name` - (Required) User-defined name of the IP.
- `subnet_id` - (Optional) Reference to a subnet in which this NIC has been created. Required when `private_ip_address_version` is IPv4.
- `private_ip_address` - (Optional) Static IP Address.
- `private_ip_address_allocation` - (Required) Defines how a private IP address is assigned. Options are Static or Dynamic.
- `private_ip_address_version` - (Optional) The IP Version to use. Possible values are `IPv4` or `IPv6`. Defaults to `IPv4`.
- `public_ip_address_id` - (Optional) Reference to a Public IP Address to associate with this NIC
- `application_gateway_backend_address_pools_ids` - (Optional / **Deprecated**) List of Application Gateway Backend Address Pool IDs references to which this NIC belongs

NOTE: At this time Network Interface <Application Gateway Backend Address Pool associations need to be configured both using this field (which is now Deprecated) and using the `azurerm_network_interface_application_gateway_backend_address_pool_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `load_balancer_backend_address_pools_ids` - (Optional / **Deprecated**) List of Load Balancer Backend Address Pool IDs references to which this NIC belongs

NOTE: At this time Network Interface <Load Balancer Backend Address Pool associations need to be configured both using this field (which is now Deprecated) and using the

`azurerm_network_interface_backend_address_pool_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `load_balancer_inbound_nat_rules_ids` - (Optional / **Deprecated**) List of Load Balancer Inbound Nat Rules IDs involving this NIC

NOTE: At this time Network Interface <Load Balancer Inbound NAT Rule associations need to be configured both using this field (which is now Deprecated) and using the `azurerm_network_interface_nat_rule_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `application_security_group_ids` - (Optional / **Deprecated**) List of Application Security Group IDs which should be attached to this NIC

NOTE: At this time Network Interface <Application Security Group associations need to be configured both using this field (which is now Deprecated) and using the `azurerm_network_interface_application_security_group_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `primary` - (Optional) Is this the Primary Network Interface? If set to `true` this should be the first `ip_configuration` in the array.

Attributes Reference

The following attributes are exported:

- `id` - The Virtual Network Interface ID.
- `mac_address` - The media access control (MAC) address of the network interface.
- `private_ip_address` - The first private IP address of the network interface.
- `private_ip_addresses` - The private IP addresses of the network interface.
- `virtual_machine_id` - Reference to a VM with which this NIC has been associated.
- `applied_dns_servers` - If the VM that uses this NIC is part of an Availability Set, then this list will have the union of all DNS servers from all NICs that are part of the Availability Set

Import

Network Interfaces can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_interface.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.network/networkInterfaces/nic1
```

azurerm_network_interface_nat_rule_association

Manages the association between a Network Interface and a Load Balancer's NAT Rule.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "internal"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name                = "example-pip"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
}

resource "azurerm_lb" "test" {
  name                = "example-lb"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name                = "primary"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_nat_rule" "test" {
  resource_group_name      = "${azurerm_resource_group.test.name}"
  loadbalancer_id          = "${azurerm_lb.test.id}"
  name                     = "RDPAccess"
  protocol                 = "Tcp"
  frontend_port            = 3389
  backend_port             = 3389
  frontend_ip_configuration_name = "primary"
}

resource "azurerm_network_interface" "test" {
  name = "example-nic"
```

```

location          = "${azurerm_resource_group.test.location}"
resource_group_name = "${azurerm_resource_group.test.name}"

ip_configuration {
  name              = "testconfiguration1"
  subnet_id         = "${azurerm_subnet.test.id}"
  private_ip_address_allocation = "Dynamic"
}

resource "azurerm_network_interface_nat_rule_association" "test" {
  network_interface_id = "${azurerm_network_interface.test.id}"
  ip_configuration_name = "testconfiguration1"
  nat_rule_id          = "${azurerm_lb_nat_rule.test.id}"
}

```

Argument Reference

The following arguments are supported:

- `network_interface_id` - (Required) The ID of the Network Interface. Changing this forces a new resource to be created.
- `ip_configuration_name` - (Required) The Name of the IP Configuration within the Network Interface which should be connected to the NAT Rule. Changing this forces a new resource to be created.
- `nat_rule_id` - (Required) The ID of the Load Balancer NAT Rule which this Network Interface which should be connected to. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The (Terraform specific) ID of the Association between the Network Interface and the Load Balancers NAT Rule.

Import

Associations between Network Interfaces and Load Balancer NAT Rule can be imported using the `resource id`, e.g.

```

terraform import azurerm_network_interface_nat_rule_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.network/networkInterfaces/nic1/ipConfigurations/example1/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/loadBalancers/lb1/inboundNatRules/rule1

```

NOTE: This ID is specific to Terraform - and is of the format

```
{networkInterfaceId}/ipConfigurations/{ipConfigurationName}|{natRuleId}.
```

azurerm_packet_capture

Configures Network Packet Capturing against a Virtual Machine using a Network Watcher.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "packet-capture-rg"
  location  = "West Europe"
}

resource "azurerm_network_watcher" "test" {
  name                = "network-watcher"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_virtual_network" "test" {
  name                = "production-network"
  address_space       = ["10.0.0.0/16"]
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "pctest-nic"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id           = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "pctest-vm"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_interface_ids = ["${azurerm_network_interface.test.id}"]
  vm_size             = "Standard_F2"

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }
}
```

```

}

storage_os_disk {
  name          = "osdisk"
  caching        = "ReadWrite"
  create_option  = "FromImage"
  managed_disk_type = "Standard_LRS"
}

os_profile {
  computer_name  = "pctest-vm"
  admin_username = "testadmin"
  admin_password = "Password1234!"
}

os_profile_linux_config {
  disable_password_authentication = false
}
}

resource "azurerm_virtual_machine_extension" "test" {
  name                  = "network-watcher"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  virtual_machine_name  = "${azurerm_virtual_machine.test.name}"
  publisher             = "Microsoft.Azure.NetworkWatcher"
  type                 = "NetworkWatcherAgentLinux"
  type_handler_version  = "1.4"
  auto_upgrade_minor_version = true
}

resource "azurerm_storage_account" "test" {
  name                  = "pctestsa"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  location              = "${azurerm_resource_group.test.location}"
  account_tier          = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_network_packet_capture" "test" {
  name                  = "pctestcapture"
  network_watcher_name = "${azurerm_network_watcher.test.name}"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  target_resource_id    = "${azurerm_virtual_machine.test.id}"

  storage_location {
    storage_account_id = "${azurerm_storage_account.test.id}"
  }

  depends_on = ["azurerm_virtual_machine_extension.test"]
}

```

NOTE: This Resource requires that the Network Watcher Virtual Machine Extension (<https://docs.microsoft.com/azure/network-watcher/network-watcher-packet-capture-manage-portal#before-you-begin>) is installed on the Virtual Machine before capturing can be enabled which can be installed via the `azurerm_virtual_machine_extension` resource ([/docs/providers/azurerm/r/virtual_machine_extension.html](https://docs/providers/azurerm/r/virtual_machine_extension.html)).

Argument Reference

The following arguments are supported:

- `name` - (Required) The name to use for this Network Packet Capture. Changing this forces a new resource to be created.
- `network_watcher_name` - (Required) The name of the Network Watcher. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Network Watcher exists. Changing this forces a new resource to be created.
- `target_resource_id` - (Required) The ID of the Resource to capture packets from. Changing this forces a new resource to be created.

NOTE: Currently only Virtual Machines ID's are supported.

- `maximum_bytes_per_packet` - (Optional) The number of bytes captured per packet. The remaining bytes are truncated. Defaults to `0` (Entire Packet Captured). Changing this forces a new resource to be created.
- `maximum_bytes_per_session` - (Optional) Maximum size of the capture in Bytes. Defaults to `1073741824` (1GB). Changing this forces a new resource to be created.
- `maximum_capture_duration` - (Optional) The maximum duration of the capture session in seconds. Defaults to `18000` (5 hours). Changing this forces a new resource to be created.
- `storage_location` - (Required) A `storage_location` block as defined below. Changing this forces a new resource to be created.
- `filter` - (Optional) One or more `filter` blocks as defined below. Changing this forces a new resource to be created.

A `storage_location` block contains:

- `file_path` - (Optional) A valid local path on the targeting VM. Must include the name of the capture file (*.cap). For linux virtual machine it must start with `/var/captures`.
- `storage_account_id` - (Optional) The ID of the storage account to save the packet capture session

NOTE: At least one of `file_path` or `storage_account_id` must be specified.

A `filter` block contains:

- `local_ip_address` - (Optional) The local IP Address to be filtered on. Notation: "127.0.0.1" for single address entry. "127.0.0.1-127.0.0.255" for range. "127.0.0.1;127.0.0.5" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.
- `local_port` - (Optional) The local port to be filtered on. Notation: "80" for single port entry. "80-85" for range. "80;443;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.
- `protocol` - (Required) The Protocol to be filtered on. Possible values include `Any`, `TCP` and `UDP`. Changing this forces a new resource to be created.

- `remote_ip_address` - (Optional) The remote IP Address to be filtered on. Notation: "127.0.0.1" for single address entry. "127.0.0.1-127.0.0.255" for range. "127.0.0.1;127.0.0.5;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported.. Changing this forces a new resource to be created.
- `remote_port` - (Optional) The remote port to be filtered on. Notation: "80" for single port entry."80-85" for range. "80;443;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Packet Capture ID.
- `storage_location` - (Required) A `storage_location` block as defined below.

A `storage_location` block contains:

- `storage_path` - The URI of the storage path to save the packet capture.

Import

Packet Captures can be imported using the `resource id` , e.g.

```
terraform import azurerm_network_packet_capture.capture1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/networkWatchers/watcher1/packetCaptures/capture1
```

azurerm_network_profile

Manages an Azure Network Profile.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "examplegroup"
  location = "West Europe"
}

resource "azurerm_virtual_network" "example" {
  name            = "examplevnet"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  address_space   = ["10.1.0.0/16"]
}

resource "azurerm_subnet" "example" {
  name                 = "examplesubnet"
  resource_group_name = "${azurerm_resource_group.example.name}"
  virtual_network_name = "${azurerm_virtual_network.example.name}"
  address_prefix       = "10.1.0.0/24"

  delegation {
    name = "delegation"

    service_delegation {
      name    = "Microsoft.ContainerInstance/containerGroups"
      actions = ["Microsoft.Network/virtualNetworks/subnets/action"]
    }
  }
}

resource "azurerm_network_profile" "example" {
  name            = "examplenetworkprofile"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  container_network_interface {
    name = "examplecnic"

    ip_configuration {
      name          = "exampleipconfig"
      subnet_id     = "${azurerm_subnet.example.id}"
    }
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Network Profile. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the resource. Changing this forces a new resource to be created.
- `container_network_interface` - (Required) A `container_network_interface` block as documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `container_network_interface` block supports the following:

- `name` - (Required) Specifies the name of the IP Configuration.
- `ip_configuration` - (Required) One or more `ip_configuration` blocks as documented below.

A `ip_configuration` block supports the following:

- `name` - (Required) Specifies the name of the IP Configuration.
- `subnet_id` - (Required) Reference to the subnet associated with the IP Configuration.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Azure Network Profile.
- `container_network_interface_ids` - One or more Resource IDs of Azure Container Network Interfaces.

Import

Azure Network Profile can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_profile.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/networkProfiles/examplenetworkprofile
```

azurerm_network_security_group

Manages a network security group that contains a list of network security rules. Network security groups enable inbound or outbound traffic to be enabled or denied.

NOTE on Network Security Groups and Network Security Rules: Terraform currently provides both a standalone Network Security Rule resource ([/docs/providers/azurerm/r/network_security_rule.html](/docs/providers/azurerm/r/network_security_rule.html)), and allows for Network Security Rules to be defined in-line within the Network Security Group resource ([/docs/providers/azurerm/r/network_security_group.html](/docs/providers/azurerm/r/network_security_group.html)). At this time you cannot use a Network Security Group with in-line Network Security Rules in conjunction with any Network Security Rule resources. Doing so will cause a conflict of rule settings and will overwrite rules.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_network_security_group" "test" {
  name                = "acceptanceTestSecurityGroup1"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  security_rule {
    name                = "test123"
    priority             = 100
    direction           = "Inbound"
    access              = "Allow"
    protocol             = "Tcp"
    source_port_range    = "*"
    destination_port_range = "*"
    source_address_prefix = "*"
    destination_address_prefix = "*"
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the network security group. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the network security group. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `security_rule` - (Optional) List of objects (</docs/configuration/attr-as-blocks.html>) representing security rules, as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Elements of `security_rule` support:

- `name` - (Required) The name of the security rule.
- `description` - (Optional) A description for this rule. Restricted to 140 characters.
- `protocol` - (Required) Network protocol this rule applies to. Can be `Tcp`, `Udp`, `Icmp`, or `*` to match all.
- `source_port_range` - (Optional) Source Port or Range. Integer or range between `0` and `65535` or `*` to match any. This is required if `source_port_ranges` is not specified.
- `source_port_ranges` - (Optional) List of source ports or port ranges. This is required if `source_port_range` is not specified.
- `destination_port_range` - (Optional) Destination Port or Range. Integer or range between `0` and `65535` or `*` to match any. This is required if `destination_port_ranges` is not specified.
- `destination_port_ranges` - (Optional) List of destination ports or port ranges. This is required if `destination_port_range` is not specified.
- `source_address_prefix` - (Optional) CIDR or source IP range or `*` to match any IP. Tags such as 'VirtualNetwork', 'AzureLoadBalancer' and 'Internet' can also be used. This is required if `source_address_prefixes` is not specified.
- `source_address_prefixes` - (Optional) List of source address prefixes. Tags may not be used. This is required if `source_address_prefix` is not specified.
- `source_application_security_group_ids` - (Optional) A List of source Application Security Group ID's
- `destination_address_prefix` - (Optional) CIDR or destination IP range or `*` to match any IP. Tags such as 'VirtualNetwork', 'AzureLoadBalancer' and 'Internet' can also be used. This is required if `destination_address_prefixes` is not specified.
- `destination_address_prefixes` - (Optional) List of destination address prefixes. Tags may not be used. This is required if `destination_address_prefix` is not specified.
- `destination_application_security_group_ids` - (Optional) A List of destination Application Security Group ID's
- `access` - (Required) Specifies whether network traffic is allowed or denied. Possible values are `Allow` and `Deny`.
- `priority` - (Required) Specifies the priority of the rule. The value can be between `100` and `4096`. The priority number must be unique for each rule in the collection. The lower the priority number, the higher the priority of the rule.
- `direction` - (Required) The direction specifies if rule will be evaluated on incoming or outgoing traffic. Possible values are `Inbound` and `Outbound`.

Attributes Reference

The following attributes are exported:

- `id` - The Network Security Group ID.

Import

Network Security Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_security_group.group1 /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/mygroup1/providers/Microsoft.Network/networkSecurityGroups/mySecurityGroup
```

azurerm_network_security_rule

Manages a Network Security Rule.

NOTE on Network Security Groups and Network Security Rules: Terraform currently provides both a standalone Network Security Rule resource ([/docs/providers/azurerm/r/network_security_rule.html](/docs/providers/azurerm/r/network_security_rule.html)), and allows for Network Security Rules to be defined in-line within the Network Security Group resource ([/docs/providers/azurerm/r/network_security_group.html](/docs/providers/azurerm/r/network_security_group.html)). At this time you cannot use a Network Security Group with in-line Network Security Rules in conjunction with any Network Security Rule resources. Doing so will cause a conflict of rule settings and will overwrite rules.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_network_security_group" "test" {
  name            = "acceptanceTestSecurityGroup1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_network_security_rule" "test" {
  name                = "test123"
  priority            = 100
  direction           = "Outbound"
  access              = "Allow"
  protocol            = "Tcp"
  source_port_range   = "*"
  destination_port_range = "*"
  source_address_prefix = "*"
  destination_address_prefix = "*"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  network_security_group_name = "${azurerm_network_security_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the security rule. This needs to be unique across all Rules in the Network Security Group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Network Security Rule. Changing this forces a new resource to be created.

- `network_security_group_name` - (Required) The name of the Network Security Group that we want to attach the rule to. Changing this forces a new resource to be created.
- `description` - (Optional) A description for this rule. Restricted to 140 characters.
- `protocol` - (Required) Network protocol this rule applies to. Possible values include `Tcp`, `Udp`, `Icmp`, or `*` (which matches all).
- `source_port_range` - (Optional) Source Port or Range. Integer or range between `0` and `65535` or `*` to match any. This is required if `source_port_ranges` is not specified.
- `source_port_ranges` - (Optional) List of source ports or port ranges. This is required if `source_port_range` is not specified.
- `destination_port_range` - (Optional) Destination Port or Range. Integer or range between `0` and `65535` or `*` to match any. This is required if `destination_port_ranges` is not specified.
- `destination_port_ranges` - (Optional) List of destination ports or port ranges. This is required if `destination_port_range` is not specified.
- `source_address_prefix` - (Optional) CIDR or source IP range or `*` to match any IP. Tags such as 'VirtualNetwork', 'AzureLoadBalancer' and 'Internet' can also be used. This is required if `source_address_prefixes` is not specified.
- `source_address_prefixes` - (Optional) List of source address prefixes. Tags may not be used. This is required if `source_address_prefix` is not specified.
- `source_application_security_group_ids` - (Optional) A List of source Application Security Group ID's
- `destination_address_prefix` - (Optional) CIDR or destination IP range or `*` to match any IP. Tags such as 'VirtualNetwork', 'AzureLoadBalancer' and 'Internet' can also be used. This is required if `destination_address_prefixes` is not specified.
- `destination_address_prefixes` - (Optional) List of destination address prefixes. Tags may not be used. This is required if `destination_address_prefix` is not specified.
- `destination_application_security_group_ids` - (Optional) A List of destination Application Security Group ID's
- `access` - (Required) Specifies whether network traffic is allowed or denied. Possible values are `Allow` and `Deny`.
- `priority` - (Required) Specifies the priority of the rule. The value can be between 100 and 4096. The priority number must be unique for each rule in the collection. The lower the priority number, the higher the priority of the rule.
- `direction` - (Required) The direction specifies if rule will be evaluated on incoming or outgoing traffic. Possible values are `Inbound` and `Outbound`.

Attributes Reference

The following attributes are exported:

- `id` - The Network Security Rule ID.

Import

Network Security Rules can be imported using the `resource_id`, e.g.

```
terraform import azurerm_network_security_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.Network/networkSecurityGroups/mySecurityGroup/securityRules/r  
ule1
```

azurerm_network_watcher

Manages a Network Watcher.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "production-nwwatcher"
  location  = "West US"
}

resource "azurerm_network_watcher" "test" {
  name                = "production-nwwatcher"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Network Watcher. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Network Watcher. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Network Watcher ID.

Import

Network Watchers can be imported using the `resource id`, e.g.

```
terraform import azurerm_network_watcher.watcher1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/networkWatchers/watcher1
```


azurerm_notification_hub_authorization_rule

Manages an Authorization Rule associated with a Notification Hub within a Notification Hub Namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "notificationhub-resources"
  location  = "Australia East"
}

resource "azurerm_notification_hub_namespace" "test" {
  name                = "myappnamespace"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  namespace_type      = "NotificationHub"

  sku_name = "Free"
}

resource "azurerm_notification_hub" "test" {
  name                = "mynotificationhub"
  namespace_name      = "${azurerm_notification_hub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
}

resource "azurerm_notification_hub_authorization_rule" "test" {
  name                = "management-auth-rule"
  notification_hub_name = "${azurerm_notification_hub.test.name}"
  namespace_name      = "${azurerm_notification_hub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  manage              = true
  send                = true
  listen              = true
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name to use for this Authorization Rule. Changing this forces a new resource to be created.
- `notification_hub_name` - (Required) The name of the Notification Hub for which the Authorization Rule should be created. Changing this forces a new resource to be created.
- `namespace_name` - (Required) The name of the Notification Hub Namespace in which the Notification Hub exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Notification Hub Namespace exists. Changing this forces a new resource to be created.

- `manage` - (Optional) Does this Authorization Rule have Manage access to the Notification Hub? Defaults to `false`.

NOTE: If `manage` is set to `true` then both `send` and `listen` must also be set to `true`.

- `send` - (Optional) Does this Authorization Rule have Send access to the Notification Hub? Defaults to `false`.
- `listen` - (Optional) Does this Authorization Rule have Listen access to the Notification Hub? Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Authorization Rule.
- `primary_access_key` - The Primary Access Key associated with this Authorization Rule.
- `secondary_access_key` - The Secondary Access Key associated with this Authorization Rule.

Import

Notification Hub Authorization Rule can be imported using the `resource id`, e.g.

```
terraform import azurerm_notification_hub_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.NotificationHubs/namespaces/{namespaceName}/notificationHubs/hub1/AuthorizationRules/rule1
```

azurerm_notification_hub

Manages a Notification Hub within a Notification Hub Namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "notificationhub-resources"
  location  = "Australia East"
}

resource "azurerm_notification_hub_namespace" "test" {
  name                  = "myappnamespace"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location              = "${azurerm_resource_group.test.location}"
  namespace_type       = "NotificationHub"

  sku_name = "Free"
}

resource "azurerm_notification_hub" "test" {
  name                  = "mynotificationhub"
  namespace_name       = "${azurerm_notification_hub_namespace.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location              = "${azurerm_resource_group.test.location}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name to use for this Notification Hub. Changing this forces a new resource to be created.
- `namespace_name` - (Required) The name of the Notification Hub Namespace in which to create this Notification Hub. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Notification Hub Namespace exists. Changing this forces a new resource to be created.
- `location` - (Required) The Azure Region in which this Notification Hub Namespace exists. Changing this forces a new resource to be created.
- `apns_credential` - (Optional) A `apns_credential` block as defined below.

NOTE: Removing the `apns_credential` block will currently force a recreation of this resource due to this bug in the Azure SDK for Go (<https://github.com/Azure/azure-sdk-for-go/issues/2246>) - we'll remove this limitation when the SDK bug is fixed.

- `gcm_credential` - (Optional) A `gcm_credential` block as defined below.

NOTE: Removing the `gcm_credential` block will currently force a recreation of this resource due to this bug in the Azure SDK for Go (<https://github.com/Azure/azure-sdk-for-go/issues/2246>) - we'll remove this limitation when the SDK bug is fixed.

A `apns_credential` block contains:

- `application_mode` - (Required) The Application Mode which defines which server the APNS Messages should be sent to. Possible values are `Production` and `Sandbox`.
- `bundle_id` - (Required) The Bundle ID of the iOS/macOS application to send push notifications for, such as `com.hashicorp.example`.
- `key_id` - (Required) The Apple Push Notifications Service (APNS) Key.
- `team_id` - (Required) The ID of the team the Token.
- `token` - (Required) The Push Token associated with the Apple Developer Account. This is the contents of the `key` downloaded from the Apple Developer Portal (<https://developer.apple.com/account/ios/authkey/>) between the `-----BEGIN PRIVATE KEY-----` and `-----END PRIVATE KEY-----` blocks.

A `gcm_credential` block contains:

- `api_key` - (Required) The API Key associated with the Google Cloud Messaging service.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Notification Hub.

Import

Notification Hubs can be imported using the `resource id`, e.g.

```
terraform import azurerm_notification_hub.hub1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.NotificationHubs/namespaces/{namespaceName}/notificationHubs/hub1
```

azurerm_notification_hub_namespace

Manages a Notification Hub Namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "notificationhub-resources"
  location  = "Australia East"
}

resource "azurerm_notification_hub_namespace" "test" {
  name                        = "myappnamespace"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  namespace_type             = "NotificationHub"

  sku_name = "Free"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name to use for this Notification Hub Namespace. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Notification Hub Namespace should exist. Changing this forces a new resource to be created.
- `location` - (Required) The Azure Region in which this Notification Hub Namespace should be created.
- `namespace_type` - (Required) The Type of Namespace - possible values are `Messaging` or `NotificationHub`. Changing this forces a new resource to be created.
- `sku` - (Optional **Deprecated**) A `sku` block as described below.
- `sku_name` - (Optional) The name of the SKU to use for this Notification Hub Namespace. Possible values are `Free`, `Basic` or `Standard`. Changing this forces a new resource to be created.
- `enabled` - (Optional) Is this Notification Hub Namespace enabled? Defaults to `true`.

A `sku` block contains:

- `name` - (Required) The name of the SKU to use for this Notification Hub Namespace. Possible values are `Free`, `Basic` or `Standard`. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Notification Hub Namespace.
- `servicebus_endpoint` - The ServiceBus Endpoint for this Notification Hub Namespace.

Import

Notification Hub Namespaces can be imported using the `resource id`, e.g.

```
terraform import azurerm_notification_hub_namespace.namespace1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.NotificationHubs/namespaces/{namespaceName}
```

azurerm_packet_capture

Configures Packet Capturing against a Virtual Machine using a Network Watcher.

NOTE: This resource has been deprecated in favour of the `azurerm_network_connection_monitor` resource and will be removed in the next major version of the AzureRM Provider. The new resource shares the same fields as this one, and information on migrating across can be found in this guide (</docs/providers/azurerm/guides/migrating-between-renamed-resources.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "packet-capture-rg"
  location  = "West Europe"
}

resource "azurerm_network_watcher" "test" {
  name                = "network-watcher"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_virtual_network" "test" {
  name                = "production-network"
  address_space       = ["10.0.0.0/16"]
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "internal"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "pctest-nic"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                = "testconfiguration1"
    subnet_id           = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_virtual_machine" "test" {
  name                = "pctest-vm"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_interface_ids = ["${azurerm_network_interface.test.id}"]
}
```

```

vm_size                = "Standard_F2"

storage_image_reference {
  publisher = "Canonical"
  offer     = "UbuntuServer"
  sku       = "16.04-LTS"
  version   = "latest"
}

storage_os_disk {
  name          = "osdisk"
  caching       = "ReadWrite"
  create_option = "FromImage"
  managed_disk_type = "Standard_LRS"
}

os_profile {
  computer_name  = "pctest-vm"
  admin_username = "testadmin"
  admin_password = "Password1234!"
}

os_profile_linux_config {
  disable_password_authentication = false
}
}

resource "azurerm_virtual_machine_extension" "test" {
  name                = "network-watcher"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_machine_name = "${azurerm_virtual_machine.test.name}"
  publisher           = "Microsoft.Azure.NetworkWatcher"
  type               = "NetworkWatcherAgentLinux"
  type_handler_version = "1.4"
  auto_upgrade_minor_version = true
}

resource "azurerm_storage_account" "test" {
  name                = "pctestsa"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  account_tier        = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_packet_capture" "test" {
  name                = "pctestcapture"
  network_watcher_name = "${azurerm_network_watcher.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  target_resource_id  = "${azurerm_virtual_machine.test.id}"

  storage_location {
    storage_account_id = "${azurerm_storage_account.test.id}"
  }

  depends_on = ["azurerm_virtual_machine_extension.test"]
}

```


NOTE: This Resource requires that the Network Watcher Virtual Machine Extension (<https://docs.microsoft.com/azure/network-watcher/network-watcher-packet-capture-manage-portal#before-you-begin>) is installed on the Virtual Machine before capturing can be enabled which can be installed via the `azurerm_virtual_machine_extension` resource ([/docs/providers/azurerm/r/virtual_machine_extension.html](https://docs/providers/azurerm/r/virtual_machine_extension.html)).

Argument Reference

The following arguments are supported:

- `name` - (Required) The name to use for this Packet Capture. Changing this forces a new resource to be created.
- `network_watcher_name` - (Required) The name of the Network Watcher. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Network Watcher exists. Changing this forces a new resource to be created.
- `target_resource_id` - (Required) The ID of the Resource to capture packets from. Changing this forces a new resource to be created.

NOTE: Currently only Virtual Machines ID's are supported.

- `maximum_bytes_per_packet` - (Optional) The number of bytes captured per packet. The remaining bytes are truncated. Defaults to `0` (Entire Packet Captured). Changing this forces a new resource to be created.
- `maximum_bytes_per_session` - (Optional) Maximum size of the capture in Bytes. Defaults to `1073741824` (1GB). Changing this forces a new resource to be created.
- `maximum_capture_duration` - (Optional) The maximum duration of the capture session in seconds. Defaults to `18000` (5 hours). Changing this forces a new resource to be created.
- `storage_location` - (Required) A `storage_location` block as defined below. Changing this forces a new resource to be created.
- `filter` - (Optional) One or more `filter` blocks as defined below. Changing this forces a new resource to be created.

A `storage_location` block contains:

- `file_path` - (Optional) A valid local path on the targeting VM. Must include the name of the capture file (*.cap). For linux virtual machine it must start with `/var/captures`.
- `storage_account_id` - (Optional) The ID of the storage account to save the packet capture session

NOTE: At least one of `file_path` or `storage_account_id` must be specified.

A `filter` block contains:

- `local_ip_address` - (Optional) The local IP Address to be filtered on. Notation: "127.0.0.1" for single address entry. "127.0.0.1-127.0.0.255" for range. "127.0.0.1;127.0.0.5" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.

- `local_port` - (Optional) The local port to be filtered on. Notation: "80" for single port entry."80-85" for range. "80;443;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.
- `protocol` - (Required) The Protocol to be filtered on. Possible values include `Any` , `TCP` and `UDP` . Changing this forces a new resource to be created.
- `remote_ip_address` - (Optional) The remote IP Address to be filtered on. Notation: "127.0.0.1" for single address entry. "127.0.0.1-127.0.0.255" for range. "127.0.0.1;127.0.0.5;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported.. Changing this forces a new resource to be created.
- `remote_port` - (Optional) The remote port to be filtered on. Notation: "80" for single port entry."80-85" for range. "80;443;" for multiple entries. Multiple ranges not currently supported. Mixing ranges with multiple entries not currently supported. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The Packet Capture ID.
- `storage_location` - (Required) A `storage_location` block as defined below.

A `storage_location` block contains:

- `storage_path` - The URI of the storage path to save the packet capture.

Import

Packet Captures can be imported using the `resource id` , e.g.

```
terraform import azurerm_packet_capture.capture1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/networkWatchers/watcher1/packetCaptures/capture1
```

azurerm_policy_assignment

Configures the specified Policy Definition at the specified Scope. Also, Policy Set Definitions are supported.

Example Usage

```

resource "azurerm_policy_definition" "test" {
  name           = "my-policy-definition"
  policy_type    = "Custom"
  mode           = "All"
  display_name   = "my-policy-definition"

  policy_rule = <<POLICY_RULE
    {
      "if": {
        "not": {
          "field": "location",
          "in": "[parameters('allowedLocations')]"
        }
      },
      "then": {
        "effect": "audit"
      }
    }
  }
POLICY_RULE

  parameters = <<PARAMETERS
    {
      "allowedLocations": {
        "type": "Array",
        "metadata": {
          "description": "The list of allowed locations for resources.",
          "displayName": "Allowed locations",
          "strongType": "location"
        }
      }
    }
  }
PARAMETERS
}

resource "azurerm_resource_group" "test" {
  name     = "test-resources"
  location = "West Europe"
}

resource "azurerm_policy_assignment" "test" {
  name                = "example-policy-assignment"
  scope               = "${azurerm_resource_group.test.id}"
  policy_definition_id = "${azurerm_policy_definition.test.id}"
  description         = "Policy Assignment created via an Acceptance Test"
  display_name        = "My Example Policy Assignment"

  parameters = <<PARAMETERS
    {
      "allowedLocations": {
        "value": [ "West Europe" ]
      }
    }
  }
PARAMETERS
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Policy Assignment. Changing this forces a new resource to be created.
- `scope` - (Required) The Scope at which the Policy Assignment should be applied, which must be a Resource ID (such as Subscription e.g. `/subscriptions/00000000-0000-0000-000000000000` or a Resource Group e.g. `/subscriptions/00000000-0000-0000-000000000000/resourceGroups/myResourceGroup`). Changing this forces a new resource to be created.
- `policy_definition_id` - (Required) The ID of the Policy Definition to be applied at the specified Scope.
- `identity` - (Optional) An `identity` block.
- `location` - (Optional) The Azure location where this policy assignment should exist. This is required when an Identity is assigned. Changing this forces a new resource to be created.
- `description` - (Optional) A description to use for this Policy Assignment. Changing this forces a new resource to be created.
- `display_name` - (Optional) A friendly display name to use for this Policy Assignment. Changing this forces a new resource to be created.
- `parameters` - (Optional) Parameters for the policy definition. This field is a JSON object that maps to the `Parameters` field from the Policy Definition. Changing this forces a new resource to be created.

NOTE: This value is required when the specified Policy Definition contains the `parameters` field.

- `not_scopes` - (Optional) A list of the Policy Assignment's excluded scopes. The list must contain Resource IDs (such as Subscriptions e.g. `/subscriptions/00000000-0000-0000-000000000000` or Resource Groups e.g. `/subscriptions/00000000-0000-0000-000000000000/resourceGroups/myResourceGroup`).

An `identity` block supports the following:

- `type` - (Required) The Managed Service Identity Type of this Policy Assignment. Possible values are `SystemAssigned` (where Azure will generate a Service Principal for you), or `None` (no use of a Managed Service Identity).

NOTE: When `type` is set to `SystemAssigned`, identity the Principal ID can be retrieved after the policy has been assigned.

Attributes Reference

The following attributes are exported:

- `id` - The Policy Assignment id.
 - `identity` - An `identity` block.
-

An `identity` block exports the following:

- `principal_id` - The Principal ID of this Policy Assignment if `type` is `SystemAssigned`.
 - `tenant_id` - The Tenant ID of this Policy Assignment if `type` is `SystemAssigned`.
-

Import

Policy Assignments can be imported using the `policy_name`, e.g.

```
terraform import azurerm_policy_assignment.assignment1 /subscriptions/00000000-0000-0000-000000000000/providers/Microsoft.Authorization/policyAssignments/assignment1
```

azurerm_policy_definition

Manages a policy rule definition on a management group or your provider subscription.

Policy definitions do not take effect until they are assigned to a scope using a Policy Assignment.

Example Usage

```
resource "azurerm_policy_definition" "policy" {
  name          = "accTestPolicy"
  policy_type   = "Custom"
  mode          = "Indexed"
  display_name  = "acceptance test policy definition"

  metadata = <<METADATA
    {
      "category": "General"
    }
  METADATA

  policy_rule = <<POLICY_RULE
    {
      "if": {
        "not": {
          "field": "location",
          "in": "[parameters('allowedLocations')]"
        }
      },
      "then": {
        "effect": "audit"
      }
    }
  POLICY_RULE

  parameters = <<PARAMETERS
    {
      "allowedLocations": {
        "type": "Array",
        "metadata": {
          "description": "The list of allowed locations for resources.",
          "displayName": "Allowed locations",
          "strongType": "location"
        }
      }
    }
  PARAMETERS
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the policy definition. Changing this forces a new resource to be created.
- `policy_type` - (Required) The policy type. The value can be "BuiltIn", "Custom" or "NotSpecified". Changing this forces a new resource to be created.
- `mode` - (Required) The policy mode that allows you to specify which resource types will be evaluated. The value can be "All", "Indexed" or "NotSpecified". Changing this resource forces a new resource to be created.
- `display_name` - (Required) The display name of the policy definition.
- `description` - (Optional) The description of the policy definition.
- `management_group_id` - (Optional) The ID of the Management Group where this policy should be defined. Changing this forces a new resource to be created.

Note: if you are using `azurerm_management_group` to assign a value to `management_group_id`, be sure to use `.group_id` and not `.id`.

- `policy_rule` - (Optional) The policy rule for the policy definition. This is a json object representing the rule that contains an if and a then block.
- `metadata` - (Optional) The metadata for the policy definition. This is a json object representing additional metadata that should be stored with the policy definition.
- `parameters` - (Optional) Parameters for the policy definition. This field is a json object that allows you to parameterize your policy definition.

Attributes Reference

The following attributes are exported:

- `id` - The policy definition id.

Import

Policy Definitions can be imported using the `policy_name`, e.g.

```
terraform import azurerm_policy_definition.testPolicy /subscriptions/<SUBSCRIPTION_ID>/providers/Microsoft.Authorization/policyDefinitions/<POLICY_NAME>
```

or

```
terraform import azurerm_policy_definition.testPolicy /providers/Microsoft.Management/managementgroups/<MANAGEMENT_GROUP_ID>/providers/Microsoft.Authorization/policyDefinitions/<POLICY_NAME>
```


azurerm_policy_set_definition

Manages a policy set definition.

NOTE: Policy set definitions (also known as policy initiatives) do not take effect until they are assigned to a scope using a Policy Set Assignment.

Example Usage

```
resource "azurerm_policy_set_definition" "test" {
  name          = "testPolicySet"
  policy_type   = "Custom"
  display_name  = "Test Policy Set"

  parameters = <<PARAMETERS
    {
      "allowedLocations": {
        "type": "Array",
        "metadata": {
          "description": "The list of allowed locations for resources.",
          "displayName": "Allowed locations",
          "strongType": "location"
        }
      }
    }
  PARAMETERS

  policy_definitions = <<POLICY_DEFINITIONS
    [
      {
        "parameters": {
          "listOfAllowedLocations": {
            "value": "[parameters('allowedLocations')]"
          }
        },
        "policyDefinitionId": "/providers/Microsoft.Authorization/policyDefinitions/e765b5de-1225-4ba3-bd56-1ac6695af988"
      }
    ]
  POLICY_DEFINITIONS
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the policy set definition. Changing this forces a new resource to be created.
- `policy_type` - (Required) The policy set type. Possible values are `BuiltIn` or `Custom`. Changing this forces a new resource to be created.

- `display_name` - (Required) The display name of the policy set definition.
- `policy_definitions` - (Required) The policy definitions for the policy set definition. This is a json object representing the bundled policy definitions .
- `description` - (Optional) The description of the policy set definition.
- `management_group_id` - (Optional) The ID of the Management Group where this policy should be defined. Changing this forces a new resource to be created.

Note: if you are using `azurerm_management_group` to assign a value to `management_group_id` , be sure to use `.group_id` and not `.id` .

- `metadata` - (Optional) The metadata for the policy set definition. This is a json object representing additional metadata that should be stored with the policy definition.
- `parameters` - (Optional) Parameters for the policy set definition. This field is a json object that allows you to parameterize your policy definition.

Attributes Reference

The following attributes are exported:

- `id` - The policy set definition id.

Import

Policy Set Definitions can be imported using the Resource ID, e.g.

```
terraform import azurerm_policy_set_definition.test /subscriptions/00000000-0000-0000-0000-000000000000/providers/Microsoft.Authorization/policySetDefinitions/testPolicySet
```

```
or shell terraform import azurerm_policy_set_definition.test
/providers/Microsoft.Management/managementgroups/my-mgmt-group-
id/providers/Microsoft.Authorization/policySetDefinitions/testPolicySet
```

azurerm_postgresql_configuration

Sets a PostgreSQL Configuration value on a PostgreSQL Server.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_postgresql_server" "test" {
  name                = "postgresql-server-1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name      = "B_Gen5_2"
    capacity  = 2
    tier       = "Basic"
    family    = "Gen4"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "psqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "9.5"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_postgresql_configuration" "test" {
  name                = "backslash_quote"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name         = "${azurerm_postgresql_server.test.name}"
  value               = "on"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the PostgreSQL Configuration, which needs to be a valid PostgreSQL configuration name (<https://www.postgresql.org/docs/current/static/sql-syntax-lexical.html#SQL-SYNTAX-IDENTIFIER>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the PostgreSQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the PostgreSQL Server exists. Changing this forces a new resource to be created.
- `value` - (Required) Specifies the value of the PostgreSQL Configuration. See the PostgreSQL documentation for valid values.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the PostgreSQL Configuration.

Import

PostgreSQL Configurations can be imported using the `resource id`, e.g.

```
terraform import azurerm_postgresql_configuration.backslash_quote /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforPostgreSQL/servers/server1/configurations/backslash_quote
```

azurerm_postgresql_database

Manages a PostgreSQL Database within a PostgreSQL Server

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "api-rg-pro"
  location = "West Europe"
}

resource "azurerm_postgresql_server" "test" {
  name                = "postgresql-server-1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name       = "B_Gen5_2"
    capacity   = 2
    tier        = "Basic"
    family     = "Gen5"
  }

  storage_profile {
    storage_mb           = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
  }

  administrator_login          = "psqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "9.5"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_postgresql_database" "test" {
  name                = "exampleldb"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name         = "${azurerm_postgresql_server.test.name}"
  charset             = "UTF8"
  collation            = "English_United States.1252"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the PostgreSQL Database, which needs to be a valid PostgreSQL identifier (<https://www.postgresql.org/docs/current/static/sql-syntax-lexical.html#SQL-SYNTAX-IDENTIFIERS>). Changing this forces a new resource to be created.

- `server_name` - (Required) Specifies the name of the PostgreSQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the PostgreSQL Server exists. Changing this forces a new resource to be created.
- `charset` - (Required) Specifies the Charset for the PostgreSQL Database, which needs to be a valid PostgreSQL Charset (<https://www.postgresql.org/docs/current/static/multibyte.html>). Changing this forces a new resource to be created.
- `collation` - (Required) Specifies the Collation for the PostgreSQL Database, which needs to be a valid PostgreSQL Collation (<https://www.postgresql.org/docs/current/static/collation.html>). Note that Microsoft uses different notation (<https://msdn.microsoft.com/library/windows/desktop/dd373814.aspx>) - en-US instead of en_US. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the PostgreSQL Database.

Import

PostgreSQL Database's can be imported using the `resource id`, e.g.

```
terraform import azurerm_postgresql_database.database1 /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/mygroup1/providers/Microsoft.DBforPostgreSQL/servers/server1/databases/database1
```

azurerm_postgresql_firewall_rule

Manages a Firewall Rule for a PostgreSQL Server

Example Usage (Single IP Address)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_postgresql_server" "test" {
  # ...
}

resource "azurerm_postgresql_firewall_rule" "test" {
  name                        = "office"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  server_name                 = "${azurerm_postgresql_server.test.name}"
  start_ip_address            = "40.112.8.12"
  end_ip_address              = "40.112.8.12"
}
```

Example Usage (IP Range)

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_postgresql_server" "test" {
  # ...
}

resource "azurerm_postgresql_firewall_rule" "test" {
  name                        = "office"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  server_name                 = "${azurerm_postgresql_server.test.name}"
  start_ip_address            = "40.112.0.0"
  end_ip_address              = "40.112.255.255"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the PostgreSQL Firewall Rule. Changing this forces a new resource to be created.
- `server_name` - (Required) Specifies the name of the PostgreSQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the PostgreSQL Server exists. Changing this forces a new resource to be created.
- `start_ip_address` - (Required) Specifies the Start IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.
- `end_ip_address` - (Required) Specifies the End IP Address associated with this Firewall Rule. Changing this forces a new resource to be created.

NOTE: The Azure feature `Allow access to Azure services` can be enabled by setting `start_ip_address` and `end_ip_address` to `0.0.0.0` which (is documented in the Azure API Docs (<https://docs.microsoft.com/en-us/rest/api/sql/firewallrules/createorupdate>)).

Attributes Reference

The following attributes are exported:

- `id` - The ID of the PostgreSQL Firewall Rule.

Import

PostgreSQL Firewall Rule's can be imported using the `resource id`, e.g.

```
terraform import azurerm_postgresql_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/00/resourceGroups/mygroup1/providers/Microsoft.DBforPostgreSQL/servers/server1/firewallRules/rule1
```


azurerm_postgresql_server

Manages a PostgreSQL Server.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "api-rg-pro"
  location  = "West Europe"
}

resource "azurerm_postgresql_server" "test" {
  name                  = "postgresql-server-1"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name      = "B_Gen5_2"
    capacity = 2
    tier       = "Basic"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb          = 5120
    backup_retention_days = 7
    geo_redundant_backup = "Disabled"
    auto_grow           = "Enabled"
  }

  administrator_login          = "psqladminun"
  administrator_login_password = "H@Sh1CoR3!"
  version                     = "9.5"
  ssl_enforcement              = "Enabled"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the PostgreSQL Server. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the PostgreSQL Server. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) A `sku` block as defined below.
- `storage_profile` - (Required) A `storage_profile` block as defined below.

- `administrator_login` - (Required) The Administrator Login for the PostgreSQL Server. Changing this forces a new resource to be created.
 - `administrator_login_password` - (Required) The Password associated with the `administrator_login` for the PostgreSQL Server.
 - `version` - (Required) Specifies the version of PostgreSQL to use. Valid values are `9.5`, `9.6`, `10`, `10.0`, and `11`. Changing this forces a new resource to be created.
 - `ssl_enforcement` - (Required) Specifies if SSL should be enforced on connections. Possible values are `Enabled` and `Disabled`.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

`sku` supports the following:

- `name` - (Required) Specifies the SKU Name for this PostgreSQL Server. The name of the SKU, follows the `tier + family + cores` pattern (e.g. `B_Gen4_1`, `GP_Gen5_8`). For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/postgresql/servers/create#sku>).
 - `capacity` - (Required) The scale up/out capacity, representing server's compute units.
 - `tier` - (Required) The tier of the particular SKU. Possible values are `Basic`, `GeneralPurpose`, and `MemoryOptimized`. For more information see the product documentation (<https://docs.microsoft.com/en-us/azure/postgresql/concepts-pricing-tiers>).
 - `family` - (Required) The family of hardware `Gen4` or `Gen5`, before selecting your family check the product documentation (<https://docs.microsoft.com/en-us/azure/postgresql/concepts-pricing-tiers#compute-generations-vcores-and-memory>) for availability in your region.
-

`storage_profile` supports the following:

- `storage_mb` - (Required) Max storage allowed for a server. Possible values are between `5120 MB(5GB)` and `1048576 MB(1TB)` for the Basic SKU and between `5120 MB(5GB)` and `4194304 MB(4TB)` for General Purpose/Memory Optimized SKUs. For more information see the product documentation (<https://docs.microsoft.com/en-us/rest/api/postgresql/servers/create#StorageProfile>).
- `backup_retention_days` - (Optional) Backup retention days for the server, supported values are between `7` and `35` days.
- `geo_redundant_backup` - (Optional) Enable/Disable Geo-redundant for server backup. Valid values for this property are `Enabled` or `Disabled`, not supported for the `basic` tier. This allows you to choose between locally redundant or geo-redundant backup storage in the General Purpose and Memory Optimized tiers. When the backups are stored in geo-redundant backup storage, they are not only stored within the region in which your server is hosted, but are also replicated to a paired data center. This provides better protection and ability to restore your server in a different region in the event of a disaster. The Basic tier only offers locally redundant backup storage.
- `auto_grow` - (Optional) Enable/Disable auto-growing of the storage. Valid values for this property are `Enabled` or `Disabled`. Storage auto-grow prevents your server from running out of storage and becoming read-only. If storage auto grow is enabled, the storage automatically grows without impacting the workload. The default value if not explicitly specified is `Enabled`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the PostgreSQL Server.
- `fqdn` - The FQDN of the PostgreSQL Server.

Import

PostgreSQL Server's can be imported using the `resource id`, e.g.

```
terraform import azurerm_postgresql_server.server1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.DBforPostgreSQL/servers/server1
```

azurerm_postgresql_virtual_network_rule

Manages a PostgreSQL Virtual Network Rule.

NOTE: PostgreSQL Virtual Network Rules can only be used with SKU Tiers of `GeneralPurpose` or `MemoryOptimized` (<https://docs.microsoft.com/en-us/azure/postgresql/concepts-data-access-and-security-vnet>)

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West US"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-vnet"
  address_space   = ["10.7.29.0/29"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "internal" {
  name                 = "internal"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.7.29.0/29"
  service_endpoints    = ["Microsoft.Sql"]
}

resource "azurerm_postgresql_server" "test" {
  name            = "postgresql-server-1"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku {
    name      = "GP_Gen5_2"
    capacity = 2
    tier       = "GeneralPurpose"
    family    = "Gen5"
  }

  storage_profile {
    storage_mb            = 5120
    backup_retention_days = 7
    geo_redundant_backup  = "Disabled"
  }

  administrator_login          = "psqladminun"
  administrator_login_password = "HqSh1CoR3!"
  version                     = "9.5"
  ssl_enforcement              = "Enabled"
}

resource "azurerm_postgresql_virtual_network_rule" "test" {
  name                        = "postgresql-vnet-rule"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  server_name                 = "${azurerm_postgresql_server.test.name}"
  subnet_id                   = "${azurerm_subnet.internal.id}"
  ignore_missing_vnet_service_endpoint = true
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the PostgreSQL virtual network rule. Cannot be empty and must only contain alphanumeric characters and hyphens. Cannot start with a number, and cannot start or end with a hyphen. Changing this forces a new resource to be created.

NOTE: **name** must be between 1-128 characters long and must satisfy all of the requirements below: 1. Contains only alphanumeric and hyphen characters 2. Cannot start with a number or hyphen 3. Cannot end with a hyphen

- **resource_group_name** - (Required) The name of the resource group where the PostgreSQL server resides. Changing this forces a new resource to be created.
- **server_name** - (Required) The name of the SQL Server to which this PostgreSQL virtual network rule will be applied to. Changing this forces a new resource to be created.
- **subnet_id** - (Required) The ID of the subnet that the PostgreSQL server will be connected to.
- **ignore_missing_vnet_service_endpoint** - (Optional) Should the Virtual Network Rule be created before the Subnet has the Virtual Network Service Endpoint enabled? Defaults to `false`.

Attributes Reference

The following attributes are exported:

- **id** - The ID of the PostgreSQL Virtual Network Rule.

Import

PostgreSQL Virtual Network Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_postgresql_virtual_network_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.DBforPostgreSQL/servers/myserver/virtualNetworkRules/vnetrulename
```

azurerm_private_dns_a_record

Enables you to manage DNS A Records within Azure Private DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_private_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_private_dns_a_record" "test" {
  name           = "test"
  zone_name      = "${azurerm_private_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl            = 300
  records        = ["10.0.180.17"]
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the DNS A Record.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `zone_name` - (Required) Specifies the Private DNS Zone where the resource exists. Changing this forces a new resource to be created.
- `TTL` - (Required) The Time To Live (TTL) of the DNS record in seconds.
- `records` - (Required) List of IPv4 Addresses.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Private DNS A Record ID.

Import

Private DNS A Records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_private_dns_a_record.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/privateDnsZones/zone1/A/myrecord1
```


azurerm_private_dns_cname_record

Enables you to manage DNS CNAME Records within Azure Private DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_private_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_private_dns_cname_record" "test" {
  name           = "test"
  zone_name      = "${azurerm_private_dns_zone.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  ttl            = 300
  record         = "contoso.com"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the DNS CNAME Record.
- **resource_group_name** - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- **zone_name** - (Required) Specifies the Private DNS Zone where the resource exists. Changing this forces a new resource to be created.
- **TTL** - (Required) The Time To Live (TTL) of the DNS record in seconds.
- **record** - (Required) The target of the CNAME.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- **id** - The Private DNS CNAME Record ID.

Import

Private DNS CName Records can be imported using the `resource_id`, e.g.

```
terraform import azurerm_private_dns_cname_record.test /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/mygroup1/providers/Microsoft.Network/privateDnsZones/zone1/CName/myrecord1
```

azurerm_private_dns_zone

Enables you to manage Private DNS zones within Azure DNS. These zones are hosted on Azure's name servers.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_private_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Private DNS Zone. Must be a valid domain name.
- `resource_group_name` - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Private DNS Zone ID.
- `number_of_record_sets` - The current number of record sets in this Private DNS zone.
- `max_number_of_record_sets` - The maximum number of record sets that can be created in this Private DNS zone.
- `max_number_of_virtual_network_links` - The maximum number of virtual networks that can be linked to this Private DNS zone.
- `max_number_of_virtual_network_links_with_registration` - The maximum number of virtual networks that can be linked to this Private DNS zone with registration enabled.

Import

Private DNS Zones can be imported using the `resource id`, e.g.

```
terraform import azurerm_private_dns_zone.zone1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/privateDnsZones/zone1
```

azurerm_private_dns_zone_virtual_network_link

Enables you to manage Private DNS zone Virtual Network Links. These Links enable DNS resolution and registration inside Azure Virtual Networks using Azure Private DNS.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_private_dns_zone" "test" {
  name                = "mydomain.com"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_private_dns_zone_virtual_network_link" "test" {
  name                = "test"
  resource_group_name = azurerm_resource_group.test.name
  private_dns_zone_name = azurerm_private_dns_zone.test.name
  virtual_network_id   = azurerm_virtual_network.test.id
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Private DNS Zone Virtual Network Link. Changing this forces a new resource to be created.
- **private_dns_zone_name** - (Required) The name of the Private DNS zone (without a terminating dot). Changing this forces a new resource to be created.
- **resource_group_name** - (Required) Specifies the resource group where the resource exists. Changing this forces a new resource to be created.
- **virtual_network_id** - (Required) The Resource ID of the Virtual Network that should be linked to the DNS Zone. Changing this forces a new resource to be created.
- **registration_enabled** - (Optional) Is auto-registration of virtual machine records in the virtual network in the Private DNS zone enabled? Defaults to `false`.
- **tags** - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Resource ID of the Private DNS Zone Virtual Network Link.

Import

Private DNS Zone Virtual Network Links can be imported using the `resource id`, e.g.

```
terraform import azurerm_private_dns_zone_virtual_network_link.link1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/privateDnsZones/zone1.com/virtualNetworkLinks/myVnetLink1
```

azurerm_proximity_placement_group

Manages a proximity placement group for virtual machines, virtual machine scale sets and availability sets.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "resourceGroup1"
  location  = "West US"
}

resource "azurerm_proximity_placement_group" "example" {
  name                  = "exampleProximityPlacementGroup"
  location              = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the availability set. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the availability set. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Proximity Placement Group ID.
- `name` - The name of the Proximity Placement Group.
- `location` - The location of the Proximity Placement Group.
- `resource_group_name` - The name of the resource group in which the Proximity Placement Group exists.
- `tags` - The tags attached to the Proximity Placement Group.

Import

Proximity Placement Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_proximity_placement_group.example /subscriptions/00000000-0000-0000-0000-00000000  
00000/resourceGroups/example-rg/providers/Microsoft.Compute/proximityPlacementGroups/example-ppg
```


azurerm_public_ip

Manages a Public IP Address.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "resourceGroup1"
  location = "West US"
}

resource "azurerm_public_ip" "test" {
  name                = "acceptanceTestPublicIp1"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Public IP resource . Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the public ip.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Optional) The SKU of the Public IP. Accepted values are `Basic` and `Standard` . Defaults to `Basic` .

Note Public IP Standard SKUs require `allocation_method` to be set to `Static` .

- `allocation_method` - (Required) Defines the allocation method for this IP address. Possible values are `Static` or `Dynamic` .

Note `Dynamic` Public IP Addresses aren't allocated until they're assigned to a resource (such as a Virtual Machine or a Load Balancer) by design within Azure - more information is available below.

- `ip_version` - (Optional) The IP Version to use, IPv6 or IPv4.

Note Only `dynamic` IP address allocation is supported for IPv6.

- `idle_timeout_in_minutes` - (Optional) Specifies the timeout for the TCP idle connection. The value can be set between 4 and 30 minutes.
- `domain_name_label` - (Optional) Label for the Domain Name. Will be used to make up the FQDN. If a domain name label is specified, an A DNS record is created for the public IP in the Microsoft Azure DNS system.
- `reverse_fqdn` - (Optional) A fully qualified domain name that resolves to this public IP address. If the `reverseFqdn` is specified, then a PTR DNS record is created pointing from the IP address in the `in-addr.arpa` domain to the reverse FQDN.
- `public_ip_prefix_id` - (Optional) If specified then public IP address allocated will be provided from the public IP prefix resource.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zones` - (Optional) A collection containing the availability zone to allocate the Public IP in.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

Attributes Reference

The following attributes are exported:

- `id` - The Public IP ID.
- `ip_address` - The IP address value that was allocated.

Note Dynamic Public IP Addresses aren't allocated until they're attached to a device (e.g. a Virtual Machine/Load Balancer). Instead you can obtain the IP Address once the the Public IP has been assigned via the `azurerm_public_ip` Data Source ([/docs/providers/azurerm/d/public_ip.html](/docs/providers/azurerm/d/public_ip.html)).

- `fqdn` - Fully qualified domain name of the A DNS record associated with the public IP. `domain_name_label` must be specified to get the `fqdn`. This is the concatenation of the `domain_name_label` and the regionalized DNS zone

Import

Public IPs can be imported using the `resource id`, e.g.

```
terraform import azurerm_public_ip.myPublicIp /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/publicIPAddresses/myPublicIpAddress1
```

azurerm_public_ip_prefix

Manages a Public IP Prefix.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "resourceGroup1"
  location  = "West Europe"
}

resource "azurerm_public_ip_prefix" "test" {
  name                  = "acceptanceTestPublicIpPrefix1"
  location              = "${azurerm_resource_group.test.location}"
  resource_group_name  = "${azurerm_resource_group.test.name}"

  prefix_length = 31

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Public IP resource . Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the public IP.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **sku** - (Optional) The SKU of the Public IP Prefix. Accepted values are `Standard` . Defaults to `Standard` . Changing this forces a new resource to be created.

Note Public IP Prefix can only be created with Standard SKUs at this time.

- **prefix_length** - (Optional) Specifies the number of bits of the prefix. The value can be set between 24 (256 addresses) and 31 (2 addresses). Changing this forces a new resource to be created.

Please Note:: There may be Public IP address limits on the subscription . More information available here (<https://docs.microsoft.com/en-us/azure/azure-subscription-service-limits?toc=%2fazure%2fvirtual-network%2ftoc.json#publicip-address>)

- **tags** - (Optional) A mapping of tags to assign to the resource.

- `zones` - (Optional) A collection containing the availability zone to allocate the Public IP in.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

Attributes Reference

The following attributes are exported:

- `id` - The Public IP Prefix ID.
- `ip_prefix` - The IP address prefix value that was allocated.

Import

Public IP Prefixes can be imported using the `resource id`, e.g.

```
terraform import azurerm_public_ip_prefix.myPublicIpPrefix /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/publicIPFixes/myPublicIpPrefix1
```

azurerm_recovery_network_mapping

Manages a site recovery network mapping on Azure. A network mapping decides how to translate connected networks when a VM is migrated from one region to another.

Example Usage

```

resource "azurerm_resource_group" "primary" {
  name      = "tfex-network-mapping-primary"
  location  = "West US"
}

resource "azurerm_resource_group" "secondary" {
  name      = "tfex-network-mapping-secondary"
  location  = "East US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name                = "example-recovery-vault"
  location            = "${azurerm_resource_group.secondary.location}"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  sku                 = "Standard"
}

resource "azurerm_recovery_services_fabric" "primary" {
  name                = "primary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_recovery_services_fabric" "secondary" {
  name                = "secondary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.secondary.location}"
  depends_on         = ["azurerm_recovery_services_fabric.primary"] # Avoids issues with creating fabric
s simultaneously
}

resource "azurerm_virtual_network" "primary" {
  name                = "network1"
  resource_group_name = "${azurerm_resource_group.primary.name}"
  address_space       = ["192.168.1.0/24"]
  location            = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_virtual_network" "secondary" {
  name                = "network2"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  address_space       = ["192.168.2.0/24"]
  location            = "${azurerm_resource_group.secondary.location}"
}

resource "azurerm_recovery_network_mapping" "recovery-mapping" {
  name                = "recovery-network-mapping-1"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  source_recovery_fabric_name = "primary-fabric"
  target_recovery_fabric_name = "secondary-fabric"
  source_network_id    = "${azurerm_virtual_network.primary.id}"
  target_network_id    = "${azurerm_virtual_network.secondary.id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.
- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `source_recovery_fabric_name` - (Required) Specifies the ASR fabric where mapping should be created.
- `target_recovery_fabric_name` - (Required) The Azure Site Recovery fabric object corresponding to the recovery Azure region.
- `source_network_id` - (Required) The id of the primary network.
- `target_network_id` - (Required) The id of the recovery network.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery network mapping can be imported using the `resource id`, e.g.

```
terraform import azurerm_recovery_network_mapping.mymapping /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/replicationFabrics/primary-fabric-name/replicationNetworks/azureNetwork/replicationNetworkMappings/mapping-name
```

azurerm_recovery_services_fabric

Manages a Azure recovery vault fabric.

Example Usage

```
resource "azurerm_resource_group" "primary" {
  name      = "tfex-network-mapping-primary"
  location  = "West US"
}

resource "azurerm_resource_group" "secondary" {
  name      = "tfex-network-mapping-secondary"
  location  = "East US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name                = "example-recovery-vault"
  location            = "${azurerm_resource_group.secondary.location}"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  sku                 = "Standard"
}

resource "azurerm_recovery_services_fabric" "fabric" {
  name                = "primary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.primary.location}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.
- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `location` - (Required) In what region should the fabric be located.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery recovery vault fabric can be imported using the `resource_id`, e.g.

```
terraform import azurerm_recovery_services_fabric.myfabric /subscriptions/00000000-0000-0000-0000-00000000  
000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/  
replicationFabrics/fabric-name
```

azurerm_recovery_services_protected_vm

Manages an Recovery Protected VM.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-recovery_vault"
  location  = "West US"
}

resource "azurerm_recovery_services_vault" "example" {
  name            = "tfex-recovery-vault"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"
}

resource "azurerm_recovery_services_protection_policy_vm" "example" {
  name                = "tfex-recovery-vault-policy"
  resource_group_name = "${azurerm_resource_group.example.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.example.name}"

  backup {
    frequency = "Daily"
    time      = "23:00"
  }
}

resource "azurerm_recovery_services_protected_vm" "vm1" {
  resource_group_name = "${azurerm_resource_group.example.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.example.name}"
  source_vm_id        = "${azurerm_virtual_machine.example.id}"
  backup_policy_id     = "${azurerm_recovery_services_protection_policy_vm.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `resource_group_name` - (Required) The name of the resource group in which to create the Recovery Services Protected VM. Changing this forces a new resource to be created.
- `recovery_vault_name` - (Required) Specifies the name of the Recovery Services Vault to use. Changing this forces a new resource to be created.
- `source_vm_id` - (Required) Specifies the ID of the VM to backup. Changing this forces a new resource to be created.
- `backup_policy_id` - (Required) Specifies the id of the backup policy to use.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Recovery Services Vault.

Import

Recovery Services Protected VMs can be imported using the `resource id`, e.g.

```
terraform import azurerm_recovery_services_protected_vm.item1 "/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.RecoveryServices/vaults/example-recovery-vault/backupFabrics/Azure/protectionContainers/iaasvmcontainer;iaasvmcontainerv2;group1;vm1/protectedItems/vm;iaasvmcontainerv2;group1;vm1"
```

Note the ID requires quoting as there are semicolons

azurerm_recovery_services_protection_container

Manages a Azure recovery vault protection container.

Example Usage

```
resource "azurerm_resource_group" "primary" {
  name      = "tfex-network-mapping-primary"
  location  = "West US"
}

resource "azurerm_resource_group" "secondary" {
  name      = "tfex-network-mapping-secondary"
  location  = "East US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name                = "example-recovery-vault"
  location            = "${azurerm_resource_group.secondary.location}"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  sku                = "Standard"
}

resource "azurerm_recovery_services_fabric" "fabric" {
  name                = "primary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_recovery_services_protection_container" "protection-container" {
  name                = "protection-container"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.fabric.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.
- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `recovery_fabric_name` - (Required) Name of fabric that should contain this protection container.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery recovery vault fabric can be imported using the `resource id`, e.g.

```
terraform import azurerm_recovery_services_protection_container.mycontainer /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/replicationFabrics/fabric-name/replicationProtectionContainers/protection-container-name
```

azurerm_recovery_services_protection_container_mapping

Manages a Azure recovery vault protection container mapping. A network protection container mapping decides how to translate the protection container when a VM is migrated from one region to another.

Example Usage

```
resource "azurerm_resource_group" "primary" {
  name      = "tfex-network-mapping-primary"
  location  = "West US"
}

resource "azurerm_resource_group" "secondary" {
  name      = "tfex-network-mapping-secondary"
  location  = "East US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name                = "example-recovery-vault"
  location            = "${azurerm_resource_group.secondary.location}"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  sku                = "Standard"
}

resource "azurerm_recovery_services_fabric" "primary" {
  name                = "primary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_recovery_services_fabric" "secondary" {
  name                = "secondary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location            = "${azurerm_resource_group.secondary.location}"
}

resource "azurerm_recovery_services_protection_container" "primary" {
  name                = "primary-protection-container"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.primary.name}"
}

resource "azurerm_recovery_services_protection_container" "secondary" {
  name                = "secondary-protection-container"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.secondary.name}"
}

resource "azurerm_recovery_services_replication_policy" "policy" {
  name                = "policy"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
}
```

```

resource_group_name      = "${azurerm_resource_group.secondary.name}"
recovery_vault_name      = "${azurerm_recovery_services_vault.vault.name}"
recovery_point_retention_in_minutes = "${24 * 60}"
application_consistent_snapshot_frequency_in_minutes = "${4 * 60}"
}

resource "azurerm_recovery_services_protection_container_mapping" "container-mapping" {
  name = "container-mapping"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.primary.name}"
  recovery_source_protection_container_name = "${azurerm_recovery_services_protection_container.primary.name}"
  recovery_target_protection_container_id = "${azurerm_recovery_services_protection_container.secondary.id}"
  recovery_replication_policy_id = "${azurerm_recovery_services_replication_policy.policy.id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.
- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `recovery_fabric_name` - (Required) Name of fabric that should contains the protection container to map.
- `recovery_source_protection_container_name` - (Required) Name of the protection container to map.
- `recovery_target_protection_container_id` - (Required) Id of protection container to map to.
- `recovery_replication_policy_id` - (Required) Id of the policy to use for this mapping.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery recovery vault fabric can be imported using the `resource id`, e.g.

```

terraform import azurerm_recovery_services_protection_container_mapping.mymapping /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/

```

azurerm_recovery_services_protection_policy_vm

Manages an Recovery Services VM Protection Policy.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-recovery_vault"
  location  = "West US"
}

resource "azurerm_recovery_services_vault" "example" {
  name            = "tfex-recovery-vault"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"
}

resource "azurerm_recovery_services_protection_policy_vm" "test" {
  name                = "tfex-recovery-vault-policy"
  resource_group_name = "${azurerm_resource_group.example.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.example.name}"

  timezone = "UTC"

  backup {
    frequency = "Daily"
    time      = "23:00"
  }

  retention_daily {
    count = 10
  }

  retention_weekly {
    count      = 42
    weekdays = ["Sunday", "Wednesday", "Friday", "Saturday"]
  }

  retention_monthly {
    count      = 7
    weekdays = ["Sunday", "Wednesday"]
    weeks      = ["First", "Last"]
  }

  retention_yearly {
    count      = 77
    weekdays = ["Sunday"]
    weeks      = ["Last"]
    months     = ["January"]
  }
}
```


Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Recovery Services Vault Policy. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the resource group in which to create the Recovery Services Protected VM. Changing this forces a new resource to be created.
 - `recovery_vault_name` - (Required) Specifies the name of the Recovery Services Vault to use. Changing this forces a new resource to be created.
 - `backup` - (Required) Configures the Policy backup frequent, times & days as documented in the `backup` block below.
 - `timezone` - (Optional) Specifies the timezone. Defaults to `UTC`
 - `retention_daily` - (Optional) Configures the policy daily retention as documented in the `retention_daily` block below. Required when backup frequency is `Daily`.
 - `retention_weekly` - (Optional) Configures the policy weekly retention as documented in the `retention_weekly` block below. Required when backup frequency is `Weekly`.
 - `retention_monthly` - (Optional) Configures the policy monthly retention as documented in the `retention_monthly` block below.
 - `retention_yearly` - (Optional) Configures the policy yearly retention as documented in the `retention_yearly` block below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

The `backup` block supports:

- `frequency` - (Required) Sets the backup frequency. Must be either `Daily` or `Weekly`.
 - `times` - (Required) The time of day to perform the backup in 24hour format.
 - `weekdays` - (Optional) The days of the week to perform backups on. Must be one of `Sunday`, `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday` or `Saturday`.
-

The `retention_daily` block supports:

- `count` - (Required) The number of daily backups to keep. Must be between `1` and `9999`
-

The `retention_weekly` block supports:

- `count` - (Required) The number of weekly backups to keep. Must be between `1` and `9999`
 - `weekdays` - (Required) The weekday backups to retain. Must be one of `Sunday`, `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday` or `Saturday`.
-

The `retention_monthly` block supports:

- **count** - (Required) The number of monthly backups to keep. Must be between 1 and 9999
 - **weekdays** - (Required) The weekday backups to retain . Must be one of Sunday , Monday , Tuesday , Wednesday , Thursday , Friday or Saturday .
 - **weeks** - (Required) The weeks of the month to retain backups of. Must be one of First , Second , Third , Fourth , Last .
-

The `retention_yearly` block supports:

- **count** - (Required) The number of yearly backups to keep. Must be between 1 and 9999
 - **weekdays** - (Required) The weekday backups to retain . Must be one of Sunday , Monday , Tuesday , Wednesday , Thursday , Friday or Saturday .
 - **weeks** - (Required) The weeks of the month to retain backups of. Must be one of First , Second , Third , Fourth , Last .
 - **months** - (Required) The months of the year to retain backups of. Must be one of January , February , March , April , May , June , July , August , September , October , November and December .
-

Attributes Reference

The following attributes are exported:

- **id** - The ID of the Recovery Services VM Protection Policy.

Import

Recovery Services VM Protection Policy can be imported using the `resource id` , e.g.

```
terraform import azurerm_recovery_services_protection_policy_vm.policy1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.RecoveryServices/vaults/example-recovery-vault/backupPolicies/policy1
```

azurerm_recovery_replicated_vm

Manages a Azure recovery replicated vms (Azure to Azure). An replicated VM keeps a copiously updated image of the vm in another region in order to be able to start the VM in that region in case of a disaster.

Example Usage

```
resource "azurerm_resource_group" "primary" {
  name      = "tfex-replicated-vm-primary"
  location  = "West US"
}

resource "azurerm_resource_group" "secondary" {
  name      = "tfex-replicated-vm-secondary"
  location  = "East US"
}

resource "azurerm_virtual_machine" "vm" {
  name                        = "vm"
  location                   = "${azurerm_resource_group.primary.location}"
  resource_group_name        = "${azurerm_resource_group.primary.name}"
  vm_size                    = "Standard_B1s"
  network_interface_ids      = ["${azurerm_network_interface.vm.id}"]

  storage_image_reference {
    publisher = "OpenLogic"
    offer     = "CentOS"
    sku       = "7.5"
    version   = "latest"
  }

  storage_os_disk {
    name          = "vm-os-disk"
    os_type       = "Linux"
    caching       = "ReadWrite"
    create_option = "FromImage"
    managed_disk_type = "Premium_LRS"
  }

  os_profile {
    admin_username = "test-admin-123"
    admin_password = "test-pwd-123"
    computer_name  = "vm"
  }

  os_profile_linux_config {
    disable_password_authentication = false
  }
}

resource "azurerm_recovery_services_vault" "vault" {
  name                        = "example-recovery-vault"
  location                   = "${azurerm_resource_group.secondary.location}"
  resource_group_name        = "${azurerm_resource_group.secondary.name}"
  sku                        = "Standard"
```

```

sku = "Standard"
}

resource "azurerm_recovery_services_fabric" "primary" {
  name = "primary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_recovery_services_fabric" "secondary" {
  name = "secondary-fabric"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  location = "${azurerm_resource_group.secondary.location}"
}

resource "azurerm_recovery_services_protection_container" "primary" {
  name = "primary-protection-container"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.primary.name}"
}

resource "azurerm_recovery_services_protection_container" "secondary" {
  name = "secondary-protection-container"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.secondary.name}"
}

resource "azurerm_recovery_services_replication_policy" "policy" {
  name = "policy"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_point_retention_in_minutes = "${24 * 60}"
  application_consistent_snapshot_frequency_in_minutes = "${4 * 60}"
}

resource "azurerm_recovery_services_protection_container_mapping" "container-mapping" {
  name = "container-mapping"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  recovery_fabric_name = "${azurerm_recovery_services_fabric.primary.name}"
  recovery_source_protection_container_name = "${azurerm_recovery_services_protection_container.primary.name}"
  recovery_target_protection_container_id = "${azurerm_recovery_services_protection_container.secondary.id}"
  recovery_replication_policy_id = "${azurerm_recovery_services_replication_policy.policy.id}"
}

resource "azurerm_storage_account" "primary" {
  name = "primaryrecoverycache"
  location = "${azurerm_resource_group.primary.location}"
  resource_group_name = "${azurerm_resource_group.primary.name}"
  account_tier = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_virtual_network" "primary" {

```

```

name                = "network1"
resource_group_name = "${azurerm_resource_group.primary.name}"
address_space       = ["192.168.1.0/24"]
location            = "${azurerm_resource_group.primary.location}"
}

resource "azurerm_subnet" "primary" {
  name                = "network1-subnet"
  resource_group_name = "${azurerm_resource_group.primary.name}"
  virtual_network_name = "${azurerm_virtual_network.primary.name}"
  address_prefix       = "192.168.1.0/24"
}

resource "azurerm_network_interface" "vm" {
  name                = "vm-nic"
  location            = "${azurerm_resource_group.primary.location}"
  resource_group_name = "${azurerm_resource_group.primary.name}"

  ip_configuration {
    name                = "vm"
    subnet_id           = "${azurerm_subnet.primary.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_recovery_replicated_vm" "vm-replication" {
  name                = "vm-replication"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name = "${azurerm_recovery_services_vault.vault.name}"
  source_recovery_fabric_name = "${azurerm_recovery_services_fabric.primary.name}"
  source_vm_id        = "${azurerm_virtual_machine.vm.id}"
  recovery_replication_policy_id = "${azurerm_recovery_services_replication_policy.policy.id}"
  source_recovery_protection_container_name = "${azurerm_recovery_services_protection_container.primary.name}"

  target_resource_group_id = "${azurerm_resource_group.secondary.id}"
  target_recovery_fabric_id = "${azurerm_recovery_services_fabric.secondary.id}"
  target_recovery_protection_container_id = "${azurerm_recovery_services_protection_container.secondary.id}"

  managed_disk {
    disk_id                = "${azurerm_virtual_machine.vm.storage_os_disk.0.managed_disk_id}"
    staging_storage_account_id = "${azurerm_storage_account.primary.id}"
    target_resource_group_id = "${azurerm_resource_group.secondary.id}"
    target_disk_type        = "Premium_LRS"
    target_replica_disk_type = "Premium_LRS"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.

- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `source_recovery_fabric_name` - (Required) Name of fabric that should contains this replication.
- `source_vm_id` - (Required) Id of the VM to replicate
- `source_recovery_protection_container_name` - (Required) Name of the protection container to use.
- `target_resource_group_id` - (Required) Id of resource group where the VM should be created when a failover is done.
- `target_recovery_fabric_id` - (Required) Id of fabric where the VM replication should be handled when a failover is done.
- `target_recovery_protection_container_id` - (Required) Id of protection container where the VM replication should be created when a failover is done.
- `target_availability_set_id` - (Optional) Id of availability set that the new VM should belong to when a failover is done.
- `managed_disk` - (Required) One or more `managed_disk` block.

A `managed_disk` block supports the following:

- `disk_id` - (Required) Id of disk that should be replicated.
- `staging_storage_account_id` - (Required) Storage account that should be used for caching.
- `target_resource_group_id` - (Required) Resource group disk should belong to when a failover is done.
- `target_disk_type` - (Required) What type should the disk be when a failover is done.
- `target_replica_disk_type` - (Required) What type should the disk be that holds the replication data.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery recovery vault fabric can be imported using the `resource id`, e.g.

```
terraform import azurerm_recovery_replicated_vm.vmreplication /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/replicationFabrics/fabric-name/replicationProtectionContainers/protection-container-name/replicationProtectedItems/vm-replication-name
```

azurerm_recovery_services_replication_policy

Manages a Azure recovery vault replication policy.

Example Usage

```
resource "azurerm_resource_group" "secondary" {
  name      = "tfex-network-mapping-secondary"
  location  = "East US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name            = "example-recovery-vault"
  location        = "${azurerm_resource_group.secondary.location}"
  resource_group_name = "${azurerm_resource_group.secondary.name}"
  sku             = "Standard"
}

resource "azurerm_recovery_services_replication_policy" "policy" {
  name                        = "policy"
  resource_group_name        = "${azurerm_resource_group.secondary.name}"
  recovery_vault_name        = "${azurerm_recovery_services_vault.vault.name}"
  recovery_point_retention_in_minutes = "${24 * 60}"
  application_consistent_snapshot_frequency_in_minutes = "${4 * 60}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the network mapping.
- `resource_group_name` - (Required) Name of the resource group where the vault that should be updated is located.
- `recovery_vault_name` - (Required) The name of the vault that should be updated.
- `recovery_point_retention_in_minutes` - (Required) Retain the recovery points for given time in minutes.
- `application_consistent_snapshot_frequency_in_minutes` - (Required) Specifies the frequency(in minutes) at which to create application consistent recovery points.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource ID.

Import

Site recovery recovery vault fabric can be imported using the `resource_id`, e.g.

```
terraform import azurerm_recovery_services_protection_container.mycontainer /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resource-group-name/providers/Microsoft.RecoveryServices/vaults/recovery-vault-name/replicationPolicies/policy-name
```


azurerm_recovery_services_vault

Manages an Recovery Services Vault.

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "tfex-recovery_vault"
  location  = "West US"
}

resource "azurerm_recovery_services_vault" "vault" {
  name                = "example_recovery_vault"
  location            = "${azurerm_resource_group.rg.location}"
  resource_group_name = "${azurerm_resource_group.rg.name}"
  sku                 = "Standard"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Recovery Services Vault. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Recovery Services Vault. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `sku` - (Required) Sets the vault's SKU. Possible values include: `Standard` , `RS0` .

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Recovery Services Vault.

Import

Recovery Services Vaults can be imported using the `resource id` , e.g.

```
terraform import azurerm_recovery_services_vault.vault1 /subscriptions/00000000-0000-0000-0000-000000000000  
00/resourceGroups/group1/providers/Microsoft.RecoveryServices/vaults/vault1
```

azurerm_redis_cache

Manages a Redis Cache.

Example Usage

This example provisions a Standard Redis Cache. Other examples of the `azurerm_redis_cache` resource can be found in the `./examples/redis-cache` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/redis-cache>)

```
resource "azurerm_resource_group" "example" {
  name     = "example-resources"
  location = "West Europe"
}

# NOTE: the Name used for Redis needs to be globally unique
resource "azurerm_redis_cache" "example" {
  name                = "example-cache"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  capacity             = 2
  family              = "C"
  sku_name             = "Standard"
  enable_non_ssl_port = false
  minimum_tls_version = "1.2"

  redis_configuration {}
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Redis instance. Changing this forces a new resource to be created.
- `location` - (Required) The location of the resource group.
- `resource_group_name` - (Required) The name of the resource group in which to create the Redis instance.
- `capacity` - (Required) The size of the Redis cache to deploy. Valid values for a `SKU` family of `C` (Basic/Standard) are 0, 1, 2, 3, 4, 5, 6, and for `P` (Premium) family are 1, 2, 3, 4.
- `family` - (Required) The SKU family/pricing group to use. Valid values are `C` (for Basic/Standard SKU family) and `P` (for Premium)
- `sku_name` - (Required) The SKU of Redis to use. Possible values are `Basic`, `Standard` and `Premium`.

- `enable_non_ssl_port` - (Optional) Enable the non-SSL port (6379) - disabled by default.
- `minimum_tls_version` - (Optional) The minimum TLS version. Defaults to `1.0`.
- `patch_schedule` - (Optional) A list of `patch_schedule` blocks as defined below - only available for Premium SKU's.
- `private_static_ip_address` - (Optional) The Static IP Address to assign to the Redis Cache when hosted inside the Virtual Network. Changing this forces a new resource to be created.
- `redis_configuration` - (Optional) A `redis_configuration` as defined below - with some limitations by SKU - defaults/details are shown below.
- `shard_count` - (Optional) *Only available when using the Premium SKU* The number of Shards to create on the Redis Cluster.
- `subnet_id` - (Optional) *Only available when using the Premium SKU* The ID of the Subnet within which the Redis Cache should be deployed. This Subnet must only contain Azure Cache for Redis instances without any other type of resources. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zones` - (Optional) A list of a single item of the Availability Zone which the Redis Cache should be allocated in.

Please Note: Availability Zones are in Preview and only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>) - as such you must be opted into the Preview to use this functionality. You can opt into the Availability Zones Preview in the Azure Portal (<http://aka.ms/azenroll>).

A `redis_configuration` block supports the following:

- `enable_authentication` - (Optional) If set to `false`, the Redis instance will be accessible without authentication. Defaults to `true`.

NOTE: `enable_authentication` can only be set to `false` if a `subnet_id` is specified; and only works if there aren't existing instances within the subnet with `enable_authentication` set to `true`.

- `maxmemory_reserved` - (Optional) Value in megabytes reserved for non-cache usage e.g. failover. Defaults are shown below.
- `maxmemory_delta` - (Optional) The max-memory delta for this Redis instance. Defaults are shown below.
- `maxmemory_policy` - (Optional) How Redis will select what to remove when `maxmemory` is reached. Defaults are shown below.
- `maxfragmentationmemory_reserved` - (Optional) Value in megabytes reserved to accommodate for memory fragmentation. Defaults are shown below.
- `rdp_backup_enabled` - (Optional) Is Backup Enabled? Only supported on Premium SKU's.

- `rdp_backup_frequency` - (Optional) The Backup Frequency in Minutes. Only supported on Premium SKU's. Possible values are: 15 , 30 , 60 , 360 , 720 and 1440 .
- `rdp_backup_max_snapshot_count` - (Optional) The maximum number of snapshots to create as a backup. Only supported for Premium SKU's.
- `rdp_storage_connection_string` - (Optional) The Connection String to the Storage Account. Only supported for Premium SKU's. In the format: `DefaultEndpointsProtocol=https;BlobEndpoint=${azurerm_storage_account.test.primary_blob_endpoint};AccountName=${azurerm_storage_account.test.name};AccountKey=${azur`

NOTE: There's a bug in the Redis API where the original storage connection string isn't being returned, which is being tracked in this issue (<https://github.com/Azure/azure-rest-api-specs/issues/3037>). In the interim you can use the `ignore_changes` attribute to ignore changes to this field (https://www.terraform.io/docs/configuration/resources.html#ignore_changes) e.g.:

```
resource "azurerm_redis_cache" "test" {
  # ...
  ignore_changes = ["redis_configuration.0.rdp_storage_connection_string"]
}
```

- `notify_keyspace_events` - (Optional) Keyspace notifications allows clients to subscribe to Pub/Sub channels in order to receive events affecting the Redis data set in some way. Reference (<https://redis.io/topics/notifications#configuration>)

```
redis_configuration {
  maxmemory_reserved = 10
  maxmemory_delta    = 2
  maxmemory_policy    = "allkeys-lru"
}
```

Default Redis Configuration Values

Redis Value	Basic	Standard	Premium
<code>enable_authentication</code>	true	true	true
<code>maxmemory_reserved</code>	2	50	200
<code>maxfragmentationmemory_reserved</code>	2	50	200
<code>maxmemory_delta</code>	2	50	200
<code>maxmemory_policy</code>	volatile-lru	volatile-lru	volatile-lru

NOTE: The `maxmemory_reserved`, `maxmemory_delta` and `maxfragmentationmemory_reserved` settings are only available for Standard and Premium caches. More details are available in the Relevant Links section below.

A `patch_schedule` block supports the following:

- `day_of_week` (Required) the Weekday name - possible values include `Monday` , `Tuesday` , `Wednesday` etc.
- `start_hour_utc` - (Optional) the Start Hour for maintenance in UTC - possible values range from 0 - 23 .

Note: The Patch Window lasts for 5 hours from the `start_hour_utc` .

Attributes Reference

The following attributes are exported:

- `id` - The Route ID.
- `hostname` - The Hostname of the Redis Instance
- `ssl_port` - The SSL Port of the Redis Instance
- `port` - The non-SSL Port of the Redis Instance
- `primary_access_key` - The Primary Access Key for the Redis Instance
- `secondary_access_key` - The Secondary Access Key for the Redis Instance
- `redis_configuration` - A `redis_configuration` block as defined below:

A `redis_configuration` block exports the following:

- `maxclients` - Returns the max number of connected clients at the same time.

Relevant Links

- Azure Redis Cache: SKU specific configuration limitations (<https://azure.microsoft.com/en-us/documentation/articles/cache-configure/#advanced-settings>)
- Redis: Available Configuration Settings (<http://redis.io/topics/config>)

Import

Redis Cache's can be imported using the `resource id`, e.g.

```
terraform import azurerm_redis_cache.cache1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Cache/Redis/cache1
```

azurerm_redis_firewall_rule

Manages a Firewall Rule associated with a Redis Cache.

Example Usage

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name      = "redis-resourcegroup"
  location  = "West Europe"
}

resource "azurerm_redis_cache" "test" {
  name                = "redis${random_id.server.hex}"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  capacity            = 1
  family              = "P"
  sku_name            = "Premium"
  enable_non_ssl_port = false

  redis_configuration {
    maxclients          = 256
    maxmemory_reserved = 2
    maxmemory_delta    = 2
    maxmemory_policy   = "allkeys-lru"
  }
}

resource "azurerm_redis_firewall_rule" "test" {
  name                = "someIPrange"
  redis_cache_name    = "${azurerm_redis_cache.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  start_ip            = "1.2.3.4"
  end_ip              = "2.3.4.5"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Firewall Rule. Changing this forces a new resource to be created.
- `redis_cache_name` - (Required) The name of the Redis Cache. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which this Redis Cache exists.
- `start_ip` - (Required) The lowest IP address included in the range
- `end_ip` - (Required) The highest IP address included in the range.

Attributes Reference

The following attributes are exported:

- `id` - The Redis Firewall Rule ID.

Import

Redis Firewall Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_redis_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Cache/Redis/cache1/firewallRules/rule1
```

azurerm_relay_namespace

Manages an Azure Relay Namespace.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_relay_namespace" "test" {
  name                = "example-relay"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  sku_name = "Standard"

  tags = {
    source = "terraform"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Azure Relay Namespace. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Azure Relay Namespace.
- `location` - (Required) Specifies the supported Azure location where the Azure Relay Namespace exists. Changing this forces a new resource to be created.
- `sku` - (Optional **Deprecated**) A `sku` block as described below.
- `sku_name` - (Optional) The name of the SKU to use. At this time the only supported value is `Standard`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `sku` block contains:

- `name` - (Optional) The name of the SKU to use. At this time the only supported value is `Standard`.

Attributes Reference

The following attributes are exported:

- `id` - The Azure Relay Namespace ID.

The following attributes are exported only if there is an authorization rule named `RootManageSharedAccessKey` which is created automatically by Azure.

- `primary_connection_string` - The primary connection string for the authorization rule `RootManageSharedAccessKey`.
- `secondary_connection_string` - The secondary connection string for the authorization rule `RootManageSharedAccessKey`.
- `primary_key` - The primary access key for the authorization rule `RootManageSharedAccessKey`.
- `secondary_key` - The secondary access key for the authorization rule `RootManageSharedAccessKey`.
- `metric_id` - The Identifier for Azure Insights metrics.

Import

Azure Relay Namespace's can be imported using the `resource id`, e.g.

```
terraform import azurerm_relay_namespace.relay1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Relay/namespaces/relay1
```

azurerm_resource_group

Manages a resource group on Azure.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "testResourceGroup1"
  location  = "West US"

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the resource group. Must be unique on your Azure subscription.
- `location` - (Required) The location where the resource group should be created. For a list of all Azure locations, please consult this link (<http://azure.microsoft.com/en-us/regions/>) or run `az account list-locations --output table`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

In addition to the arguments above, the following attributes are exported:

- `id` - The resource group ID.

Import

Resource Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_resource_group.mygroup /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup
```

azurerm_role_assignment

Assigns a given Principal (User or Application) to a given Role.

Example Usage (using a built-in Role)

```
data "azurerm_subscription" "primary" {}

data "azurerm_client_config" "test" {}

resource "azurerm_role_assignment" "test" {
  scope                = "${data.azurerm_subscription.primary.id}"
  role_definition_name = "Reader"
  principal_id         = "${data.azurerm_client_config.test.service_principal_object_id}"
}
```

Example Usage (Custom Role & Service Principal)

```
data "azurerm_subscription" "primary" {}

data "azurerm_client_config" "test" {}

resource "azurerm_role_definition" "test" {
  role_definition_id = "00000000-0000-0000-0000-000000000000"
  name               = "my-custom-role-definition"
  scope              = "${data.azurerm_subscription.primary.id}"

  permissions {
    actions      = ["Microsoft.Resources/subscriptions/resourceGroups/read"]
    not_actions = []
  }

  assignable_scopes = [
    "${data.azurerm_subscription.primary.id}",
  ]
}

resource "azurerm_role_assignment" "test" {
  name                = "00000000-0000-0000-0000-000000000000"
  scope               = "${data.azurerm_subscription.primary.id}"
  role_definition_id = "${azurerm_role_definition.test.id}"
  principal_id        = "${data.azurerm_client_config.test.service_principal_object_id}"
}
```

Example Usage (Custom Role & User)

```

data "azurerm_subscription" "primary" {}

data "azurerm_client_config" "test" {}

resource "azurerm_role_definition" "test" {
  role_definition_id = "00000000-0000-0000-0000-000000000000"
  name               = "my-custom-role-definition"
  scope              = "${data.azurerm_subscription.primary.id}"

  permissions {
    actions      = ["Microsoft.Resources/subscriptions/resourceGroups/read"]
    not_actions = []
  }

  assignable_scopes = [
    "${data.azurerm_subscription.primary.id}",
  ]
}

resource "azurerm_role_assignment" "test" {
  name                = "00000000-0000-0000-0000-000000000000"
  scope               = "${data.azurerm_subscription.primary.id}"
  role_definition_id = "${azurerm_role_definition.test.id}"
  principal_id        = "${data.azurerm_client_config.test.client_id}"
}

```

Example Usage (Custom Role & Management Group)

```

data "azurerm_subscription" "primary" {}

data "azurerm_client_config" "test" {}

data "azurerm_management_group" "test" {}

resource "azurerm_role_definition" "test" {
  role_definition_id = "000000000-0000-0000-0000-000000000000"
  name               = "my-custom-role-definition"
  scope              = "${data.azurerm_subscription.primary.id}"

  permissions {
    actions      = ["Microsoft.Resources/subscriptions/resourceGroups/read"]
    not_actions = []
  }

  assignable_scopes = [
    "${data.azurerm_subscription.primary.id}",
  ]
}

resource "azurerm_role_assignment" "test" {
  name                = "000000000-0000-0000-0000-000000000000"
  scope               = "${data.azurerm_management_group.primary.id}"
  role_definition_id = "${azurerm_role_definition.test.id}"
  principal_id        = "${data.azurerm_client_config.test.client_id}"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Optional) A unique UUID/GUID for this Role Assignment - one will be generated if not specified. Changing this forces a new resource to be created.
- `scope` - (Required) The scope at which the Role Assignment applies to, such as `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333`, `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup`, or `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup/providers/Microsoft.Compute/virtualMachines/myVM`, or `/providers/Microsoft.Management/managementGroups/myMG`. Changing this forces a new resource to be created.
- `role_definition_id` - (Optional) The Scoped-ID of the Role Definition. Changing this forces a new resource to be created. Conflicts with `role_definition_name`.
- `role_definition_name` - (Optional) The name of a built-in Role. Changing this forces a new resource to be created. Conflicts with `role_definition_id`.
- `principal_id` - (Required) The ID of the Principal (User, Group, Service Principal, or Application) to assign the Role Definition to. Changing this forces a new resource to be created.

NOTE: The Principal ID is also known as the Object ID (ie not the "Application ID" for applications).

- `skip_service_principal_aad_check` - (Optional) If the `principal_id` is a newly provisioned Service Principal set this value to `true` to skip the Azure Active Directory check which may fail due to replication lag. This argument is only valid if the `principal_id` is a Service Principal identity. If it is not a Service Principal identity it will cause the role assignment to fail. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The Role Assignment ID.
- `principal_type` - The type of the `principal_id`, e.g. User, Group, Service Principal, Application, etc.

Import

Role Assignments can be imported using the `resource id`, e.g.

```
terraform import azurerm_role_assignment.test /subscriptions/00000000-0000-0000-0000-000000000000/providers/Microsoft.Authorization/roleAssignments/00000000-0000-0000-0000-000000000000
```

azurerm_role_definition

Manages a custom Role Definition, used to assign Roles to Users/Principals. See 'Understand role definitions' (<https://docs.microsoft.com/en-us/azure/role-based-access-control/role-definitions>) in the Azure documentation for more details.

Example Usage

```
data "azurerm_subscription" "primary" {}

resource "azurerm_role_definition" "test" {
  name          = "my-custom-role"
  scope         = "${data.azurerm_subscription.primary.id}"
  description   = "This is a custom role created via Terraform"

  permissions {
    actions      = ["*"]
    not_actions = []
  }

  assignable_scopes = [
    "${data.azurerm_subscription.primary.id}", # /subscriptions/00000000-0000-0000-0000-000000000000
  ]
}
```

Argument Reference

The following arguments are supported:

- `role_definition_id` - (Optional) A unique UUID/GUID which identifies this role - one will be generated if not specified. Changing this forces a new resource to be created.
- `name` - (Required) The name of the Role Definition. Changing this forces a new resource to be created.
- `scope` - (Required) The scope at which the Role Definition applies too, such as `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333`, `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup`, or `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup/providers/Microsoft.Compute/virtualMachines/myVM`. Changing this forces a new resource to be created.
- `description` - (Optional) A description of the Role Definition.
- `permissions` - (Required) A `permissions` block as defined below.
- `assignable_scopes` - (Required) One or more assignable scopes for this Role Definition, such as `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333`, `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup`, or `/subscriptions/0b1f6471-1bf0-4dda-aec3-111122223333/resourceGroups/myGroup/providers/Microsoft.Compute/virtualMachines/myVM`.

A permissions block as the following properties:

- `actions` - (Optional) One or more Allowed Actions, such as `*` ,
`Microsoft.Resources/subscriptions/resourceGroups/read` . See 'Azure Resource Manager resource provider operations' (<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>) for details.
- `data_actions` - (Optional) One or more Allowed Data Actions, such as `*` ,
`Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read` . See 'Azure Resource Manager resource provider operations' (<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>) for details.
- `not_actions` - (Optional) One or more Disallowed Actions, such as `*` ,
`Microsoft.Resources/subscriptions/resourceGroups/read` . See 'Azure Resource Manager resource provider operations' (<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>) for details.
- `not_data_actions` - (Optional) One or more Disallowed Data Actions, such as `*` ,
`Microsoft.Resources/subscriptions/resourceGroups/read` . See 'Azure Resource Manager resource provider operations' (<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>) for details.

Attributes Reference

The following attributes are exported:

- `id` - The Role Definition ID.

Import

Role Definitions can be imported using the `resource id` , e.g.

```
terraform import azurerm_role_definition.test /subscriptions/00000000-0000-0000-0000-000000000000/providers/Microsoft.Authorization/roleDefinitions/00000000-0000-0000-0000-000000000000
```


azurerm_route

Manages a Route within a Route Table.

NOTE on Route Tables and Routes: Terraform currently provides both a standalone Route resource (</docs/providers/azurerm/r/route.html>), and allows for Routes to be defined in-line within the Route Table resource ([/docs/providers/azurerm/r/route_table.html](/docs/providers/azurerm/r/route_table.html)). At this time you cannot use a Route Table with in-line Routes in conjunction with any Route resources. Doing so will cause a conflict of Route configurations and will overwrite Routes.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_route_table" "test" {
  name                = "acceptanceTestRouteTable1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_route" "test" {
  name                = "acceptanceTestRoute1"
  resource_group_name = "${azurerm_resource_group.test.name}"
  route_table_name    = "${azurerm_route_table.test.name}"
  address_prefix      = "10.1.0.0/16"
  next_hop_type       = "vnetlocal"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the route. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the route. Changing this forces a new resource to be created.
- `route_table_name` - (Required) The name of the route table within which create the route. Changing this forces a new resource to be created.
- `address_prefix` - (Required) The destination CIDR to which the route applies, such as `10.1.0.0/16`
- `next_hop_type` - (Required) The type of Azure hop the packet should be sent to. Possible values are `VirtualNetworkGateway`, `VnetLocal`, `Internet`, `VirtualAppliance` and `None`
- `next_hop_in_ip_address` - (Optional) Contains the IP address packets should be forwarded to. Next hop values are only allowed in routes where the next hop type is `VirtualAppliance`.

Attributes Reference

The following attributes are exported:

- `id` - The Route ID.

Import

Routes can be imported using the `resource id`, e.g.

```
terraform import azurerm_route.testRoute /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/routeTables/mytable1/routes/myroute1
```

azurerm_route_table

Manages a Route Table

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_route_table" "test" {
  name                        = "acceptanceTestSecurityGroup1"
  location                   = "${azurerm_resource_group.test.location}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  disable_bgp_route_propagation = false

  route {
    name          = "route1"
    address_prefix = "10.1.0.0/16"
    next_hop_type = "vnetlocal"
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the route table. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the route table. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `route` - (Optional) List of objects (</docs/configuration/attr-as-blocks.html>) representing routes. Each object accepts the arguments documented below.
- `disable_bgp_route_propagation` - (Optional) Boolean flag which controls propagation of routes learned by BGP on that route table. True means disable.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Elements of `route` support:

- `name` - (Required) The name of the route.

- `address_prefix` - (Required) The destination CIDR to which the route applies, such as 10.1.0.0/16
- `next_hop_type` - (Required) The type of Azure hop the packet should be sent to. Possible values are `VirtualNetworkGateway`, `VnetLocal`, `Internet`, `VirtualAppliance` and `None`.
- `next_hop_in_ip_address` - (Optional) Contains the IP address packets should be forwarded to. Next hop values are only allowed in routes where the next hop type is `VirtualAppliance`.

Attributes Reference

The following attributes are exported:

- `id` - The Route Table ID.
- `subnets` - The collection of Subnets associated with this route table.

Import

Route Tables can be imported using the `resource id`, e.g.

```
terraform import azurerm_route_table.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/routeTables/mytable1
```

azurerm_scheduler_job_collection

Manages a Scheduler Job Collection.

NOTE: Support for Scheduler Job Collections has been deprecated by Microsoft in favour of Logic Apps (more information can be found at this link (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such we plan to remove support for this resource as a part of version 2.0 of the AzureRM Provider.

Example Usage

```
resource "azurerm_resource_group" "rg" {
  name      = "tfex-job_collection"
  location  = "West US"
}

resource "azurerm_scheduler_job_collection" "jobs" {
  name                  = "example_job_collection"
  location              = "${azurerm_resource_group.rg.location}"
  resource_group_name  = "${azurerm_resource_group.rg.name}"
  sku                  = "free"
  state                = "enabled"

  quota {
    max_job_count          = 5
    max_recurrence_interval = 24
    max_recurrence_frequency = "hour"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Scheduler Job Collection. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Scheduler Job Collection. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **tags** - (Optional) A mapping of tags to assign to the resource.
- **sku** - (Required) Sets the Job Collection's pricing level's SKU. Possible values include: `Standard`, `Free`, `P10Premium`, `P20Premium`.
- **state** - (Optional) Sets Job Collection's state. Possible values include: `Enabled`, `Disabled`, `Suspended`.

- `quota` - (Optional) Configures the Job collection quotas as documented in the `quota` block below.

The `quota` block supports:

- `max_job_count` - (Optional) Sets the maximum number of jobs in the collection.
- `max_recurrence_frequency` - (Required) The maximum frequency of recurrence. Possible values include: `Minute`, `Hour`, `Day`, `Week`, `Month`
- `max_recurrence_interval` - (Optional) The maximum interval between recurrence.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Scheduler Job Collection.

Import

Scheduler Job Collections can be imported using the `resource id`, e.g.

```
terraform import azurerm_scheduler_job_collection.jobcollection1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Scheduler/jobCollections/jobcollection1
```

azurerm_scheduler_job

Manages a Scheduler Job.

NOTE: Support for Scheduler Job has been deprecated by Microsoft in favour of Logic Apps (more information can be found at this link (<https://docs.microsoft.com/en-us/azure/scheduler/migrate-from-scheduler-to-logic-apps>)) - as such we plan to remove support for this resource as a part of version 2.0 of the AzureRM Provider.

Example Usage (single web get now)

```
resource "azurerm_scheduler_job" "web-once-now" {
  name = "tfex-web-once-now"
  resource_group_name = "${azurerm_resource_group.example.name}"
  job_collection_name = "${azurerm_scheduler_job_collection.example.name}"

  # re-enable it each run
  state = "enabled"

  action_web {
    # defaults to get
    url = "http://this.url.fails"
  }

  # default start time is now
}
```

Example Usage (recurring daily with retry and basic authentication)

```

resource "azurerm_scheduler_job" "web-recurring-daily" {
  name                        = "tfex-web-recurring-daily"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  job_collection_name        = "${azurerm_scheduler_job_collection.example.name}"

  action_web {
    url      = "https://this.url.fails"
    method   = "put"
    body     = "this is some text"

    headers = {
      Content-Type = "text"
    }

    authentication_basic {
      username = "login"
      password = "apassword"
    }
  }

  retry {
    # retry every 5 min a maximum of 10 times
    interval = "00:05:00"
    count    = 10
  }

  recurrence {
    frequency = "day"
    count     = 1000

    # run 4 times an hour every 12 hours
    hours      = [0, 12]
    minutes    = [0, 15, 30, 45]
  }

  start_time = "2018-07-07T07:07:07-07:00"
}

```

Example Usage (recurring monthly with an error action and client certificate authentication)

```

resource "azurerm_scheduler_job" "web-recurring-daily" {
  name                        = "tfex-web-recurring-daily"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  job_collection_name       = "${azurerm_scheduler_job_collection.example.name}"

  action_web {
    url = "https://this.url.fails"

    authentication_certificate {
      pfx      = "${filebase64("your_cert.pfx")}"
      password = "cert_password"
    }
  }

  error_action_web {
    url      = "https://this.url.fails"
    method   = "put"
    body     = "The job failed"

    headers = {
      "Content-Type" = "text"
    }

    authentication_basic {
      username = "login"
      password = "apassword"
    }
  }

  recurrence {
    frequency = "monthly"
    count     = 1000

    monthly_occurrences = [
      {
        # first Sunday
        day      = "Sunday"
        occurrence = 1
      },
      {
        # third Sunday
        day      = "Sunday"
        occurrence = 3
      },
      {
        # last Sunday
        day      = "Sunday"
        occurrence = -1
      },
    ]
  }

  start_time = "2018-07-07T07:07:07-07:00"
}

```

Example Usage (storage queue action)

```
resource "azurerm_storage_account" "example" {
  name                        = "tfexstorageaccount"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier                = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_queue" "example" {
  name                        = "tfex-schedulerjob-storagequeue"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  storage_account_name       = "${azurerm_storage_account.example.name}"
}

resource "azurerm_scheduler_job" "storage-once-now" {
  name                        = "tfex-storage-once-now"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  job_collection_name        = "${azurerm_scheduler_job_collection.example.name}"

  action_storage_queue {
    storage_account_name = "${azurerm_storage_account.example.name}"
    storage_queue_name   = "${azurerm_storage_queue.example.name}"
    sas_token             = "${azurerm_storage_account.example.primary_access_key}"
    message               = "storage message"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Scheduler Job. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Scheduler Job. Changing this forces a new resource to be created.
- `job_collection_name` - (Required) Specifies the name of the Scheduler Job Collection in which the Job should exist. Changing this forces a new resource to be created.
- `action_web` - (Optional) A `action_web` block defining the job action as described below. Note this is identical to an `error_action_web` block.

NOTE At least one of `error_action_web` or `action_storage_queue` needs to be set.

- `action_storage_queue` - (Optional) A `action_storage_queue` block defining a storage queue job action as described below. Note this is identical to an `error_action_storage_queue` block.
- `error_action_web` - (Optional) A `error_action_web` block defining the action to take on an error as described below. Note this is identical to an `action_web` block.

- `error_action_storage_queue` - (Optional) A `error_action_storage_queue` block defining the a web action to take on an error as described below. Note this is identical to an `action_storage_queue` block.
- `retry` - (Optional) A `retry` block defining how to retry as described below.
- `recurrence` - (Optional) A `recurrence` block defining a job occurrence schedule.
- `start_time` - (Optional) The time the first instance of the job is to start running at.
- `state` - (Optional) The sets or gets the current state of the job. Can be set to either `Enabled` or `Completed`

`web_action` & `error_web_action` block supports the following:

- `url` - (Required) Specifies the URL of the web request. Must be HTTPS for authenticated requests.
- `method` - (Optional) Specifies the method of the request. Defaults to `Get` and must be one of `Get`, `Put`, `Post`, `Delete`.
- `body` - (Optional) Specifies the request body.
- `headers` - (Optional) A map specifying the headers sent with the request.
- `authentication_basic` - (Optional) An `authentication_active_directory` block which defines the Active Directory oauth configuration to use.
- `authentication_certificate` - (Optional) An `authentication_certificate` block which defines the client certificate information to be use.
- `authentication_active_directory` - (Optional) An `authentication_active_directory` block which defines the OAUTH Active Directory information to use.

`authentication_basic` block supports the following:

- `username` - (Required) Specifies the username to use.
- `password` - (Required) Specifies the password to use.

`authentication_certificate` block supports the following:

- `pfx` - (Required) Specifies the pfx certificate in base-64 format.
- `password` - (Required) Specifies the certificate password.

`authentication_active_directory` block supports the following:

- `client_id` - (Required) Specifies the client ID to use.
- `tenant_id` - (Required) Specifies the tenant ID to use.
- `client_secret` - (Required) Specifies the secret to use.
- `audience` - (Optional) Specifies the audience.

`action_storage_queue` & `error_action_storage_queue` block supports the following:

- `storage_account_name` - (Required) Specifies the the storage account name.
- `storage_queue_name` - (Required) Specifies the the storage account queue.
- `sas_token` - (Required) Specifies a SAS token/key to authenticate with.

- `message` - (Required) The message to send into the queue.

`retry` block supports the following:

- `interval` - (Required) Specifies the duration between retries.
- `count` - (Required) Specifies the number of times a retry should be attempted.

`recurrence` block supports the following:

- `frequency` - (Required) Specifies the frequency of recurrence. Must be one of `Minute`, `Hour`, `Day`, `Week`, `Month`.
- `interval` - (Optional) Specifies the interval between executions. Defaults to `1`.
- `count` - (Optional) Specifies the maximum number of times that the job should run.
- `end_time` - (Optional) Specifies the time at which the job will cease running. Must be less than 500 days into the future.
- `minutes` - (Optional) Specifies the minutes of the hour that the job should execute at. Must be between `0` and `59`.
- `hours` - (Optional) Specifies the hours of the day that the job should execute at. Must be between `0` and `23`.
- `week_days` - (Optional) Specifies the days of the week that the job should execute on. Must be one of `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday`, `Sunday`. Only applies when `Week` is used for frequency.
- `month_days` - (Optional) Specifies the days of the month that the job should execute on. Must be non zero and between `-1` and `31`. Only applies when `Month` is used for frequency.
- `monthly_occurrences` - (Optional) Specifies specific monthly occurrences like "last sunday of the month" with `monthly_occurrences` blocks. Only applies when `Month` is used for frequency.

`monthly_occurrences` block supports the following:

- `day` - (Optional) Specifies the day of the week that the job should execute on. Must be one of one of `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday`, `Sunday`.
- `occurrence` - (Optional) Specifies the week the job should run on. For example `1` for the first week, `-1` for the last week of the month. Must be between `-5` and `5`.

Attributes Reference

The following attributes are exported:

- `id` - The Scheduler Job ID.

`authentication_certificate` block exports the following:

- `thumbprint` - (Computed) The certificate thumbprint.
- `expiration` - (Computed) The certificate expiration date.
- `subject_name` - (Computed) The certificate's certificate subject name.

Import

Scheduler Job can be imported using a resource id, e.g.

```
terraform import azurerm_scheduler_job.job1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/resourceGroup1/providers/Microsoft.Scheduler/jobCollections/jobCollection1/jobs/job1
```

azurerm_search_service

Allows you to manage an Azure Search Service.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_search_service" "test" {
  name                        = "acceptanceTestSearchService1"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  sku                        = "standard"

  tags = {
    environment = "staging"
    database    = "test"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Search Service. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Search Service. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) Valid values are `basic`, `free` and `standard`. `standard2` and `standard3` are also valid, but can only be used when it's enabled on the backend by Microsoft support. `free` provisions the service in shared clusters. `standard` provisions the service in dedicated clusters. Changing this forces a new resource to be created.
- `replica_count` - (Optional) Default is 1. Valid values include 1 through 12. Valid only when `sku` is `standard`. Changing this forces a new resource to be created.
- `partition_count` - (Optional) Default is 1. Valid values include 1, 2, 3, 4, 6, or 12. Valid only when `sku` is `standard`. Changing this forces a new resource to be created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Search Service ID.
- `primary_key` - The Search Service Administration primary key.
- `secondary_key` - The Search Service Administration secondary key.

Import

Search Services can be imported using the `resource id`, e.g.

```
terraform import azurerm_search_service.service1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Search/searchServices/service1
```

azurerm_security_center_contact

Manages the subscription's Security Center Contact.

NOTE: Owner access permission is required.

Example Usage

```
resource "azurerm_security_center_contact" "example" {
  email = "contact@example.com"
  phone = "+1-555-555-5555"

  alert_notifications = true
  alerts_to_admins    = true
}
```

Argument Reference

The following arguments are supported:

- `email` - (Required) The email of the Security Center Contact.
- `phone` - (Optional) The phone number of the Security Center Contact.
- `alert_notifications` - (Required) Whether to send security alerts notifications to the security contact.
- `alerts_to_admins` - (Required) Whether to send security alerts notifications to subscription admins.

Attributes Reference

The following attributes are exported:

- `id` - The Security Center Contact ID.

Import

The contact can be imported using the `resource id`, e.g.

```
terraform import azurerm_security_center_contact.example /subscriptions/00000000-0000-0000-0000-000000000000/providers/Microsoft.Security/securityContacts/default1
```


azurerm_security_center_subscription_pricing

Manages the Pricing Tier for Azure Security Center in the current subscription.

NOTE: This resource requires the `Owner` permission on the Subscription.

NOTE: Deletion of this resource does not change or reset the pricing tier to `Free`

Example Usage

```
resource "azurerm_security_center_subscription_pricing" "example" {
  tier = "Standard"
}
```

Argument Reference

The following arguments are supported:

- `tier` - (Required) The pricing tier to use. Possible values are `Free` and `Standard`.

NOTE: Changing the pricing tier to `Standard` affects all resources in the subscription and could be quite costly.

Attributes Reference

The following attributes are exported:

- `id` - The subscription pricing ID.

Import

The pricing tier can be imported using the `resource id`, e.g.

```
terraform import azurerm_security_center_subscription_pricing.example /subscriptions/00000000-0000-0000-0000-000000000000/providers/Microsoft.Security/pricings/default
```

azurerm_security_center_workspace

Manages the subscription's Security Center Workspace.

NOTE: Owner access permission is required.

NOTE: The subscription's pricing model can not be `Free` for this to have any affect.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-security-workspace"
  location  = "westus"
}

resource "azurerm_log_analytics_workspace" "example" {
  name                = "tfex-security-workspace"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                = "PerGB2018"
}

resource "azurerm_security_center_workspace" "example" {
  scope      = "/subscriptions/00000000-0000-0000-0000-000000000000"
  workspace_id = "${azurerm_log_analytics_workspace.example.id}"
}
```

Argument Reference

The following arguments are supported:

- `scope` - (Required) The scope of VMs to send their security data to the desired workspace, unless overridden by a setting with more specific scope.
- `workspace_id` - (Required) The ID of the Log Analytics Workspace to save the data in.

Attributes Reference

The following attributes are exported:

- `id` - The Security Center Workspace ID.

Import

The contact can be imported using the `resource id`, e.g.

```
terraform import azurerm_security_center_workspace.example /subscriptions/00000000-0000-0000-0000-00000000  
00000/providers/Microsoft.Security/workspaceSettings/default
```

azurerm_servicebus_namespace_authorization_rule

Manages a ServiceBus Namespace authorization Rule within a ServiceBus.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "terraform-servicebus"
  location = "West US"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "tfex_sevicebus_namespace"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_namespace_authorization_rule" "example" {
  name                = "examplerule"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  listen = true
  send   = true
  manage = false
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the ServiceBus Namespace Authorization Rule resource. Changing this forces a new resource to be created.
- **namespace_name** - (Required) Specifies the name of the ServiceBus Namespace. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which the ServiceBus Namespace exists. Changing this forces a new resource to be created.

NOTE At least one of the 3 permissions below needs to be set.

- **listen** - (Optional) Grants listen access to this this Authorization Rule. Defaults to `false`.
- **send** - (Optional) Grants send access to this this Authorization Rule. Defaults to `false`.

- `manage` - (Optional) Grants manage access to this this Authorization Rule. When this property is `true` - both `listen` and `send` must be too. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Topic ID.
- `primary_key` - The Primary Key for the ServiceBus Namespace authorization Rule.
- `primary_connection_string` - The Primary Connection String for the ServiceBus Namespace authorization Rule.
- `secondary_key` - The Secondary Key for the ServiceBus Namespace authorization Rule.
- `secondary_connection_string` - The Secondary Connection String for the ServiceBus Namespace authorization Rule.

Import

ServiceBus Namespace authorization rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_namespace_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ServiceBus/namespaces/namespace1/AuthorizationRules/rule1
```

azurerm_servicebus_namespace

Manages a ServiceBus Namespace.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "terraform-servicebus"
  location = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "tfex_sevicebus_namespace"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "Standard"

  tags = {
    source = "terraform"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the ServiceBus Namespace resource . Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the namespace.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `sku` - (Required) Defines which tier to use. Options are basic, standard or premium.
- `capacity` - (Optional) Specifies the capacity. When `sku` is `Premium`, capacity can be `1`, `2`, `4` or `8`. When `sku` is `Basic` or `Standard`, capacity can be `0` only.
- `zone_redundant` - (Optional) Whether or not this resource is zone redundant. `sku` needs to be `Premium`. Defaults to `false`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Namespace ID.

The following attributes are exported only if there is an authorization rule named `RootManageSharedAccessKey` which is created automatically by Azure.

- `default_primary_connection_string` - The primary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_connection_string` - The secondary connection string for the authorization rule `RootManageSharedAccessKey` .
- `default_primary_key` - The primary access key for the authorization rule `RootManageSharedAccessKey` .
- `default_secondary_key` - The secondary access key for the authorization rule `RootManageSharedAccessKey` .

Import

Service Bus Namespace can be imported using the `resource_id` , e.g.

```
terraform import azurerm_servicebus_namespace.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.servicebus/namespaces/sbns1
```

azurerm_servicebus_queue_authorization_rule

Manages an Authorization Rule for a ServiceBus Queue.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "terraform-servicebus"
  location  = "West US"
}

resource "azurerm_servicebus_namespace" "example" {
  name            = "tfex_sevicebus_namespace"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_queue" "example" {
  name                = "tfex_servicebus_queue"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}

resource "azurerm_servicebus_queue_authorization_rule" "example" {
  name                = "examplerule"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  queue_name          = "${azurerm_servicebus_queue.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  listen = true
  send   = true
  manage = false
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Authorization Rule. Changing this forces a new resource to be created.
- `namespace_name` - (Required) Specifies the name of the ServiceBus Namespace in which the Queue exists. Changing this forces a new resource to be created.

- `queue_name` - (Required) Specifies the name of the ServiceBus Queue. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the ServiceBus Namespace exists. Changing this forces a new resource to be created.

NOTE At least one of the 3 permissions below needs to be set.

- `listen` - (Optional) Does this Authorization Rule have Listen permissions to the ServiceBus Queue? Defaults to `false`.
- `send` - (Optional) Does this Authorization Rule have Send permissions to the ServiceBus Queue? Defaults to `false`.
- `manage` - (Optional) Does this Authorization Rule have Manage permissions to the ServiceBus Queue? When this property is `true` - both `listen` and `send` must be too. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Authorization Rule.
- `primary_key` - The Primary Key for the Authorization Rule.
- `primary_connection_string` - The Primary Connection String for the Authorization Rule.
- `secondary_key` - The Secondary Key for the Authorization Rule.
- `secondary_connection_string` - The Secondary Connection String for the Authorization Rule.

Import

ServiceBus Queue Authorization Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_queue_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ServiceBus/namespaces/namespace1/queues/queue1/authorizationRules/rule1
```

azurerm_servicebus_queue

Manages a ServiceBus Queue.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "terraform-servicebus"
  location = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name            = "tfex_sevicebus_namespace"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_queue" "example" {
  name                = "tfex_servicebus_queue"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the ServiceBus Queue resource. Changing this forces a new resource to be created.
- `namespace_name` - (Required) The name of the ServiceBus Namespace to create this queue in. Changing this forces a new resource to be created.
- `location` - (Optional / **Deprecated**) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the namespace. Changing this forces a new resource to be created.
- `auto_delete_on_idle` - (Optional) The ISO 8601 timespan duration of the idle interval after which the Queue is automatically deleted, minimum of 5 minutes.

- `default_message_ttl` - (Optional) The ISO 8601 timespan duration of the TTL of messages sent to this queue. This is the default value used when TTL is not set on message itself.
- `duplicate_detection_history_time_window` - (Optional) The ISO 8601 timespan duration during which duplicates can be detected. Default value is 10 minutes. (PT10M)
- `enable_express` - (Optional) Boolean flag which controls whether Express Entities are enabled. An express queue holds a message in memory temporarily before writing it to persistent storage. Defaults to `false` for Basic and Standard. For Premium, it MUST be set to `false` .

NOTE: Service Bus Premium namespaces do not support Express Entities, so `enable_express` MUST be set to `false` .

- `enable_partitioning` - (Optional) Boolean flag which controls whether to enable the queue to be partitioned across multiple message brokers. Changing this forces a new resource to be created. Defaults to `false` for Basic and Standard. For Premium, it MUST be set to `true` .

NOTE: Partitioning is available at entity creation for all queues and topics in Basic or Standard SKUs. It is not available for the Premium messaging SKU, but any previously existing partitioned entities in Premium namespaces continue to work as expected. Please see the documentation (<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-partitioning>) for more information.

- `lock_duration` - (Optional) The ISO 8601 timespan duration of a peek-lock; that is, the amount of time that the message is locked for other receivers. Maximum value is 5 minutes. Defaults to 1 minute. (PT1M)
- `max_size_in_megabytes` - (Optional) Integer value which controls the size of memory allocated for the queue. For supported values see the "Queue/topic size" section of this document (<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quotas>).
- `requires_duplicate_detection` - (Optional) Boolean flag which controls whether the Queue requires duplicate detection. Changing this forces a new resource to be created. Defaults to `false` .
- `requires_session` - (Optional) Boolean flag which controls whether the Queue requires sessions. This will allow ordered handling of unbounded sequences of related messages. With sessions enabled a queue can guarantee first-in-first-out delivery of messages. Changing this forces a new resource to be created. Defaults to `false` .
- `dead_lettering_on_message_expiration` - (Optional) Boolean flag which controls whether the Queue has dead letter support when a message expires. Defaults to `false` .
- `max_delivery_count` - (Optional) Integer value which controls when a message is automatically deadlettered. Defaults to 10 .

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Queue ID.

Import

Service Bus Queue can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_queue.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.servicebus/namespaces/sbns1/queues/snqueue1
```

azurerm_servicebus_subscription

Manages a ServiceBus Subscription.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "tfex-servicebus-subscription"
  location = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name            = "tfex_sevicebus_namespace"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_topic" "example" {
  name            = "tfex_sevicebus_topic"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}

resource "azurerm_servicebus_subscription" "example" {
  name            = "tfex_sevicebus_subscription"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"
  topic_name       = "${azurerm_servicebus_topic.example.name}"
  max_delivery_count = 1
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the ServiceBus Subscription resource. Changing this forces a new resource to be created.
- `namespace_name` - (Required) The name of the ServiceBus Namespace to create this Subscription in. Changing this forces a new resource to be created.
- `topic_name` - (Required) The name of the ServiceBus Topic to create this Subscription in. Changing this forces a new resource to be created.

- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the namespace. Changing this forces a new resource to be created.
- `max_delivery_count` - (Required) The maximum number of deliveries.
- `auto_delete_on_idle` - (Optional) The idle interval after which the Subscription is automatically deleted, minimum of 5 minutes. Provided in the TimeSpan format.
- `default_message_ttl` - (Optional) The TTL of messages sent to this Subscription if no TTL value is set on the message itself. Provided in the TimeSpan format.
- `lock_duration` - (Optional) The lock duration for the subscription, maximum supported value is 5 minutes. Defaults to 1 minute.
- `dead_lettering_on_message_expiration` - (Optional) Boolean flag which controls whether the Subscription has dead letter support when a message expires. Defaults to false.
- `enable_batched_operations` - (Optional) Boolean flag which controls whether the Subscription supports batched operations. Defaults to false.
- `requires_session` - (Optional) Boolean flag which controls whether this Subscription supports the concept of a session. Defaults to false. Changing this forces a new resource to be created.
- `forward_to` - (Optional) The name of a Queue or Topic to automatically forward messages to.

TimeSpan Format

Some arguments for this resource are required in the TimeSpan format which is used to represent a length of time. The supported format is documented here ([https://msdn.microsoft.com/en-us/library/se73z7b9\(v=vs.110\).aspx#Anchor_2](https://msdn.microsoft.com/en-us/library/se73z7b9(v=vs.110).aspx#Anchor_2))

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Subscription ID.

Import

Service Bus Subscriptions can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_subscription.test /subscriptions/00000000-0000-0000-0000-000000000000
/resourceGroups/mygroup1/providers/microsoft.servicebus/namespaces/sbns1/topics/sntopic1/subscriptions/sb
sub1
```

azurerm_servicebus_subscription_rule

Manages a ServiceBus Subscription Rule.

Example Usage (SQL Filter)

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-servicebus-subscription-rule-sql"
  location  = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "tfex_sevicebus_namespace"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_topic" "example" {
  name                = "tfex_sevicebus_topic"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}

resource "azurerm_servicebus_subscription" "example" {
  name                = "tfex_sevicebus_subscription"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  topic_name          = "${azurerm_servicebus_topic.example.name}"
  max_delivery_count  = 1
}

resource "azurerm_servicebus_subscription_rule" "example" {
  name                = "tfex_sevicebus_rule"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  topic_name          = "${azurerm_servicebus_topic.example.name}"
  subscription_name    = "${azurerm_servicebus_subscription.example.name}"
  filter_type          = "SqlFilter"
  sql_filter           = "color = 'red'"
}
```

Example Usage (Correlation Filter)

```

resource "azurerm_resource_group" "example" {
  name      = "tfex-servicebus-subscription-rule-cor"
  location = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name            = "tfex_sevicebus_namespace"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_topic" "example" {
  name            = "tfex_sevicebus_topic"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}

resource "azurerm_servicebus_subscription" "example" {
  name            = "tfex_sevicebus_subscription"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"
  topic_name       = "${azurerm_servicebus_topic.example.name}"
  max_delivery_count = 1
}

resource "azurerm_servicebus_subscription_rule" "example" {
  name            = "tfex_sevicebus_rule"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"
  topic_name       = "${azurerm_servicebus_topic.example.name}"
  subscription_name = "${azurerm_servicebus_subscription.example.name}"
  filter_type      = "CorrelationFilter"

  correlation_filter {
    correlation_id = "high"
    label         = "red"
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the ServiceBus Subscription Rule. Changing this forces a new resource to be created.
- **namespace_name** - (Required) The name of the ServiceBus Namespace in which the ServiceBus Topic exists. Changing this forces a new resource to be created.

- `topic_name` - (Required) The name of the ServiceBus Topic in which the ServiceBus Subscription exists. Changing this forces a new resource to be created.
- `subscription_name` - (Required) The name of the ServiceBus Subscription in which this Rule should be created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in the ServiceBus Namespace exists. Changing this forces a new resource to be created.
- `filter_type` - (Required) Type of filter to be applied to a BrokeredMessage. Possible values are `SqlFilter` and `CorrelationFilter`.
- `sql_filter` - (Optional) Represents a filter written in SQL language-based syntax that to be evaluated against a BrokeredMessage. Required when `filter_type` is set to `SqlFilter`.
- `correlation_filter` - (Optional) A `correlation_filter` block as documented below to be evaluated against a BrokeredMessage. Required when `filter_type` is set to `CorrelationFilter`.
- `action` - (Optional) Represents set of actions written in SQL language-based syntax that is performed against a BrokeredMessage.

`correlation_filter` supports the following:

- `content_type` - (Optional) Content type of the message.
- `correlation_id` - (Optional) Identifier of the correlation.
- `label` - (Optional) Application specific label.
- `message_id` - (Optional) Identifier of the message.
- `reply_to` - (Optional) Address of the queue to reply to.
- `reply_to_session_id` - (Optional) Session identifier to reply to.
- `session_id` - (Optional) Session identifier.
- `to` - (Optional) Address to send to.

NOTE: When creating a subscription rule of type `CorrelationFilter` at least one property must be set in the `correlation_filter` block.

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Subscription Rule ID.

Import

Service Bus Subscription Rule can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_subscription_rule.test /subscriptions/00000000-0000-0000-0000-00000000  
0000/resourceGroups/mygroup1/providers/microsoft.servicebus/namespaces/sbns1/topics/sntopic1/subscriptio  
ns/sbsub1/rules/sbrule1
```

azurerm_servicebus_topic_authorization_rule

Manages a ServiceBus Topic authorization Rule within a ServiceBus Topic.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "tfex-servicebus"
  location  = "West US"
}

resource "azurerm_servicebus_namespace" "example" {
  name            = "tfex_servicebus_namespace"
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku             = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_topic" "example" {
  name            = "tfex_servicebus_topic"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"
}

resource "azurerm_servicebus_topic_authorization_rule" "example" {
  name            = "tfex_servicebus_topic_sasPolicy"
  namespace_name   = "${azurerm_servicebus_namespace.example.name}"
  topic_name       = "${azurerm_servicebus_topic.example.name}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  listen           = true
  send             = false
  manage           = false
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the ServiceBus Topic Authorization Rule resource. Changing this forces a new resource to be created.
- `namespace_name` - (Required) Specifies the name of the ServiceBus Namespace. Changing this forces a new resource to be created.
- `topic_name` - (Required) Specifies the name of the ServiceBus Topic. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which the ServiceBus Namespace exists. Changing this forces a new resource to be created.

NOTE At least one of the 3 permissions below needs to be set.

- `listen` - (Optional) Grants listen access to this this Authorization Rule. Defaults to `false`.
- `send` - (Optional) Grants send access to this this Authorization Rule. Defaults to `false`.
- `manage` - (Optional) Grants manage access to this this Authorization Rule. When this property is `true` - both `listen` and `send` must be too. Defaults to `false`.

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Topic ID.
- `primary_key` - The Primary Key for the ServiceBus Topic authorization Rule.
- `primary_connection_string` - The Primary Connection String for the ServiceBus Topic authorization Rule.
- `secondary_key` - The Secondary Key for the ServiceBus Topic authorization Rule.
- `secondary_connection_string` - The Secondary Connection String for the ServiceBus Topic authorization Rule.

Import

ServiceBus Topic authorization rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_topic_authorization_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.ServiceBus/namespaces/namespace1/topics/topic1/authorizationRules/rule1
```

azurerm_servicebus_topic

Manages a ServiceBus Topic.

Note Topics can only be created in Namespaces with an SKU of `standard` or higher.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "tfex-servicebus-topic"
  location = "West Europe"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "tfex_sevicebus_namespace"
  location            = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
  sku                 = "Standard"

  tags = {
    source = "terraform"
  }
}

resource "azurerm_servicebus_topic" "example" {
  name                = "tfex_sevicebus_topic"
  resource_group_name = "${azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"

  enable_partitioning = true
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the ServiceBus Topic resource. Changing this forces a new resource to be created.
- `namespace_name` - (Required) The name of the ServiceBus Namespace to create this topic in. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the namespace. Changing this forces a new resource to be created.
- `status` - (Optional) The Status of the Service Bus Topic. Acceptable values are `Active` or `Disabled`. Defaults to `Active`.

- `auto_delete_on_idle` - (Optional) The ISO 8601 timespan duration of the idle interval after which the Topic is automatically deleted, minimum of 5 minutes.
- `default_message_ttl` - (Optional) The ISO 8601 timespan duration of TTL of messages sent to this topic if no TTL value is set on the message itself.
- `duplicate_detection_history_time_window` - (Optional) The ISO 8601 timespan duration during which duplicates can be detected. Defaults to 10 minutes. (PT10M)
- `enable_batched_operations` - (Optional) Boolean flag which controls if server-side batched operations are enabled. Defaults to false.
- `enable_express` - (Optional) Boolean flag which controls whether Express Entities are enabled. An express topic holds a message in memory temporarily before writing it to persistent storage. Defaults to false.
- `enable_partitioning` - (Optional) Boolean flag which controls whether to enable the topic to be partitioned across multiple message brokers. Defaults to false. Changing this forces a new resource to be created.

NOTE: Partitioning is available at entity creation for all queues and topics in Basic or Standard SKUs. It is not available for the Premium messaging SKU, but any previously existing partitioned entities in Premium namespaces continue to work as expected. Please see the documentation (<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-partitioning>) for more information.

- `max_size_in_megabytes` - (Optional) Integer value which controls the size of memory allocated for the topic. For supported values see the "Queue/topic size" section of this document (<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quotas>).
- `requires_duplicate_detection` - (Optional) Boolean flag which controls whether the Topic requires duplicate detection. Defaults to false. Changing this forces a new resource to be created.
- `support_ordering` - (Optional) Boolean flag which controls whether the Topic supports ordering. Defaults to false.

Attributes Reference

The following attributes are exported:

- `id` - The ServiceBus Topic ID.

Import

Service Bus Topics can be imported using the `resource id`, e.g.

```
terraform import azurerm_servicebus_topic.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.servicebus/namespaces/sbns1/topics/sntopic1
```

azurerm_service_fabric_cluster

Manages a Service Fabric Cluster.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_service_fabric_cluster" "test" {
  name                        = "example-servicefabric"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  reliability_level           = "Bronze"
  upgrade_mode                = "Manual"
  cluster_code_version        = "6.3.176.9494"
  vm_image                   = "Windows"
  management_endpoint         = "https://example:80"

  node_type {
    name              = "first"
    instance_count    = 3
    is_primary        = true
    client_endpoint_port = 2020
    http_endpoint_port  = 80
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Service Fabric Cluster. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Service Fabric Cluster exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the Azure Region where the Service Fabric Cluster should exist. Changing this forces a new resource to be created.
- `reliability_level` - (Required) Specifies the Reliability Level of the Cluster. Possible values include `None`, `Bronze`, `Silver`, `Gold` and `Platinum`.

NOTE: The Reliability Level of the Cluster depends on the number of nodes in the Cluster: `Platinum` requires at least 9 VM's, `Gold` requires at least 7 VM's, `Silver` requires at least 5 VM's, `Bronze` requires at least 3 VM's.

- `management_endpoint` - (Required) Specifies the Management Endpoint of the cluster such as `http://example.com`. Changing this forces a new resource to be created.
- `node_type` - (Required) One or more `node_type` blocks as defined below.
- `upgrade_mode` - (Required) Specifies the Upgrade Mode of the cluster. Possible values are `Automatic` or `Manual`.
- `vm_image` - (Required) Specifies the Image expected for the Service Fabric Cluster, such as `Windows`. Changing this forces a new resource to be created.

-
- `cluster_code_version` - (Optional) Required if Upgrade Mode set to `Manual`, Specifies the Version of the Cluster Code of the cluster.
 - `add_on_features` - (Optional) A List of one or more features which should be enabled, such as `DnsService`.
 - `azure_active_directory` - (Optional) An `azure_active_directory` block as defined below.
 - `certificate_common_names` - (Optional) A `certificate_common_names` block as defined below. Conflicts with `certificate`.
 - `certificate` - (Optional) A `certificate` block as defined below. Conflicts with `certificate_common_names`.
 - `reverse_proxy_certificate` - (Optional) A `reverse_proxy_certificate` block as defined below.
 - `client_certificate_thumbprint` - (Optional) One or two `client_certificate_thumbprint` blocks as defined below.

NOTE: If Client Certificates are enabled then at a Certificate must be configured on the cluster.

- `diagnostics_config` - (Optional) A `diagnostics_config` block as defined below. Changing this forces a new resource to be created.
- `fabric_settings` - (Optional) One or more `fabric_settings` blocks as defined below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `azure_active_directory` block supports the following:

- `tenant_id` - (Required) The Azure Active Directory Tenant ID.
- `cluster_application_id` - (Required) The Azure Active Directory Cluster Application ID.
- `client_application_id` - (Required) The Azure Active Directory Client ID which should be used for the Client Application.

A `certificate_common_names` block supports the following:

- `common_names` - (Required) A `common_names` block as defined below.
- `x509_store_name` - (Required) The X509 Store where the Certificate Exists, such as `My`.

A `common_names` block supports the following:

- `certificate_common_name` - (Required) The common or subject name of the certificate.
- `certificate_issuer_thumbprint` - (Optional) The Issuer Thumbprint of the Certificate.

NOTE: Certificate Issuer Thumbprint may become required in the future, <https://docs.microsoft.com/en-us/azure/service-fabric/service-fabric-create-cluster-using-cert-cn#download-and-update-a-sample-template>.

A `certificate` block supports the following:

- `thumbprint` - (Required) The Thumbprint of the Certificate.
- `thumbprint_secondary` - (Required) The Secondary Thumbprint of the Certificate.
- `x509_store_name` - (Required) The X509 Store where the Certificate Exists, such as `My`.

A `reverse_proxy_certificate` block supports the following:

- `thumbprint` - (Required) The Thumbprint of the Certificate.
- `thumbprint_secondary` - (Required) The Secondary Thumbprint of the Certificate.
- `x509_store_name` - (Required) The X509 Store where the Certificate Exists, such as `My`.

A `client_certificate_thumbprint` block supports the following:

- `thumbprint` - (Required) The Thumbprint associated with the Client Certificate.
- `is_admin` - (Required) Does the Client Certificate have Admin Access to the cluster? Non-admin clients can only perform read only operations on the cluster.

A `diagnostics_config` block supports the following:

- `storage_account_name` - (Required) The name of the Storage Account where the Diagnostics should be sent to.
- `protected_account_key_name` - (Required) The protected diagnostics storage key name, such as `StorageAccountKey1`.
- `blob_endpoint` - (Required) The Blob Endpoint of the Storage Account.
- `queue_endpoint` - (Required) The Queue Endpoint of the Storage Account.
- `table_endpoint` - (Required) The Table Endpoint of the Storage Account.

A `fabric_settings` block supports the following:

- `name` - (Required) The name of the Fabric Setting, such as `Security` or `Federation`.
- `parameters` - (Optional) A map containing settings for the specified Fabric Setting.

A `node_type` block supports the following:

- `name` - (Required) The name of the Node Type. Changing this forces a new resource to be created.
 - `placement_properties` - (Optional) The placement tags applied to nodes in the node type, which can be used to indicate where certain services (workload) should run.
 - `capacities` - (Optional) The capacity tags applied to the nodes in the node type, the cluster resource manager uses these tags to understand how much resource a node has.
 - `instance_count` - (Required) The number of nodes for this Node Type.
 - `is_primary` - (Required) Is this the Primary Node Type? Changing this forces a new resource to be created.
 - `client_endpoint_port` - (Required) The Port used for the Client Endpoint for this Node Type. Changing this forces a new resource to be created.
 - `http_endpoint_port` - (Required) The Port used for the HTTP Endpoint for this Node Type. Changing this forces a new resource to be created.
 - `durability_level` - (Optional) The Durability Level for this Node Type. Possible values include `Bronze`, `Gold` and `Silver`. Defaults to `Bronze`. Changing this forces a new resource to be created.
 - `application_ports` - (Optional) A `application_ports` block as defined below.
 - `ephemeral_ports` - (Optional) A `ephemeral_ports` block as defined below.
 - `reverse_proxy_endpoint_port` - (Optional) The Port used for the Reverse Proxy Endpoint for this Node Type. Changing this will upgrade the cluster.
-

A `application_ports` block supports the following:

- `start_port` - (Required) The start of the Application Port Range on this Node Type.
 - `end_port` - (Required) The end of the Application Port Range on this Node Type.
-

A `ephemeral_ports` block supports the following:

- `start_port` - (Required) The start of the Ephemeral Port Range on this Node Type.
- `end_port` - (Required) The end of the Ephemeral Port Range on this Node Type.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Service Fabric Cluster.
- `cluster_endpoint` - The Cluster Endpoint for this Service Fabric Cluster.

Import

Service Fabric Clusters can be imported using the `resource id`, e.g.

```
terraform import azurerm_service_fabric_cluster.cluster1 /subscriptions/00000000-0000-0000-0000-000000000000/000/resourceGroups/mygroup1/providers/Microsoft.ServiceFabric/clusters/cluster1
```

azurerm_shared_image_gallery

Manages a Shared Image Gallery.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_shared_image_gallery" "test" {
  name                        = "example_image_gallery"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  description                 = "Shared images and things."

  tags = {
    Hello = "There"
    World = "Example"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Shared Image Gallery. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the Shared Image Gallery. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **description** - (Optional) A description for this Shared Image Gallery.
- **tags** - (Optional) A mapping of tags to assign to the Shared Image Gallery.

Attributes Reference

The following attributes are exported:

- **id** - The ID of the Shared Image Gallery.
- **unique_name** - The Unique Name for this Shared Image Gallery.

Import

Shared Image Galleries can be imported using the `resource id`, e.g.

```
terraform import azurerm_shared_image_gallery.gallery1 /subscriptions/00000000-0000-0000-0000-000000000000  
0/resourceGroups/mygroup1/providers/Microsoft.Compute/galleries/gallery1
```

azurerm_shared_image

Manages a Shared Image within a Shared Image Gallery.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_shared_image_gallery" "test" {
  name                = "example_image_gallery"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  description         = "Shared images and things."

  tags = {
    Hello = "There"
    World = "Example"
  }
}

resource "azurerm_shared_image" "test" {
  name                = "my-image"
  gallery_name       = "${azurerm_shared_image_gallery.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location           = "${azurerm_resource_group.test.location}"
  os_type            = "Linux"

  identifier {
    publisher = "PublisherName"
    offer     = "OfferName"
    sku       = "ExampleSku"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Shared Image. Changing this forces a new resource to be created.
- `gallery_name` - (Required) Specifies the name of the Shared Image Gallery in which this Shared Image should exist. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which the Shared Image Gallery exists. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the Shared Image Gallery exists. Changing this forces a new resource to be created.

- `identifier` - (Required) An `identifier` block as defined below.
- `os_type` - (Required) The type of Operating System present in this Shared Image. Possible values are `Linux` and `Windows`.

-
- `description` - (Optional) A description of this Shared Image.
 - `eula` - (Optional) The End User Licence Agreement for the Shared Image.
 - `privacy_statement_uri` - (Optional) The URI containing the Privacy Statement associated with this Shared Image.
 - `release_note_uri` - (Optional) The URI containing the Release Notes associated with this Shared Image.
 - `tags` - (Optional) A mapping of tags to assign to the Shared Image.
-

A `identifier` block supports the following:

- `offer` - (Required) The Offer Name for this Shared Image.
- `publisher` - (Required) The Publisher Name for this Gallery Image.
- `sku` - (Required) The Name of the SKU for this Gallery Image.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Shared Image.

Import

Shared Images can be imported using the `resource id`, e.g.

```
terraform import azurerm_shared_image.image1 /subscriptions/00000000-0000-0000-0000-000000000000/resource  
Groups/mygroup1/providers/Microsoft.Compute/galleries/gallery1/images/image1
```

azurerm_shared_image_version

Manages a Version of a Shared Image within a Shared Image Gallery.

Example Usage

```
data "azurerm_image" "existing" {
  name                = "search-api"
  resource_group_name = "packerimages"
}

data "azurerm_shared_image" "existing" {
  name                = "existing-image"
  gallery_name        = "existing_gallery"
  resource_group_name = "existing-resources"
}

resource "azurerm_shared_image_version" "test" {
  name                = "0.0.1"
  gallery_name        = "${data.azurerm_shared_image.existing.gallery_name}"
  image_name          = "${data.azurerm_shared_image.existing.name}"
  resource_group_name = "${data.azurerm_shared_image.existing.resource_group_name}"
  location            = "${data.azurerm_shared_image.existing.location}"
  managed_image_id    = "${data.azurerm_image.existing.id}"

  target_region {
    name                = "${data.azurerm_shared_image.existing.location}"
    regional_replica_count = "5"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The version number for this Image Version, such as `1.0.0` . Changing this forces a new resource to be created.
- **gallery_name** - (Required) The name of the Shared Image Gallery in which the Shared Image exists. Changing this forces a new resource to be created.
- **image_name** - (Required) The name of the Shared Image within the Shared Image Gallery in which this Version should be created. Changing this forces a new resource to be created.
- **location** - (Required) The Azure Region in which the Shared Image Gallery exists. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the Resource Group in which the Shared Image Gallery exists. Changing this forces a new resource to be created.
- **managed_image_id** - (Required) The ID of the Managed Image which should be used for this Shared Image Version.

Changing this forces a new resource to be created.

NOTE: The ID can be sourced from the `azurerm_image` Data Source (<https://www.terraform.io/docs/providers/azurerm/d/image.html>) or Resource (<https://www.terraform.io/docs/providers/azurerm/r/image.html>).

- `target_region` - (Required) One or more `target_region` blocks as documented below.
- `exclude_from_latest` - (Optional) Should this Image Version be excluded from the `latest` filter? If set to `true` this Image Version won't be returned for the `latest` version. Defaults to `false`.
- `tags` - (Optional) A collection of tags which should be applied to this resource.

The `target_region` block exports the following:

- `name` - (Required) The Azure Region in which this Image Version should exist.
- `regional_replica_count` - (Required) The number of replicas of the Image Version to be created per region.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Shared Image Version.

Import

Shared Image Versions can be imported using the `resource id`, e.g.

```
terraform import azurerm_shared_image_version.version /subscriptions/00000000-0000-0000-0000-000000000000  
/resourceGroups/mygroup1/providers/Microsoft.Compute/galleries/gallery1/images/image1/versions/1.2.3
```

azurerm_signalr_service

Manages an Azure SignalR service.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name     = "terraform-signalr"
  location = "West US"
}

resource "azurerm_signalr_service" "example" {
  name                = "tfex-signalr"
  location             = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"

  sku {
    name     = "Free_F1"
    capacity = 1
  }

  cors {
    allowed_origins = ["http://www.example.com"]
  }

  features {
    flag = "ServiceMode"
    value = "Default"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the SignalR service. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the resource group in which to create the SignalR service. Changing this forces a new resource to be created.
 - `location` - (Required) Specifies the supported Azure location where the SignalR service exists. Changing this forces a new resource to be created.
 - `sku` - A `sku` block as documented below.
 - `cors` - (Optional) A `cors` block as documented below.
 - `features` - (Optional) A `features` block as documented below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

A `cors` block supports the following:

- `allowed_origins` - (Required) A list of origins which should be able to make cross-origin calls. `*` can be used to allow all calls.

A `features` block supports the following:

- `flag` - (Required) A kind of feature. At this time the only supported value is `ServiceMode`.
- `value` - (Required) A value of a feature flag. Possible values are `Classic`, `Default` and `Serverless`.

A `sku` block supports the following:

- `name` - (Required) Specifies which tier to use. Valid values are `Free_F1` and `Standard_S1`.
- `capacity` - (Required) Specifies the number of units associated with this SignalR service. Valid values are `1`, `2`, `5`, `10`, `20`, `50` and `100`.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the SignalR service.
- `hostname` - The FQDN of the SignalR service.
- `ip_address` - The publicly accessible IP of the SignalR service.
- `public_port` - The publicly accessible port of the SignalR service which is designed for browser/client use.
- `server_port` - The publicly accessible port of the SignalR service which is designed for customer server side use.
- `primary_access_key` - The primary access key for the SignalR service.
- `primary_connection_string` - The primary connection string for the SignalR service.
- `secondary_access_key` - The secondary access key for the SignalR service.
- `secondary_connection_string` - The secondary connection string for the SignalR service.

Import

SignalR services can be imported using the `resource id`, e.g.

```
terraform import azurerm_signalr_service.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/terraform-signalr/providers/Microsoft.SignalRService/SignalR/tfex-signalr
```

azurerm_snapshot

Manages a Disk Snapshot.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "snapshot-rg"
  location  = "West Europe"
}

resource "azurerm_managed_disk" "test" {
  name                = "managed-disk"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_type = "Standard_LRS"
  create_option       = "Empty"
  disk_size_gb        = "10"
}

resource "azurerm_snapshot" "test" {
  name                = "snapshot"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  create_option       = "Copy"
  source_uri          = "${azurerm_managed_disk.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Snapshot resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Snapshot. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `create_option` - (Required) Indicates how the snapshot is to be created. Possible values are `Copy` or `Import` . Changing this forces a new resource to be created.

Note: One of `source_uri` , `source_resource_id` or `storage_account_id` must be specified.

- `source_uri` - (Optional) Specifies the URI to a Managed or Unmanaged Disk. Changing this forces a new resource to be created.

- `source_resource_id` - (Optional) Specifies a reference to an existing snapshot, when `create_option` is `Copy` . Changing this forces a new resource to be created.
- `storage_account_id` - (Optional) Specifies the ID of an storage account. Used with `source_uri` to allow authorization during import of unmanaged blobs from a different subscription. Changing this forces a new resource to be created.
- `disk_size_gb` - (Optional) The size of the Snapshotted Disk in GB.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Snapshot ID.
- `disk_size_gb` - The Size of the Snapshotted Disk in GB.

Import

Snapshots can be imported using the `resource id` , e.g.

```
terraform import azurerm_snapshot.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Compute/snapshots/snapshot1
```

azurerm_sql_active_directory_administrator

Allows you to set a user or group as the AD administrator for an Azure SQL server

Example Usage

```
data "azurerm_client_config" "current" {}

resource "azurerm_resource_group" "test" {
  name     = "acceptanceTestResourceGroup1"
  location = "West US"
}

resource "azurerm_sql_server" "test" {
  name                        = "mysqlserver"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  version                    = "12.0"
  administrator_login        = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_sql_active_directory_administrator" "test" {
  server_name      = "${azurerm_sql_server.test.name}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  login            = "sqladmin"
  tenant_id        = "${data.azurerm_client_config.current.tenant_id}"
  object_id        = "${data.azurerm_client_config.current.service_principal_object_id}"
}
```

Argument Reference

The following arguments are supported:

- `server_name` - (Required) The name of the SQL Server on which to set the administrator. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group for the SQL server. Changing this forces a new resource to be created.
- `login` - (Required) The login name of the principal to set as the server administrator
- `object_id` - (Required) The ID of the principal to set as the server administrator
- `tenant_id` - (Required) The Azure Tenant ID

Attributes Reference

The following attributes are exported:

- `id` - The SQL Active Directory Administrator ID.

Import

A SQL Active Directory Administrator can be imported using the `resource id`, e.g.

```
terraform import azurerm_sql_active_directory_administrator.administrator /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/administrators/activeDirectory
```

azurerm_sql_database

Allows you to manage an Azure SQL Database

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_sql_server" "test" {
  name                        = "mysqlserver"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "West US"
  version                    = "12.0"
  administrator_login        = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_sql_database" "test" {
  name                = "mysqldatabase"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "West US"
  server_name          = "${azurerm_sql_server.test.name}"

  tags = {
    environment = "production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the database.
- `resource_group_name` - (Required) The name of the resource group in which to create the database. This must be the same as Database Server resource group currently.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server on which to create the database.
- `create_mode` - (Optional) Specifies how to create the database. Must be either `Default` to create a new database or `PointInTimeRestore` to restore from a snapshot. Defaults to `Default`.
- `import` - (Optional) A Database Import block as documented below. `create_mode` must be set to `Default`.
- `source_database_id` - (Optional) The URI of the source database if `create_mode` value is not `Default`.

- `restore_point_in_time` - (Optional) The point in time for the restore. Only applies if `create_mode` is `PointInTimeRestore` e.g. 2013-11-08T22:00:40Z
- `edition` - (Optional) The edition of the database to be created. Applies only if `create_mode` is `Default`. Valid values are: `Basic`, `Standard`, `Premium`, `DataWarehouse`, `Business`, `BusinessCritical`, `Free`, `GeneralPurpose`, `Hyperscale`, `Premium`, `PremiumRS`, `Standard`, `Stretch`, `System`, `System2`, or `Web`. Please see Azure SQL Database Service Tiers (<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-service-tiers/>).
- `collation` - (Optional) The name of the collation. Applies only if `create_mode` is `Default`. Azure default is `SQL_LATIN1_GENERAL_CP1_CI_AS`. Changing this forces a new resource to be created.
- `max_size_bytes` - (Optional) The maximum size that the database can grow to. Applies only if `create_mode` is `Default`. Please see Azure SQL Database Service Tiers (<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-service-tiers/>).
- `requested_service_objective_id` - (Optional) Use `requested_service_objective_id` or `requested_service_objective_name` to set the performance level for the database. Please see Azure SQL Database Service Tiers (<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-service-tiers/>).
- `requested_service_objective_name` - (Optional) Use `requested_service_objective_name` or `requested_service_objective_id` to set the performance level for the database. Valid values are: `S0`, `S1`, `S2`, `S3`, `P1`, `P2`, `P4`, `P6`, `P11` and `ElasticPool`. Please see Azure SQL Database Service Tiers (<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-service-tiers/>).
- `source_database_deletion_date` - (Optional) The deletion date time of the source database. Only applies to deleted databases where `create_mode` is `PointInTimeRestore`.
- `elastic_pool_name` - (Optional) The name of the elastic database pool.
- `threat_detection_policy` - (Optional) Threat detection policy configuration. The `threat_detection_policy` block supports fields documented below.
- `read_scale` - (Optional) Read-only connections will be redirected to a high-available replica. Please see Use read-only replicas to load-balance read-only query workloads (<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-read-scale-out>).
- `tags` - (Optional) A mapping of tags to assign to the resource.

`import` supports the following:

- `storage_uri` - (Required) Specifies the blob URI of the .bacpac file.
- `storage_key` - (Required) Specifies the access key for the storage account.
- `storage_key_type` - (Required) Specifies the type of access key for the storage account. Valid values are `StorageAccessKey` or `SharedAccessKey`.
- `administrator_login` - (Required) Specifies the name of the SQL administrator.
- `administrator_login_password` - (Required) Specifies the password of the SQL administrator.
- `authentication_type` - (Required) Specifies the type of authentication used to access the server. Valid values are `SQL` or `ADPassword`.
- `operation_mode` - (Optional) Specifies the type of import operation being performed. The only allowable value is `Import`.

`threat_detection_policy` supports the following:

- `state` - (Required) The State of the Policy. Possible values are `Enabled` , `Disabled` or `New` .
- `disabled_alerts` - (Optional) Specifies a list of alerts which should be disabled. Possible values include `Access_Anomaly` , `Sql_Injection` and `Sql_Injection_Vulnerability` .
- `email_account_admins` - (Optional) Should the account administrators be emailed when this alert is triggered?
- `email_addresses` - (Optional) A list of email addresses which alerts should be sent to.
- `retention_days` - (Optional) Specifies the number of days to keep in the Threat Detection audit logs.
- `storage_account_access_key` - (Optional) Specifies the identifier key of the Threat Detection audit storage account. Required if `state` is `Enabled` .
- `storage_endpoint` - (Optional) Specifies the blob storage endpoint (e.g. `https://MyAccount.blob.core.windows.net` (`https://MyAccount.blob.core.windows.net`)). This blob storage will hold all Threat Detection audit logs. Required if `state` is `Enabled` .
- `use_server_default` - (Optional) Should the default server policy be used? Defaults to `Disabled` .

Attributes Reference

The following attributes are exported:

- `id` - The SQL Database ID.
- `creation_date` - The creation date of the SQL Database.
- `default_secondary_location` - The default secondary location of the SQL Database.

Import

SQL Databases can be imported using the `resource id` , e.g.

```
terraform import azurerm_sql_database.database1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/databases/database1
```

azurerm_sql_elasticpool

Allows you to manage an Azure SQL Elastic Pool.

NOTE: - This version of the `Elasticpool` resource is being **deprecated** and should no longer be used. Please use the `azurerm_mssql_elasticpool` (/docs/providers/azurerm/r/mssql_elasticpool.html) version instead.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "my-resource-group"
  location  = "West US"
}

resource "azurerm_sql_server" "test" {
  name                        = "my-sql-server" # NOTE: needs to be globally unique
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  version                    = "12.0"
  administrator_login        = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_sql_elasticpool" "test" {
  name                = "test"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"
  server_name          = "${azurerm_sql_server.test.name}"
  edition              = "Basic"
  dtu                  = 50
  db_dtu_min           = 0
  db_dtu_max           = 5
  pool_size            = 5000
}
```

NOTE on `azurerm_sql_elasticpool`: - The values of `edition`, `dtu`, and `pool_size` must be consistent with the Azure SQL Database Service Tiers (<https://docs.microsoft.com/en-gb/azure/sql-database/sql-database-service-tiers#elastic-pool-service-tiers-and-performance-in-edtus>). Any inconsistent argument configuration will be rejected.

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the elastic pool. This needs to be globally unique. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the elastic pool. This must be the same as the resource group of the underlying SQL server.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server on which to create the elastic pool. Changing this forces a new resource to be created.
- `edition` - (Required) The edition of the elastic pool to be created. Valid values are `Basic`, `Standard`, and `Premium`. Refer to Azure SQL Database Service Tiers (<https://docs.microsoft.com/en-gb/azure/sql-database/sql-database-service-tiers#elastic-pool-service-tiers-and-performance-in-edtus>) for details. Changing this forces a new resource to be created.
- `dtu` - (Required) The total shared DTU for the elastic pool. Valid values depend on the `edition` which has been defined. Refer to Azure SQL Database Service Tiers (<https://docs.microsoft.com/en-gb/azure/sql-database/sql-database-service-tiers#elastic-pool-service-tiers-and-performance-in-edtus>) for valid combinations.
- `db_dtu_min` - (Optional) The minimum DTU which will be guaranteed to all databases in the elastic pool to be created.
- `db_dtu_max` - (Optional) The maximum DTU which will be guaranteed to all databases in the elastic pool to be created.
- `pool_size` - (Optional) The maximum size in MB that all databases in the elastic pool can grow to. The maximum size must be consistent with combination of `edition` and `dtu` and the limits documented in Azure SQL Database Service Tiers (<https://docs.microsoft.com/en-gb/azure/sql-database/sql-database-service-tiers#elastic-pool-service-tiers-and-performance-in-edtus>). If not defined when creating an elastic pool, the value is set to the size implied by `edition` and `dtu`.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The SQL Elastic Pool ID.
- `creation_date` - The creation date of the SQL Elastic Pool.

azurerm_sql_failover_group

Create a failover group of databases on a collection of Azure SQL servers.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example"
  location  = "uksouth"
}

resource "azurerm_sql_server" "primary" {
  name                        = "sql-primary"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  version                    = "12.0"
  administrator_login        = "sqladmin"
  administrator_login_password = "pa$$$$word"
}

resource "azurerm_sql_server" "secondary" {
  name                        = "sql-secondary"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "northeurope"
  version                    = "12.0"
  administrator_login        = "sqladmin"
  administrator_login_password = "pa$$$$word"
}

resource "azurerm_sql_database" "db1" {
  name                = "db1"
  resource_group_name = "${azurerm_sql_server.primary.resource_group_name}"
  location            = "${azurerm_sql_server.primary.location}"
  server_name         = "${azurerm_sql_server.primary.name}"
}

resource "azurerm_sql_failover_group" "example" {
  name                = "example-failover-group"
  resource_group_name = "${azurerm_sql_server.primary.resource_group_name}"
  server_name         = "${azurerm_sql_server.primary.name}"
  databases           = ["${azurerm_sql_database.db1.id}"]
  partner_servers {
    id = "${azurerm_sql_server.secondary.id}"
  }
}

read_write_endpoint_failover_policy {
  mode          = "Automatic"
  grace_minutes = 60
}
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the failover group. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group containing the SQL server
- `server_name` - (Required) The name of the primary SQL server. Changing this forces a new resource to be created.
- `databases` - A list of database ids to add to the failover group

NOTE: The failover group will create a secondary database for each database listed in `databases` . If the secondary databases need to be managed through Terraform, they should be defined as resources and a dependency added to the failover group to ensure the secondary databases are created first.

- `partner_servers` - (Required) A list of secondary servers as documented below
- `read_write_endpoint_failover_policy` - (Required) A read/write policy as documented below
- `readonly_endpoint_failover_policy` - (Optional) a read-only policy as documented below
- `tags` - (Optional) A mapping of tags to assign to the resource.

`partner_servers` supports the following:

- `id` - (Required) the SQL server ID

`read_write_endpoint_failover_policy` supports the following:

- `mode` - (Required) the failover mode. Possible values are `Manual` , `Automatic`
- `grace_minutes` - Applies only if `mode` is `Automatic` . The grace period in minutes before failover with data loss is attempted

`readonly_endpoint_failover_policy` supports the following:

- `mode` - Failover policy for the read-only endpoint. Possible values are `Enabled` , and `Disabled`

Attribute Reference

The following attributes are exported:

- `id` - The failover group ID.
- `location` - the location of the failover group.
- `server_name` - the name of the primary SQL Database Server.
- `role` - local replication role of the failover group instance.
- `databases` - list of databases in the failover group.
- `partner_servers` - list of partner server information for the failover group.

Import

SQL Failover Groups can be imported using the `resource id`, e.g.

```
terraform import azurerm_sql_failover_group.example /subscriptions/00000000-0000-0000-0000-000000000000/r
esourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/failovergroups/group1
```

azurerm_sql_firewall_rule

Allows you to manage an Azure SQL Firewall Rule

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_sql_server" "test" {
  name                        = "mysqlserver"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "West US"
  version                    = "12.0"
  administrator_login        = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_sql_firewall_rule" "test" {
  name                = "FirewallRule1"
  resource_group_name = "${azurerm_resource_group.test.name}"
  server_name         = "${azurerm_sql_server.test.name}"
  start_ip_address    = "10.0.17.62"
  end_ip_address      = "10.0.17.62"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the firewall rule.
- `resource_group_name` - (Required) The name of the resource group in which to create the sql server.
- `server_name` - (Required) The name of the SQL Server on which to create the Firewall Rule.
- `start_ip_address` - (Required) The starting IP address to allow through the firewall for this rule.
- `end_ip_address` - (Required) The ending IP address to allow through the firewall for this rule.

NOTE: The Azure feature `Allow access to Azure services` can be enabled by setting `start_ip_address` and `end_ip_address` to `0.0.0.0` which (is documented in the Azure API Docs (<https://docs.microsoft.com/en-us/rest/api/sql/firewallrules/createorupdate>)).

Attributes Reference

The following attributes are exported:

- `id` - The SQL Firewall Rule ID.

Import

SQL Firewall Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_sql_firewall_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/firewallRules/rule1
```

azurerm_sql_server

Manages a SQL Azure Database Server.

Note: All arguments including the administrator login and password will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "database-rg"
  location  = "West US"
}

resource "azurerm_sql_server" "test" {
  name                        = "mysqlserver"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  version                    = "12.0"
  administrator_login        = "mradministrator"
  administrator_login_password = "thisIsDog11"

  tags = {
    environment = "production"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the SQL Server. This needs to be globally unique within Azure.
- `resource_group_name` - (Required) The name of the resource group in which to create the SQL Server.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `version` - (Required) The version for the new server. Valid values are: 2.0 (for v11 server) and 12.0 (for v12 server).
- `administrator_login` - (Required) The administrator login name for the new server. Changing this forces a new resource to be created.
- `administrator_login_password` - (Required) The password associated with the `administrator_login` user. Needs to comply with Azure's Password Policy (<https://msdn.microsoft.com/library/ms161959.aspx>)
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The SQL Server ID.
- `fully_qualified_domain_name` - The fully qualified domain name of the Azure SQL Server (e.g. `myServerName.database.windows.net`)

Import

SQL Servers can be imported using the `resource id`, e.g.

```
terraform import azurerm_sql_server.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver
```

azurerm_sql_virtual_network_rule

Allows you to add, update, or remove an Azure SQL server to a subnet of a virtual network.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-sql-server-vnet-rule"
  location  = "West US"
}

resource "azurerm_virtual_network" "vnet" {
  name            = "example-vnet"
  address_space   = ["10.7.29.0/29"]
  location        = "${azurerm_resource_group.example.location}"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_subnet" "subnet" {
  name                 = "example-subnet"
  resource_group_name = "${azurerm_resource_group.example.name}"
  virtual_network_name = "${azurerm_virtual_network.vnet.name}"
  address_prefix       = "10.7.29.0/29"
  service_endpoints    = ["Microsoft.Sql"]
}

resource "azurerm_sql_server" "sqlserver" {
  name                  = "uniqueazuresqlserver"
  resource_group_name   = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"
  version               = "12.0"
  administrator_login   = "4dm1n157r470r"
  administrator_login_password = "4-v3ry-53cr37-p455w0rd"
}

resource "azurerm_sql_virtual_network_rule" "sqlvnetrule" {
  name            = "sql-vnet-rule"
  resource_group_name = "${azurerm_resource_group.example.name}"
  server_name     = "${azurerm_sql_server.sqlserver.name}"
  subnet_id       = "${azurerm_subnet.subnet.id}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the SQL virtual network rule. Changing this forces a new resource to be created. Cannot be empty and must only contain alphanumeric characters and hyphens. Cannot start with a number, and cannot start or end with a hyphen.

NOTE: name must be between 1-128 characters long and must satisfy all of the requirements below: 1. Contains only alphanumeric and hyphen characters 2. Cannot start with a number or hyphen 3. Cannot end with a hyphen

- `resource_group_name` - (Required) The name of the resource group where the SQL server resides. Changing this forces a new resource to be created.
- `server_name` - (Required) The name of the SQL Server to which this SQL virtual network rule will be applied to. Changing this forces a new resource to be created.
- `subnet_id` - (Required) The ID of the subnet that the SQL server will be connected to.
- `ignore_missing_vnet_service_endpoint` - (Optional) Create the virtual network rule before the subnet has the virtual network service endpoint enabled. The default value is false.

NOTE: If `ignore_missing_vnet_service_endpoint` is false, and the target subnet does not contain the `Microsoft.Sql` endpoint in the `service_endpoints` array, the deployment will fail when it tries to create the SQL virtual network rule.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the SQL virtual network rule.

Import

SQL Virtual Network Rules can be imported using the `resource id`, e.g.

```
terraform import azurerm_sql_virtual_network_rule.rule1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myresourcegroup/providers/Microsoft.Sql/servers/myserver/virtualNetworkRules/vnetrulename
```

azurerm_storage_account

Manages an Azure Storage Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "storageaccountname"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "GRS"

  tags = {
    environment = "staging"
  }
}
```

Example Usage with Network Rules

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "virtnetname"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "subnetname"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix      = "10.0.2.0/24"
  service_endpoints   = ["Microsoft.Sql", "Microsoft.Storage"]
}

resource "azurerm_storage_account" "testsa" {
  name                = "storageaccountname"
  resource_group_name = "${azurerm_resource_group.test.name}"

  location                = "${azurerm_resource_group.test.location}"
  account_tier            = "Standard"
  account_replication_type = "LRS"

  network_rules {
    default_action      = "Deny"
    ip_rules            = ["100.0.0.1"]
    virtual_network_subnet_ids = ["${azurerm_subnet.test.id}"]
  }

  tags = {
    environment = "staging"
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the storage account. Changing this forces a new resource to be created. This must be unique across the entire Azure service, not just within the resource group.
- **resource_group_name** - (Required) The name of the resource group in which to create the storage account. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- **account_kind** - (Optional) Defines the Kind of account. Valid options are `BlobStorage`, `BlockBlobStorage`, `FileStorage`, `Storage` and `StorageV2`. Changing this forces a new resource to be created. Defaults to `Storage`.

- `account_tier` - (Required) Defines the Tier to use for this storage account. Valid options are `Standard` and `Premium`. For `FileStorage` accounts only `Premium` is valid. Changing this forces a new resource to be created.
- `account_replication_type` - (Required) Defines the type of replication to use for this storage account. Valid options are `LRS`, `GRS`, `RAGRS` and `ZRS`.
- `access_tier` - (Optional) Defines the access tier for `BlobStorage`, `FileStorage` and `StorageV2` accounts. Valid options are `Hot` and `Cool`, defaults to `Hot`.
- `enable_blob_encryption` - (Optional) Boolean flag which controls if Encryption Services are enabled for Blob storage, see here (<https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/>) for more information. Defaults to `true`.
- `enable_file_encryption` - (Optional) Boolean flag which controls if Encryption Services are enabled for File storage, see here (<https://azure.microsoft.com/en-us/documentation/articles/storage-service-encryption/>) for more information. Defaults to `true`.
- `enable_https_traffic_only` - (Optional) Boolean flag which forces HTTPS if enabled, see here (<https://docs.microsoft.com/en-us/azure/storage/storage-require-secure-transfer/>) for more information.
- `is_hns_enabled` - (Optional) Is Hierarchical Namespace enabled? This can be used with Azure Data Lake Storage Gen 2 (see here for more information (<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-quickstart-create-account/>)). Changing this forces a new resource to be created.
- `account_encryption_source` - (Optional) The Encryption Source for this Storage Account. Possible values are `Microsoft.Keyvault` and `Microsoft.Storage`. Defaults to `Microsoft.Storage`.
- `custom_domain` - (Optional) A `custom_domain` block as documented below.
- `enable_advanced_threat_protection` (Optional) Boolean flag which controls if advanced threat protection is enabled, see here (<https://docs.microsoft.com/en-us/azure/storage/common/storage-advanced-threat-protection>) for more information. Defaults to `false`.

Note: `enable_advanced_threat_protection` is not supported in all regions.

- `identity` - (Optional) A `identity` block as defined below.
- `queue_properties` - (Optional) A `queue_properties` block as defined below.

NOTE: `queue_properties` cannot be set when the `access_tier` is set to `BlobStorage`

- `network_rules` - (Optional) A `network_rules` block as documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

A `cors_rule` block supports the following:

- `allowed_headers` - (Required) A list of headers that are allowed to be a part of the cross-origin request.
- `allowed_methods` - (Required) A list of http headers that are allowed to be executed by the origin. Valid options are `DELETE`, `GET`, `HEAD`, `MERGE`, `POST`, `OPTIONS` or `PUT`.
- `allowed_origins` - (Required) A list of origin domains that will be allowed by CORS.

- `exposed_headers` - (Required) A list of response headers that are exposed to CORS clients.
 - `max_age_in_seconds` - (Required) The number of seconds the client should cache a preflight response.
-

A `custom_domain` block supports the following:

- `name` - (Optional) The Custom Domain Name to use for the Storage Account, which will be validated by Azure.
 - `use_subdomain` - (Optional) Should the Custom Domain Name be validated by using indirect CNAME validation?
-

A `hour_metrics` block supports the following:

- `enabled` - (Required) Indicates whether hour metrics are enabled for the Queue service. Changing this forces a new resource.
 - `version` - (Required) The version of storage analytics to configure. Changing this forces a new resource.
 - `include_apis` - (Optional) Indicates whether metrics should generate summary statistics for called API operations.
 - `retention_policy_days` - (Optional) Specifies the number of days that logs will be retained. Changing this forces a new resource.
-

A `identity` block supports the following:

- `type` - (Required) Specifies the identity type of the Storage Account. At this time the only allowed value is `SystemAssigned`.

The assigned `principal_id` and `tenant_id` can be retrieved after the `identity type` has been set to `SystemAssigned` and Storage Account has been created. More details are available below.

A `logging` block supports the following:

- `delete` - (Required) Indicates whether all delete requests should be logged. Changing this forces a new resource.
 - `read` - (Required) Indicates whether all read requests should be logged. Changing this forces a new resource.
 - `version` - (Required) The version of storage analytics to configure. Changing this forces a new resource.
 - `write` - (Required) Indicates whether all write requests should be logged. Changing this forces a new resource.
 - `retention_policy_days` - (Optional) Specifies the number of days that logs will be retained. Changing this forces a new resource.
-

A `minute_metrics` block supports the following:

- `enabled` - (Required) Indicates whether minute metrics are enabled for the Queue service. Changing this forces a new resource.
- `version` - (Required) The version of storage analytics to configure. Changing this forces a new resource.
- `include_apis` - (Optional) Indicates whether metrics should generate summary statistics for called API operations.

- `retention_policy_days` - (Optional) Specifies the number of days that logs will be retained. Changing this forces a new resource.

A `network_rules` block supports the following:

- `default_action` - (Required) Specifies the default action of allow or deny when no other rules match. Valid options are `Deny` or `Allow`.
- `bypass` - (Optional) Specifies whether traffic is bypassed for Logging/Metrics/AzureServices. Valid options are any combination of `Logging`, `Metrics`, `AzureServices`, or `None`.
- `ip_rules` - (Optional) List of public IP or IP ranges in CIDR Format. Only IPV4 addresses are allowed. Private IP address ranges (as defined in RFC 1918 (<https://tools.ietf.org/html/rfc1918#section-3>)) are not allowed.
- `virtual_network_subnet_ids` - (Optional) A list of resource ids for subnets.

Note: If specifying `network_rules`, one of either `ip_rules` or `virtual_network_subnet_ids` must be specified and `default_action` must be set to `Deny`.

Note: More information on Validation is available here (<https://docs.microsoft.com/en-gb/azure/storage/blobs/storage-custom-domain-name>)

A `queue_properties` block supports the following:

- `cors_rule` - (Optional) A `cors_rule` block as defined below.
- `logging` - (Optional) A `logging` block as defined below.
- `minute_metrics` - (Optional) A `minute_metrics` block as defined below.
- `hour_metrics` - (Optional) A `hour_metrics` block as defined below.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The storage account Resource ID.
- `primary_location` - The primary location of the storage account.
- `secondary_location` - The secondary location of the storage account.
- `primary_blob_endpoint` - The endpoint URL for blob storage in the primary location.
- `primary_blob_host` - The hostname with port if applicable for blob storage in the primary location.
- `secondary_blob_endpoint` - The endpoint URL for blob storage in the secondary location.
- `secondary_blob_host` - The hostname with port if applicable for blob storage in the secondary location.
- `primary_queue_endpoint` - The endpoint URL for queue storage in the primary location.

- `primary_queue_host` - The hostname with port if applicable for queue storage in the primary location.
- `secondary_queue_endpoint` - The endpoint URL for queue storage in the secondary location.
- `secondary_queue_host` - The hostname with port if applicable for queue storage in the secondary location.
- `primary_table_endpoint` - The endpoint URL for table storage in the primary location.
- `primary_table_host` - The hostname with port if applicable for table storage in the primary location.
- `secondary_table_endpoint` - The endpoint URL for table storage in the secondary location.
- `secondary_table_host` - The hostname with port if applicable for table storage in the secondary location.
- `primary_file_endpoint` - The endpoint URL for file storage in the primary location.
- `primary_file_host` - The hostname with port if applicable for file storage in the primary location.
- `secondary_file_endpoint` - The endpoint URL for file storage in the secondary location.
- `secondary_file_host` - The hostname with port if applicable for file storage in the secondary location.
- `primary_dfs_endpoint` - The endpoint URL for DFS storage in the primary location.
- `primary_dfs_host` - The hostname with port if applicable for DFS storage in the primary location.
- `secondary_dfs_endpoint` - The endpoint URL for DFS storage in the secondary location.
- `secondary_dfs_host` - The hostname with port if applicable for DFS storage in the secondary location.
- `primary_web_endpoint` - The endpoint URL for web storage in the primary location.
- `primary_web_host` - The hostname with port if applicable for web storage in the primary location.
- `secondary_web_endpoint` - The endpoint URL for web storage in the secondary location.
- `secondary_web_host` - The hostname with port if applicable for web storage in the secondary location.
- `primary_access_key` - The primary access key for the storage account.
- `secondary_access_key` - The secondary access key for the storage account.
- `primary_connection_string` - The connection string associated with the primary location.
- `secondary_connection_string` - The connection string associated with the secondary location.
- `primary_blob_connection_string` - The connection string associated with the primary blob location.
- `secondary_blob_connection_string` - The connection string associated with the secondary blob location.

NOTE: If there's a Write Lock on the Storage Account, or the account doesn't have permission then these fields will have an empty value due to a bug in the Azure API (<https://github.com/Azure/azure-rest-api-specs/issues/6363>)

- `identity` - An `identity` block as defined below, which contains the Identity information for this Storage Account.

`identity` exports the following:

- `principal_id` - The Principal ID for the Service Principal associated with the Identity of this Storage Account.

- `tenant_id` - The Tenant ID for the Service Principal associated with the Identity of this Storage Account.

You can access the Principal ID via `${azurerm_storage_account.test.identity.0.principal_id}` and the Tenant ID via `${azurerm_storage_account.test.identity.0.tenant_id}`

Import

Storage Accounts can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_account.storageAcc1 /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/myresourcegroup/providers/Microsoft.Storage/storageAccounts/myaccount
```

azurerm_storage_blob

Manages a Blob within a Storage Container.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "examplestoracc"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_container" "test" {
  name                  = "content"
  resource_group_name   = "${azurerm_resource_group.test.name}"
  storage_account_name  = "${azurerm_storage_account.test.name}"
  container_access_type = "private"
}

resource "azurerm_storage_blob" "test" {
  name                = "my-awesome-content.zip"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
  storage_container_name = "${azurerm_storage_container.test.name}"
  type                = "Block"
  source              = "some-local-file.zip"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the storage blob. Must be unique within the storage container the blob is located.
- `storage_account_name` - (Required) Specifies the storage account in which to create the storage container. Changing this forces a new resource to be created.
- `storage_container_name` - (Required) The name of the storage container in which this blob should be created.
- `type` - (Required) The type of the storage blob to be created. Possible values are `Append`, `Block` or `Page`. Changing this forces a new resource to be created.
- `size` - (Optional) Used only for `page` blobs to specify the size in bytes of the blob to be created. Must be a multiple of 512. Defaults to 0.

- `access_tier` - (Optional) The access tier of the storage blob. Possible values are `Archive`, `Cool` and `Hot`.
- `content_type` - (Optional) The content type of the storage blob. Cannot be defined if `source_uri` is defined. Defaults to `application/octet-stream`.
- `source` - (Optional) An absolute path to a file on the local system. This field cannot be specified for Append blobs and cannot be specified if `source_content` or `source_uri` is specified.
- `source_content` - (Optional) The content for this blob which should be defined inline. This field can only be specified for Block blobs and cannot be specified if `source` or `source_uri` is specified.
- `source_uri` - (Optional) The URI of an existing blob, or a file in the Azure File service, to use as the source contents for the blob to be created. Changing this forces a new resource to be created. This field cannot be specified for Append blobs and cannot be specified if `source` or `source_content` is specified.
- `parallelism` - (Optional) The number of workers per CPU core to run for concurrent uploads. Defaults to `8`.

NOTE: `parallelism` is only applicable for Page blobs - support for Block Blobs is blocked on the upstream issue (<https://github.com/tombuildsstuff/giovanni/issues/15>).

- `metadata` - (Optional) A map of custom blob metadata.
- `attempts` - (Optional / **Deprecated**) The number of attempts to make per page or block when uploading. Defaults to `1`.
- `resource_group_name` - (Optional / **Deprecated**) The name of the resource group in which to create the storage container.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Storage Blob.
- `url` - The URL of the blob

Import

Storage Blob's can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_blob.blob1 https://example.blob.core.windows.net/container/blob.vhd
```

azurerm_storage_container

Manages a Container within an Azure Storage Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "examplestoraccount"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"

  tags = {
    environment = "staging"
  }
}

resource "azurerm_storage_container" "test" {
  name                = "vhds"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
  container_access_type = "private"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Container which should be created within the Storage Account.
- `storage_account_name` - (Required) The name of the Storage Account where the Container should be created.
- `container_access_type` - (Optional) The Access Level configured for this Container. Possible values are `blob` , `container` or `private` . Defaults to `private` .
- `metadata` - (Optional) A mapping of MetaData for this Container.
- `resource_group_name` - (Optional / **Deprecated**) The name of the resource group in which to create the storage container. This field is no longer used and will be removed in 2.0.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Storage Container.
- `has_immutability_policy` - Is there an Immutability Policy configured on this Storage Container?
- `has_legal_hold` - Is there a Legal Hold configured on this Storage Container?
- `properties` - **(Deprecated)** Key-value definition of additional properties associated to the Storage Container

Import

Storage Containers can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_container.container1 https://example.blob.core.windows.net/container
```


azurerm_storage_data_lake_gen2_filesystem

Manages a Data Lake Gen2 File System within an Azure Storage Account.

NOTE: This Resource requires using Azure Active Directory to connect to Azure Storage, which in turn requires the Storage specific roles - which are not granted by default.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "examplestorageacc"
  resource_group_name        = azurerm_resource_group.test.name
  location                   = azurerm_resource_group.test.location
  account_tier               = "Standard"
  account_replication_type   = "LRS"
  account_kind               = "StorageV2"
  is_hns_enabled             = "true"
}

resource "azurerm_storage_data_lake_gen2_filesystem" "test" {
  name                = "example"
  storage_account_id = azurerm_storage_account.test.id

  properties = {
    hello = "aGVsbG8="
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Data Lake Gen2 File System which should be created within the Storage Account. Must be unique within the storage account the queue is located. Changing this forces a new resource to be created.
- `storage_account_id` - (Required) Specifies the ID of the Storage Account in which the Data Lake Gen2 File System should exist. Changing this forces a new resource to be created.

NOTE: The Storage Account requires `account_kind` to be either `StorageV2` or `BlobStorage`. In addition, `is_hns_enabled` has to be set to `true`.

- `properties` - (Optional) A mapping of Key to Base64-Encoded Values which should be assigned to this Data Lake Gen2 File System.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Data Lake Gen2 File System.

Import

Data Lake Gen2 File System's can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_data_lake_gen2_filesystem.queue1 https://account1.dfs.core.windows.net/fileSystem1
```

azurerm_storage_management_policy

Manages an Azure Storage Account Management Policy.

Example Usage

```

resource "azurerm_resource_group" "testrg" {
  name      = "resourceGroupName"
  location  = "westus"
}

resource "azurerm_storage_account" "testsa" {
  name                = "storageaccountname"
  resource_group_name = "${azurerm_resource_group.testrg.name}"

  location                = "${azurerm_resource_group.testrg.location}"
  account_tier            = "Standard"
  account_replication_type = "LRS"
  account_kind            = "BlobStorage"
}

resource "azurerm_storage_management_policy" "testpolicy" {
  storage_account_id = "${azurerm_storage_account.testsa.id}"

  rule {
    name      = "rule1"
    enabled   = true
    filters {
      prefix_match = ["container1/prefix1"]
      blob_types   = ["blockBlob"]
    }
    actions {
      base_blob {
        tier_to_cool_after_days_since_modification_greater_than = 10
        tier_to_archive_after_days_since_modification_greater_than = 50
        delete_after_days_since_modification_greater_than        = 100
      }
      snapshot {
        delete_after_days_since_creation_greater_than = 30
      }
    }
  }
}

rule {
  name      = "rule2"
  enabled   = false
  filters {
    prefix_match = ["container2/prefix1", "container2/prefix2"]
    blob_types   = ["blockBlob"]
  }
  actions {
    base_blob {
      tier_to_cool_after_days_since_modification_greater_than = 11
      tier_to_archive_after_days_since_modification_greater_than = 51
      delete_after_days_since_modification_greater_than        = 101
    }
    snapshot {
      delete_after_days_since_creation_greater_than = 31
    }
  }
}

```

Argument Reference

The following arguments are supported:

- `storage_account_id` - (Required) Specifies the id of the storage account to apply the management policy to.
- `rule` - (Optional) A `rule` block as documented below.

-
- `rule` supports the following:
 - `name` - (Required) A rule name can contain any combination of alpha numeric characters. Rule name is case-sensitive. It must be unique within a policy.
 - `enabled` - (Required) Boolean to specify whether the rule is enabled.
 - `filters` - A `filter` block as documented below.
 - `actions` - An `actions` block as documented below.

`filters` supports the following:

- `prefix_match` - An array of strings for prefixes to be matched.
- `blob_types` - An array of predefined values. Only `blockBlob` is supported.

`actions` supports the following:

- `base_blob` - A `base_blob` block as documented below.
- `snapshot` - A `snapshot` block as documented below.

`base_blob` supports the following:

- `tier_to_cool_after_days_since_modification_greater_than` - The age in days after last modification to tier blobs to cool storage. Supports blob currently at Hot tier. Must be at least 0.
- `tier_to_archive_after_days_since_modification_greater_than` - The age in days after last modification to tier blobs to archive storage. Supports blob currently at Hot or Cool tier. Must be at least 0.
- `delete_after_days_since_modification_greater_than` - The age in days after last modification to delete the blob. Must be at least 0.

`snapshot` supports the following:

- `delete_after_days_since_creation_greater_than` - The age in days after create to delete the snapshot. Must be at least 0.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Storage Account Management Policy.

Import

Storage Account Management Policies can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_management_policy.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Storage/storageAccounts/myaccountname/managementPolicies/default
```

azurerm_storage_queue

Manages a Queue within an Azure Storage Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "example-resources"
  location = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "examplestorageacc"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_queue" "test" {
  name                = "mysamplequeue"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Queue which should be created within the Storage Account. Must be unique within the storage account the queue is located.
- `storage_account_name` - (Required) Specifies the Storage Account in which the Storage Queue should exist. Changing this forces a new resource to be created.
- `resource_group_name` - (Optional / **Deprecated**) The name of the resource group in which to create the storage queue.
- `metadata` - (Optional) A mapping of MetaData which should be assigned to this Storage Queue.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Storage Queue.

Import

Storage Queue's can be imported using the `resource_id`, e.g.

```
terraform import azurerm_storage_queue.queue1 https://example.queue.core.windows.net/queue1
```


azurerm_storage_share_directory

Manages a Directory within an Azure Storage File Share.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "azuretest"
  location = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "azureteststorage"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_share" "test" {
  name                = "sharename"
  storage_account_name = "${azurerm_storage_account.test.name}"
  quota               = 50
}

resource "azurerm_storage_share_directory" "test" {
  name            = "example"
  share_name      = "${azurerm_storage_share.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name (or path) of the Directory that should be created within this File Share. Changing this forces a new resource to be created.
- **share_name** - (Required) The name of the File Share where this Directory should be created. Changing this forces a new resource to be created.
- **storage_account_name** - (Required) The name of the Storage Account within which the File Share is located. Changing this forces a new resource to be created.
- **metadata** - (Optional) A mapping of metadata to assign to this Directory.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Directory within the File Share.

Import

Directories within an Azure Storage File Share can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_share_directory.test https://tomdevsa20.file.core.windows.net/share1/directory1
```

azurerm_storage_share

Manages a File Share within Azure Storage.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "azuretest"
  location = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "azureteststorage"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_share" "testshare" {
  name                = "sharename"
  storage_account_name = "${azurerm_storage_account.test.name}"
  quota               = 50
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the share. Must be unique within the storage account where the share is located.
- `storage_account_name` - (Required) Specifies the storage account in which to create the share. Changing this forces a new resource to be created.
- `acl` - (Optional) One or more `acl` blocks as defined below.
- `quota` - (Optional) The maximum size of the share, in gigabytes. Must be greater than 0, and less than or equal to 5 TB (5120 GB) for Standard storage accounts or 100 TB (102400 GB) for Premium storage accounts. Default is 5120.
- `metadata` - (Optional) A mapping of MetaData for this File Share.
- `resource_group_name` - (Optional / **Deprecated**) The name of the resource group in which to create the share. Changing this forces a new resource to be created.

A `acl` block supports the following:

- `id` - (Required) The ID which should be used for this Shared Identifier.
 - `access_policy` - (Required) An `access_policy` block as defined below.
-

A `access_policy` block supports the following:

- `expiry` - (Required) The ISO8061 UTC time at which this Access Policy should be valid until.
- `permissions` - (Required) The permissions which should associated with this Shared Identifier.
- `start` - (Required) The ISO8061 UTC time at which this Access Policy should be valid from.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the File Share.
- `url` - The URL of the File Share

Import

Storage Shares can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_share.testShare https://account1.file.core.windows.net/share1
```

azurerm_storage_table_entity

Manages an Entity within a Table in an Azure Storage Account.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "azureexample"
  location  = "westus"
}

resource "azurerm_storage_account" "example" {
  name                        = "azureexamplestorage1"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                   = "${azurerm_resource_group.example.location}"
  account_tier                = "Standard"
  account_replication_type    = "LRS"
}

resource "azurerm_storage_table" "example" {
  name                      = "myexampletable"
  resource_group_name       = "${azurerm_resource_group.example.name}"
  storage_account_name      = "${azurerm_storage_account.example.name}"
}

resource "azurerm_storage_table_entity" "example" {
  storage_account_name = "${azurerm_storage_account.example.name}"
  table_name           = "${azurerm_storage_table.example.name}"

  partition_key = "examplepartition"
  row_key       = "examamplerow"

  entity = {
    example = "sample"
  }
}
```

Argument Reference

The following arguments are supported:

- `storage_account_name` - (Required) Specifies the storage account in which to create the storage table entity. Changing this forces a new resource to be created.
- `table_name` - (Required) The name of the storage table in which to create the storage table entity. Changing this forces a new resource to be created.
- `partition_key` - (Required) The key for the partition where the entity will be inserted/merged. Changing this forces a new resource.

- `row_key` - (Required) The key for the row where the entity will be inserted/merged. Changing this forces a new resource.
- `entity` - (Required) A map of key/value pairs that describe the entity to be inserted/merged in to the storage table.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Entity within the Table in the Storage Account.

Import

Entities within a Table in an Azure Storage Account can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_table_entity.entity1 https://example.table.core.windows.net/table1(PartitionKey='samplepartition',RowKey='samplerow')
```

azurerm_storage_table

Manages a Table within an Azure Storage Account.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "azuretest"
  location = "West Europe"
}

resource "azurerm_storage_account" "test" {
  name                        = "azureteststorage1"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
  account_tier               = "Standard"
  account_replication_type   = "LRS"
}

resource "azurerm_storage_table" "test" {
  name                = "mysampletable"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the storage table. Must be unique within the storage account the table is located.
- `storage_account_name` - (Required) Specifies the storage account in which to create the storage table. Changing this forces a new resource to be created.
- `resource_group_name` - (Optional / **Deprecated**) The name of the resource group in which to create the storage table.
- `acl` - (Optional) One or more `acl` blocks as defined below.

A `acl` block supports the following:

- `id` - (Required) The ID which should be used for this Shared Identifier.
- `access_policy` - (Required) An `access_policy` block as defined below.

A `access_policy` block supports the following:

- `expiry` - (Required) The ISO8061 UTC time at which this Access Policy should be valid until.

- `permissions` - (Required) The permissions which should associated with this Shared Identifier.
- `start` - (Required) The ISO8061 UTC time at which this Access Policy should be valid from.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Table within the Storage Account.

Import

Table's within a Storage Account can be imported using the `resource id`, e.g.

```
terraform import azurerm_storage_table.table1 "https://example.table.core.windows.net/Tables('replace-with-table-name')"
```


azurerm_stream_analytics_function_javascript_udf

Manages a JavaScript UDF Function within Stream Analytics Streaming Job.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_stream_analytics_function_javascript_udf" "test" {
  name = "example-javascript-function"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name = "${data.azurerm_stream_analytics_job.example.resource_group_name}"

  script = <<SCRIPT
function getRandomNumber(in) {
  return in;
}
SCRIPT

  input {
    type = "bigint"
  }

  output {
    type = "bigint"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the JavaScript UDF Function. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job where this Function should be created. Changing this forces a new resource to be created.
- `input` - (Required) One or more `input` blocks as defined below.
- `output` - (Required) An `output` blocks as defined below.

- `script` - (Required) The JavaScript of this UDF Function.
-

A `input` block supports the following:

- `type` - The Data Type for the Input Argument of this JavaScript Function. Possible values include `array`, `any`, `bigint`, `datetime`, `float`, `nvarchar(max)` and `record`.
-

A `output` block supports the following:

- `type` - The Data Type output from this JavaScript Function. Possible values include `array`, `any`, `bigint`, `datetime`, `float`, `nvarchar(max)` and `record`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics JavaScript UDF Function.

Import

Stream Analytics JavaScript UDF Functions can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_output_blob.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_job

Manages a Stream Analytics Job.

Example Usage

```
resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_stream_analytics_job" "example" {
  name                                = "example-job"
  resource_group_name                 = "${azurerm_resource_group.example.name}"
  location                           = "${azurerm_resource_group.example.location}"
  compatibility_level                 = "1.1"
  data_locale                         = "en-GB"
  events_late_arrival_max_delay_in_seconds = 60
  events_out_of_order_max_delay_in_seconds = 50
  events_out_of_order_policy          = "Adjust"
  output_error_policy                 = "Drop"
  streaming_units                     = 3

  tags = {
    environment = "Example"
  }

  transformation_query = <<QUERY
    SELECT *
    INTO [YourOutputAlias]
    FROM [YourInputAlias]
  QUERY
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the Resource Group where the Stream Analytics Job should exist. Changing this forces a new resource to be created.
- **location** - (Required) The Azure Region in which the Resource Group exists. Changing this forces a new resource to be created.
- **compatibility_level** - (Required) Specifies the compatibility level for this job - which controls certain runtime behaviors of the streaming job. Possible values are `1.0` and `1.1`.

NOTE: Support for Compatibility Level 1.2 is dependent on a new version of the Stream Analytics API, which being tracked in this issue (<https://github.com/Azure/azure-rest-api-specs/issues/5604>).

- `data_locale` - (Optional) Specifies the Data Locale of the Job, which should be a supported .NET Culture ([https://msdn.microsoft.com/en-us/library/system.globalization.culturetypes\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/system.globalization.culturetypes(v=vs.110).aspx)).
- `events_late_arrival_max_delay_in_seconds` - (Optional) Specifies the maximum tolerable delay in seconds where events arriving late could be included. Supported range is `-1` (indefinite) to `1814399` (20d 23h 59m 59s). Default is `0`.
- `events_out_of_order_max_delay_in_seconds` - (Optional) Specifies the maximum tolerable delay in seconds where out-of-order events can be adjusted to be back in order. Supported range is `0` to `599` (9m 59s). Default is `5`.
- `events_out_of_order_policy` - (Optional) Specifies the policy which should be applied to events which arrive out of order in the input event stream. Possible values are `Adjust` and `Drop`. Default is `Adjust`.
- `output_error_policy` - (Optional) Specifies the policy which should be applied to events which arrive at the output and cannot be written to the external storage due to being malformed (such as missing column values, column values of wrong type or size). Possible values are `Drop` and `Stop`. Default is `Drop`.
- `streaming_units` - (Required) Specifies the number of streaming units that the streaming job uses. Supported values are `1`, `3`, `6` and multiples of `6` up to `120`.
- `transformation_query` - (Required) Specifies the query that will be run in the streaming job, written in Stream Analytics Query Language (SAQL) (<https://msdn.microsoft.com/library/azure/dn834998>).
- `tags` - A mapping of tags assigned to the resource.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Job.
- `job_id` - The Job ID assigned by the Stream Analytics Job.

Import

Stream Analytics Job's can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_job.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1
```

azurerm_stream_analytics_output_blob

Manages a Stream Analytics Output to Blob Storage.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_storage_account" "example" {
  name = "examplesa"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  location = "${data.azurerm_resource_group.example.location}"
  account_tier = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_storage_container" "example" {
  name = "example"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  storage_account_name = "${azurerm_storage_account.example.name}"
  container_access_type = "private"
}

resource "azurerm_stream_analytics_output_blob" "example" {
  name = "output-to-blob-storage"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  storage_account_name = "${azurerm_storage_account.example.name}"
  storage_account_key = "${azurerm_storage_account.example.primary_access_key}"
  storage_container_name = "${azurerm_storage_container.example.name}"
  path_pattern = "some-pattern"
  date_format = "yyyy-MM-dd"
  time_format = "HH"

  serialization {
    type = "Csv"
    encoding = "UTF8"
    field_delimiter = ","
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Output. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `date_format` - (Required) The date format. Wherever `{date}` appears in `path_pattern`, the value of this property is used as the date format instead.
- `path_pattern` - (Required) The blob path pattern. Not a regular expression. It represents a pattern against which blob names will be matched to determine whether or not they should be included as input or output to the job.
- `storage_account_name` - (Required) The name of the Storage Account.
- `storage_account_key` - (Required) The Access Key which should be used to connect to this Storage Account.
- `storage_container_name` - (Required) The name of the Container within the Storage Account.
- `time_format` - (Required) The time format. Wherever `{time}` appears in `path_pattern`, the value of this property is used as the time format instead.
- `serialization` - (Required) A `serialization` block as defined below.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for outgoing data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

- `format` - (Optional) Specifies the format of the JSON the output will be written in. Possible values are `Array` and `LineSeparated`.

NOTE: This is Required and can only be specified when `type` is set to `Json`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Output Blob Storage.

Import

Stream Analytics Outputs to Blob Storage can be imported using the `resource_id`, e.g.

```
terraform import azurerm_stream_analytics_output_blob.test /subscriptions/00000000-0000-0000-0000-00000000  
000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_output_eventhub

Manages a Stream Analytics Output to an EventHub.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name                = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_eventhub_namespace" "test" {
  name                = "example-ehnamespace"
  location            = "${data.azurerm_resource_group.example.location}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  sku                 = "Standard"
  capacity            = 1
  kafka_enabled       = false
}

resource "azurerm_eventhub" "test" {
  name                = "example-eventhub"
  namespace_name      = "${azurerm_eventhub_namespace.example.name}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  partition_count     = 2
  message_retention   = 1
}

resource "azurerm_stream_analytics_output_eventhub" "example" {
  name                = "output-to-eventhub"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name   = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  eventhub_name         = "${azurerm_eventhub.example.name}"
  servicebus_namespace = "${azurerm_eventhub_namespace.example.name}"
  shared_access_policy_key = "${azurerm_eventhub_namespace.example.default_primary_key}"
  shared_access_policy_name = "RootManageSharedAccessKey"

  serialization {
    type = "Avro"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Output. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `eventhub_name` - (Required) The name of the Event Hub.
- `servicebus_namespace` - (Required) The namespace that is associated with the desired Event Hub, Service Bus Queue, Service Bus Topic, etc.
- `shared_access_policy_key` - (Required) The shared access policy key for the specified shared access policy.
- `shared_access_policy_name` - (Required) The shared access policy name for the Event Hub, Service Bus Queue, Service Bus Topic, etc.
- `serialization` - (Required) A `serialization` block as defined below.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for outgoing data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

- `format` - (Optional) Specifies the format of the JSON the output will be written in. Possible values are `Array` and `LineSeparated`.

NOTE: This is Required and can only be specified when `type` is set to `Json`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Output EventHub.

Import

Stream Analytics Outputs to an EventHub can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_output_eventhub.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_output_mssql

Manages a Stream Analytics Output to Microsoft SQL Server Database.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_sql_server" "example" {
  name = "example-server"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location = "${azurerm_resource_group.example.location}"
  version = "12.0"
  administrator_login = "dbadmin"
  administrator_login_password = "example-password"
}

resource "azurerm_sql_database" "example" {
  name = "exampledb"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location = "${azurerm_resource_group.example.location}"
  server_name = "${azurerm_sql_server.example.name}"
  requested_service_objective_name = "S0"
  collation = "SQL_LATIN1_GENERAL_CP1_CI_AS"
  max_size_bytes = "268435456000"
  create_mode = "Default"
}

resource "azurerm_stream_analytics_output_mssql" "example" {
  name = "example-output-sql"
  stream_analytics_job_name = "${azurerm_stream_analytics_job.example.name}"
  resource_group_name = "${azurerm_stream_analytics_job.example.resource_group_name}"

  server = "${azurerm_sql_server.example.fully_qualified_domain_name}"
  user = "${azurerm_sql_server.example.administrator_login}"
  password = "${azurerm_sql_server.example.administrator_login_password}"
  database = "${azurerm_sql_database.example.name}"
  table = "ExampleTable"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Output. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `server` - (Required) The SQL server url. Changing this forces a new resource to be created.
- `user` - (Required) Username used to login to the Microsoft SQL Server. Changing this forces a new resource to be created.
- `password` - (Required) Password used together with username, to login to the Microsoft SQL Server. Changing this forces a new resource to be created.
- `table` - (Required) Table in the database that the output points to. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Output Microsoft SQL Server Database.

Import

Stream Analytics Outputs to Microsoft SQL Server Database can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_output_mssql.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_output_servicebus_queue

Manages a Stream Analytics Output to a ServiceBus Queue.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name                = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "example-namespace"
  location            = "${data.azurerm_resource_group.example.location}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  sku                 = "Standard"
}

resource "azurerm_servicebus_queue" "example" {
  name                = "example-queue"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  enable_partitioning = true
}

resource "azurerm_stream_analytics_output_servicebus_queue" "test" {
  name                = "blob-storage-output"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name   = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  queue_name            = "${azurerm_servicebus_queue.example.name}"
  servicebus_namespace = "${azurerm_servicebus_namespace.example.name}"
  shared_access_policy_key = "${azurerm_servicebus_namespace.example.default_primary_key}"
  shared_access_policy_name = "RootManageSharedAccessKey"

  serialization {
    format = "Avro"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Stream Output. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.

- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
 - `queue_name` - (Required) The name of the Service Bus Queue.
 - `servicebus_namespace` - (Required) The namespace that is associated with the desired Event Hub, Service Bus Queue, Service Bus Topic, etc.
 - `shared_access_policy_key` - (Required) The shared access policy key for the specified shared access policy.
 - `shared_access_policy_name` - (Required) The shared access policy name for the Event Hub, Service Bus Queue, Service Bus Topic, etc.
 - `serialization` - (Required) A `serialization` block as defined below.
-

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for outgoing data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

- `format` - (Optional) Specifies the format of the JSON the output will be written in. Possible values are `Array` and `LineSeparated`.

NOTE: This is Required and can only be specified when `type` is set to `Json`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Output ServiceBus Queue.

Import

Stream Analytics Output ServiceBus Queue's can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_output_servicebus_queue.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_output_servicebus_topic

Manages a Stream Analytics Output to a ServiceBus Topic.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name                = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_servicebus_namespace" "example" {
  name                = "example-namespace"
  location            = "${data.azurerm_resource_group.example.location}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  sku                 = "Standard"
}

resource "azurerm_servicebus_topic" "example" {
  name                = "example-topic"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  namespace_name      = "${azurerm_servicebus_namespace.example.name}"
  enable_partitioning = true
}

resource "azurerm_stream_analytics_output_servicebus_topic" "test" {
  name                = "blob-storage-output"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name   = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  topic_name            = "${azurerm_servicebus_topic.example.name}"
  servicebus_namespace = "${azurerm_servicebus_namespace.example.name}"
  shared_access_policy_key = "${azurerm_servicebus_namespace.example.default_primary_key}"
  shared_access_policy_name = "RootManageSharedAccessKey"

  serialization {
    format = "Avro"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the Stream Output. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.

- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `topic_name` - (Required) The name of the Service Bus Topic.
- `servicebus_namespace` - (Required) The namespace that is associated with the desired Event Hub, Service Bus Topic, Service Bus Topic, etc.
- `shared_access_policy_key` - (Required) The shared access policy key for the specified shared access policy.
- `shared_access_policy_name` - (Required) The shared access policy name for the Event Hub, Service Bus Queue, Service Bus Topic, etc.
- `serialization` - (Required) A `serialization` block as defined below.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for outgoing data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

- `format` - (Optional) Specifies the format of the JSON the output will be written in. Possible values are `Array` and `LineSeparated`.

NOTE: This is Required and can only be specified when `type` is set to `Json`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Output ServiceBus Topic.

Import

Stream Analytics Output ServiceBus Topic's can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_output_servicebus_topic.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/outputs/output1
```

azurerm_stream_analytics_stream_input_blob

Manages a Stream Analytics Stream Input Blob.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_storage_account" "example" {
  name = "examplestoracc"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location = "${azurerm_resource_group.example.location}"
  account_tier = "Standard"
  account_replication_type = "LRS"
}

resource "azurerm_storage_container" "example" {
  name = "example"
  resource_group_name = "${azurerm_resource_group.example.name}"
  storage_account_name = "${azurerm_storage_account.example.name}"
  container_access_type = "private"
}

resource "azurerm_stream_analytics_stream_input_blob" "test" {
  name = "eventhub-stream-input"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  storage_account_name = "${azurerm_storage_account.example.name}"
  storage_account_key = "${azurerm_storage_account.example.primary_access_key}"
  storage_container_name = "${azurerm_storage_container.example.name}"
  path_pattern = "some-random-pattern"
  date_format = "yyyy/MM/dd"
  time_format = "HH"

  serialization {
    type = "Json"
    encoding = "UTF8"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Input Blob. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `date_format` - (Required) The date format. Wherever `{date}` appears in `path_pattern`, the value of this property is used as the date format instead.
- `path_pattern` - (Required) The blob path pattern. Not a regular expression. It represents a pattern against which blob names will be matched to determine whether or not they should be included as input or output to the job.
- `storage_account_name` - (Required) The name of the Storage Account.
- `storage_account_key` - (Required) The Access Key which should be used to connect to this Storage Account.
- `storage_container_name` - (Required) The name of the Container within the Storage Account.
- `time_format` - (Required) The time format. Wherever `{time}` appears in `path_pattern`, the value of this property is used as the time format instead.
- `serialization` - (Required) A `serialization` block as defined below.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for incoming data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Stream Input Blob.

Import

Stream Analytics Stream Input Blob's can be imported using the `resource_id`, e.g.

```
terraform import azurerm_stream_analytics_stream_input_blob.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/inputs/input1
```

azurerm_stream_analytics_stream_input_eventhub

Manages a Stream Analytics Stream Input EventHub.

Example Usage

```

data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name                = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_eventhub_namespace" "example" {
  name                = "example-namespace"
  location            = "${data.azurerm_resource_group.example.location}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  sku                 = "Standard"
  capacity            = 1
  kafka_enabled       = false
}

resource "azurerm_eventhub" "example" {
  name                = "example-eventhub"
  namespace_name      = "${azurerm_eventhub_namespace.example.name}"
  resource_group_name = "${data.azurerm_resource_group.example.name}"
  partition_count     = 2
  message_retention   = 1
}

resource "azurerm_eventhub_consumer_group" "test" {
  name                = "example-consumergroup"
  namespace_name      = "${azurerm_eventhub_namespace.test.name}"
  eventhub_name       = "${azurerm_eventhub.test.name}"
  resource_group_name = "${data.azurerm_resource_group.test.name}"
}

resource "azurerm_stream_analytics_stream_input_eventhub" "test" {
  name                = "eventhub-stream-input"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  eventhub_consumer_group_name = "${azurerm_eventhub_consumer_group.example.name}"
  eventhub_name       = "${azurerm_eventhub.example.name}"
  servicebus_namespace = "${azurerm_eventhub_namespace.example.name}"
  shared_access_policy_key = "${azurerm_eventhub_namespace.example.default_primary_key}"
  shared_access_policy_name = "RootManageSharedAccessKey"

  serialization {
    type      = "Json"
    encoding = "UTF8"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Input EventHub. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.
- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `eventhub_consumer_group_name` - (Required) The name of an Event Hub Consumer Group that should be used to read events from the Event Hub. Specifying distinct consumer group names for multiple inputs allows each of those inputs to receive the same events from the Event Hub.
- `eventhub_name` - (Required) The name of the Event Hub.
- `servicebus_namespace` - (Required) The namespace that is associated with the desired Event Hub, Service Bus Queue, Service Bus Topic, etc.
- `shared_access_policy_key` - (Required) The shared access policy key for the specified shared access policy.
- `shared_access_policy_name` - (Required) The shared access policy name for the Event Hub, Service Bus Queue, Service Bus Topic, etc.
- `serialization` - (Required) A `serialization` block as defined below.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for incoming data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Stream Input EventHub.

Import

Stream Analytics Stream Input EventHub's can be imported using the `resource_id`, e.g.


```
terraform import azurerm_stream_analytics_stream_input_eventhub.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/inputs/input1
```

azurerm_stream_analytics_stream_input_iothub

Manages a Stream Analytics Stream Input IoT Hub.

Example Usage

```
data "azurerm_resource_group" "example" {
  name = "example-resources"
}

data "azurerm_stream_analytics_job" "example" {
  name                  = "example-job"
  resource_group_name = "${azurerm_resource_group.example.name}"
}

resource "azurerm_iothub" "test" {
  name                  = "example-iothub"
  resource_group_name = "${azurerm_resource_group.example.name}"
  location              = "${azurerm_resource_group.example.location}"

  sku {
    name      = "S1"
    tier       = "Standard"
    capacity  = "1"
  }
}

resource "azurerm_stream_analytics_stream_input_iothub" "test" {
  name                  = "example-iothub-input"
  stream_analytics_job_name = "${data.azurerm_stream_analytics_job.example.name}"
  resource_group_name    = "${data.azurerm_stream_analytics_job.example.resource_group_name}"
  endpoint               = "messages/events"
  eventhub_consumer_group_name = "$Default"
  iothub_namespace       = "${azurerm_iothub.example.name}"
  shared_access_policy_key = "${azurerm_iothub.example.shared_access_policy.0.primary_key}"
  shared_access_policy_name = "iothubowner"

  serialization {
    type      = "Json"
    encoding  = "UTF8"
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Stream Input IoT Hub. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group where the Stream Analytics Job exists. Changing this forces a new resource to be created.

- `stream_analytics_job_name` - (Required) The name of the Stream Analytics Job. Changing this forces a new resource to be created.
- `eventhub_consumer_group_name` - (Required) The name of an Event Hub Consumer Group that should be used to read events from the Event Hub. Specifying distinct consumer group names for multiple inputs allows each of those inputs to receive the same events from the Event Hub.
- `endpoint` - (Required) The IoT Hub endpoint to connect to (ie. `messages/events`, `messages/operationsMonitoringEvents`, etc.).
- `iothub_namespace` - (Required) The name or the URI of the IoT Hub.
- `serialization` - (Required) A `serialization` block as defined below.
- `shared_access_policy_key` - (Required) The shared access policy key for the specified shared access policy.
- `shared_access_policy_name` - (Required) The shared access policy name for the Event Hub, Service Bus Queue, Service Bus Topic, etc.

A `serialization` block supports the following:

- `type` - (Required) The serialization format used for incoming data streams. Possible values are `Avro`, `Csv` and `Json`.
- `encoding` - (Optional) The encoding of the incoming data in the case of input and the encoding of outgoing data in the case of output. It currently can only be set to `UTF8`.

NOTE: This is required when `type` is set to `Csv` or `Json`.

- `field_delimiter` - (Optional) The delimiter that will be used to separate comma-separated value (CSV) records. Possible values are (space), `,` (comma), (tab), `|` (pipe) and `;`.

NOTE: This is required when `type` is set to `Csv`.

Attributes Reference

The following attributes are exported in addition to the arguments listed above:

- `id` - The ID of the Stream Analytics Stream Input IoT Hub.

Import

Stream Analytics Stream Input IoT Hub's can be imported using the `resource id`, e.g.

```
terraform import azurerm_stream_analytics_stream_input_iothub.test /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/group1/providers/Microsoft.StreamAnalytics/streamingjobs/job1/inputs/input1
```

azurerm_subnet

Manages a subnet. Subnets represent network segments within the IP space defined by the virtual network.

NOTE on Virtual Networks and Subnet's: Terraform currently provides both a standalone Subnet resource (</docs/providers/azurerm/r/subnet.html>), and allows for Subnets to be defined in-line within the Virtual Network resource ([/docs/providers/azurerm/r/virtual_network.html](/docs/providers/azurerm/r/virtual_network.html)). At this time you cannot use a Virtual Network with in-line Subnets in conjunction with any Subnet resources. Doing so will cause a conflict of Subnet configurations and will overwrite Subnet's.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_virtual_network" "test" {
  name            = "acceptanceTestVirtualNetwork1"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "testsubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"

  delegation {
    name = "acctestdelegation"

    service_delegation {
      name  = "Microsoft.ContainerInstance/containerGroups"
      actions = ["Microsoft.Network/virtualNetworks/subnets/join/action", "Microsoft.Network/virtualNetworks/subnets/prepareNetworkPolicies/action"]
    }
  }
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the subnet. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the subnet. Changing this forces a new resource to be created.

- `virtual_network_name` - (Required) The name of the virtual network to which to attach the subnet. Changing this forces a new resource to be created.
- `address_prefix` - (Required) The address prefix to use for the subnet.
- `network_security_group_id` - (Optional / **Deprecated**) The ID of the Network Security Group to associate with the subnet.

NOTE: At this time Subnet <-> Network Security Group associations need to be configured both using this field (which is now Deprecated) and using the `azurerm_subnet_network_security_group_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `route_table_id` - (Optional / **Deprecated**) The ID of the Route Table to associate with the subnet.

NOTE: At this time Subnet <-> Route Table associations need to be configured both using this field (which is now Deprecated) and using the `azurerm_subnet_route_table_association` resource. This field is deprecated and will be removed in favour of that resource in the next major version (2.0) of the AzureRM Provider.

- `service_endpoints` - (Optional) The list of Service endpoints to associate with the subnet. Possible values include: `Microsoft.AzureActiveDirectory`, `Microsoft.AzureCosmosDB`, `Microsoft.ContainerRegistry`, `Microsoft.EventHub`, `Microsoft.KeyVault`, `Microsoft.ServiceBus`, `Microsoft.Sql`, `Microsoft.Storage` and `Microsoft.Web`.
- `delegation` - (Optional) One or more `delegation` blocks as defined below.

A `delegation` block supports the following:

- `name` (Required) A name for this delegation.
- `service_delegation` (Required) A `service_delegation` block as defined below.

A `service_delegation` block supports the following:

NOTE: Delegating to services may not be available in all regions. Check that the service you are delegating to is available in your region using the Azure CLI (<https://docs.microsoft.com/en-us/cli/azure/network/vnet/subnet?view=azure-cli-latest#az-network-vnet-subnet-list-available-delegations>). Also, `actions` is specific to each service type. The exact list of `actions` needs to be retrieved using the aforementioned Azure CLI (<https://docs.microsoft.com/en-us/cli/azure/network/vnet/subnet?view=azure-cli-latest#az-network-vnet-subnet-list-available-delegations>).

- `name` - (Required) The name of service to delegate to. Possible values include `Microsoft.BareMetal/AzureVMware`, `Microsoft.BareMetal/CrayServers`, `Microsoft.Batch/batchAccounts`, `Microsoft.ContainerInstance/containerGroups`, `Microsoft.Databricks/workspaces`, `Microsoft.DBforPostgreSQL/serversv2`, `Microsoft.HardwareSecurityModules/dedicatedHSMs`, `Microsoft.Logic/integrationServiceEnvironments`, `Microsoft.Netapp/volumes`, `Microsoft.ServiceFabricMesh/networks`, `Microsoft.Sql/managedInstances`, `Microsoft.Sql/servers`, `Microsoft.StreamAnalytics/streamingJobs`, `Microsoft.Web/hostingEnvironments` and `Microsoft.Web/serverFarms`.

- `actions` - (Optional) A list of Actions which should be delegated. This list is specific to the service to delegate to. Possible values include `Microsoft.Network/networkinterfaces/*` , `Microsoft.Network/virtualNetworks/subnets/action` , `Microsoft.Network/virtualNetworks/subnets/join/action` , `Microsoft.Network/virtualNetworks/subnets/prepareNetworkPolicies/action` and `Microsoft.Network/virtualNetworks/subnets/unprepareNetworkPolicies/action` .

Attributes Reference

The following attributes are exported:

- `id` - The subnet ID.
- `ip_configurations` - The collection of IP Configurations with IPs within this subnet.
- `name` - The name of the subnet.
- `resource_group_name` - The name of the resource group in which the subnet is created in.
- `virtual_network_name` - The name of the virtual network in which the subnet is created in
- `address_prefix` - The address prefix for the subnet

Import

Subnets can be imported using the `resource id` , e.g.

```
terraform import azurerm_subnet.testSubnet /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/virtualNetworks/myvnet1/subnets/mysubnet1
```

azurerm_subnet_network_security_group_association

Associates a Network Security Group ([/docs/providers/azurerm/r/network_security_group.html](/docs/providers/azurerm/r/network_security_group.html)) with a Subnet (</docs/providers/azurerm/r/subnet.html>) within a Virtual Network ([/docs/providers/azurerm/r/virtual_network.html](/docs/providers/azurerm/r/virtual_network.html)).

NOTE: Subnet <-> Network Security Group associations currently need to be configured on both this resource and using the `network_security_group_id` field on the `azurerm_subnet` resource. The next major version of the AzureRM Provider (2.0) will remove the `network_security_group_id` field from the `azurerm_subnet` resource such that this resource is used to link resources in future.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                    = "frontend"
  resource_group_name     = "${azurerm_resource_group.test.name}"
  virtual_network_name    = "${azurerm_virtual_network.test.name}"
  address_prefix          = "10.0.2.0/24"
  network_security_group_id = "${azurerm_network_security_group.test.id}"
}

resource "azurerm_network_security_group" "test" {
  name      = "example-nsg"
  location  = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  security_rule {
    name                = "test123"
    priority             = 100
    direction           = "Inbound"
    access               = "Allow"
    protocol             = "Tcp"
    source_port_range    = "*"
    destination_port_range = "*"
    source_address_prefix = "*"
    destination_address_prefix = "*"
  }
}

resource "azurerm_subnet_network_security_group_association" "test" {
  subnet_id                = "${azurerm_subnet.test.id}"
  network_security_group_id = "${azurerm_network_security_group.test.id}"
}

```

Argument Reference

The following arguments are supported:

- `network_security_group_id` - (Required) The ID of the Network Security Group which should be associated with the Subnet. Changing this forces a new resource to be created.
- `subnet_id` - (Required) The ID of the Subnet. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Subnet.

Import

Subnet <-> Network Security Group Associations can be imported using the `resource_id` of the Subnet, e.g.

```
terraform import azurerm_subnet_network_security_group_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/virtualNetworks/myvnet1/subnets/mysubnet1
```

azurerm_subnet_route_table_association

Associates a Route Table ([/docs/providers/azurerm/r/route_table.html](/docs/providers/azurerm/r/route_table.html)) with a Subnet (</docs/providers/azurerm/r/subnet.html>) within a Virtual Network ([/docs/providers/azurerm/r/virtual_network.html](/docs/providers/azurerm/r/virtual_network.html)).

NOTE: Subnet <-> Route Table associations currently need to be configured on both this resource and using the `route_table_id` field on the `azurerm_subnet` resource. The next major version of the AzureRM Provider (2.0) will remove the `route_table_id` field from the `azurerm_subnet` resource such that this resource is used to link resources in future.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name            = "example-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "frontend"
  resource_group_name  = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix        = "10.0.2.0/24"
  route_table_id       = "${azurerm_route_table.test.id}"
}

resource "azurerm_route_table" "test" {
  name            = "example-routetable"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  route {
    name                 = "example"
    address_prefix        = "10.100.0.0/14"
    next_hop_type         = "VirtualAppliance"
    next_hop_in_ip_address = "10.10.1.1"
  }
}

resource "azurerm_subnet_route_table_association" "test" {
  subnet_id      = "${azurerm_subnet.test.id}"
  route_table_id = "${azurerm_route_table.test.id}"
}
```

Argument Reference

The following arguments are supported:

- `route_table_id` - (Required) The ID of the Route Table which should be associated with the Subnet. Changing this forces a new resource to be created.
- `subnet_id` - (Required) The ID of the Subnet. Changing this forces a new resource to be created.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Subnet.

Import

Subnet Route Table Associations can be imported using the `resource id` of the Subnet, e.g.

```
terraform import azurerm_subnet_route_table_association.association1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/virtualNetworks/myvnet1/subnets/mysubnet1
```

azurerm_template_deployment

Manages a template deployment of resources

Note on ARM Template Deployments: Due to the way the underlying Azure API is designed, Terraform can only manage the deployment of the ARM Template - and not any resources which are created by it. This means that when deleting the `azurerm_template_deployment` resource, Terraform will only remove the reference to the deployment, whilst leaving any resources created by that ARM Template Deployment. One workaround for this is to use a unique Resource Group for each ARM Template Deployment, which means deleting the Resource Group would contain any resources created within it - however this isn't ideal. More information (https://docs.microsoft.com/en-us/rest/api/resources/deployments#Deployments_Delete).

Example Usage

Note: This example uses Storage Accounts ([/docs/providers/azurerm/r/storage_account.html](/docs/providers/azurerm/r/storage_account.html)) and Public IP's ([/docs/providers/azurerm/r/public_ip.html](/docs/providers/azurerm/r/public_ip.html)) which are natively supported by Terraform - we'd highly recommend using the Native Resources where possible instead rather than an ARM Template, for the reasons outlined above.

```
resource "azurerm_resource_group" "test" {
  name      = "acctestRG-01"
  location = "West US"
}

resource "azurerm_template_deployment" "test" {
  name                = "acctesttemplate-01"
  resource_group_name = "${azurerm_resource_group.test.name}"

  template_body = <<DEPLOY
{
  "$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {
    "storageAccountType": {
      "type": "string",
      "defaultValue": "Standard_LRS",
      "allowedValues": [
        "Standard_LRS",
        "Standard_GRS",
        "Standard_ZRS"
      ],
      "metadata": {
        "description": "Storage Account type"
      }
    }
  },
  "variables": {
    "location": "[resourceGroup().location]",
    "storageAccountName": "[concat(uniquestring(resourceGroup().id), 'storage')]",
    "publicIPAddressName": "[concat('myPublicIp', uniquestring(resourceGroup().id))]",
    "publicIPAddressType": "Dynamic",
```

```

    "apiVersion": "2015-06-15",
    "dnsLabelPrefix": "terraform-acctest"
  },
  "resources": [
    {
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[variables('storageAccountName')]",
      "apiVersion": "[variables('apiVersion')]",
      "location": "[variables('location')]",
      "properties": {
        "accountType": "[parameters('storageAccountType')]"
      }
    },
    {
      "type": "Microsoft.Network/publicIPAddresses",
      "apiVersion": "[variables('apiVersion')]",
      "name": "[variables('publicIPAddressName')]",
      "location": "[variables('location')]",
      "properties": {
        "publicIPAllocationMethod": "[variables('publicIPAllocationType')]",
        "dnsSettings": {
          "domainNameLabel": "[variables('dnsLabelPrefix')]"
        }
      }
    }
  ],
  "outputs": {
    "storageAccountName": {
      "type": "string",
      "value": "[variables('storageAccountName')]"
    }
  }
}
DEPLOY

# these key-value pairs are passed into the ARM Template's `parameters` block
parameters = {
  "storageAccountType" = "Standard_GRS"
}

deployment_mode = "Incremental"
}

output "storageAccountName" {
  value = "${lookup(azurerm_template_deployment.test.outputs, "storageAccountName")}"
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the template deployment. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the template deployment.
- **deployment_mode** - (Required) Specifies the mode that is used to deploy resources. This value could be either

`Incremental` or `Complete`. Note that you will almost *always* want this to be set to `Incremental` otherwise the deployment will destroy all infrastructure not specified within the template, and Terraform will not be aware of this.

- `template_body` - (Optional) Specifies the JSON definition for the template.

Note: There's a `file` function available (<https://www.terraform.io/docs/configuration/functions/file.html>) which allows you to read this from an external file, which helps makes this more resource more readable.

- `parameters` - (Optional) Specifies the name and value pairs that define the deployment parameters for the template.
- `parameters_body` - (Optional) Specifies a valid Azure JSON parameters file that define the deployment parameters. It can contain KeyVault references

Note: There's a `file` function available (<https://www.terraform.io/docs/configuration/functions/file.html>) which allows you to read this from an external file, which helps makes this more resource more readable.

Attributes Reference

The following attributes are exported:

- `id` - The Template Deployment ID.
- `outputs` - A map of supported scalar output types returned from the deployment (currently, Azure Template Deployment outputs of type String, Int and Bool are supported, and are converted to strings - others will be ignored) and can be accessed using `.outputs["name"]`.

Note

Terraform does not know about the individual resources created by Azure using a deployment template and therefore cannot delete these resources during a destroy. Destroying a template deployment removes the associated deployment operations, but will not delete the Azure resources created by the deployment. In order to delete these resources, the containing resource group must also be destroyed. More information (https://docs.microsoft.com/en-us/rest/api/resources/deployments#Deployments_Delete).

azurerm_traffic_manager_endpoint

Manages a Traffic Manager Endpoint.

Example Usage

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name      = "trafficmanagerendpointTest"
  location  = "West US"
}

resource "azurerm_traffic_manager_profile" "test" {
  name                        = "${random_id.server.hex}"
  resource_group_name        = "${azurerm_resource_group.test.name}"

  traffic_routing_method = "Weighted"

  dns_config {
    relative_name = "${random_id.server.hex}"
    ttl           = 100
  }

  monitor_config {
    protocol           = "http"
    port               = 80
    path               = "/"
    interval_in_seconds = 30
    timeout_in_seconds  = 9
    tolerated_number_of_failures = 3
  }

  tags = {
    environment = "Production"
  }
}

resource "azurerm_traffic_manager_endpoint" "test" {
  name                        = "${random_id.server.hex}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  profile_name                = "${azurerm_traffic_manager_profile.test.name}"
  target                     = "terraform.io"
  type                       = "externalEndpoints"
  weight                     = 100
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Traffic Manager endpoint. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Traffic Manager endpoint.
- `profile_name` - (Required) The name of the Traffic Manager Profile to attach create the Traffic Manager endpoint.
- `endpoint_status` - (Optional) The status of the Endpoint, can be set to either `Enabled` or `Disabled`. Defaults to `Enabled`.
- `type` - (Required) The Endpoint type, must be one of:
 - `azureEndpoints`
 - `externalEndpoints`
 - `nestedEndpoints`
- `target` - (Optional) The FQDN DNS name of the target. This argument must be provided for an endpoint of type `externalEndpoints`, for other types it will be computed.
- `target_resource_id` - (Optional) The resource id of an Azure resource to target. This argument must be provided for an endpoint of type `azureEndpoints` or `nestedEndpoints`.
- `weight` - (Optional) Specifies how much traffic should be distributed to this endpoint, this must be specified for Profiles using the `Weighted` traffic routing method. Supports values between 1 and 1000.
- `priority` - (Optional) Specifies the priority of this Endpoint, this must be specified for Profiles using the `Priority` traffic routing method. Supports values between 1 and 1000, with no Endpoints sharing the same value. If omitted the value will be computed in order of creation.
- `endpoint_location` - (Optional) Specifies the Azure location of the Endpoint, this must be specified for Profiles using the `Performance` routing method if the Endpoint is of either type `nestedEndpoints` or `externalEndpoints`. For Endpoints of type `azureEndpoints` the value will be taken from the location of the Azure target resource.
- `min_child_endpoints` - (Optional) This argument specifies the minimum number of endpoints that must be 'online' in the child profile in order for the parent profile to direct traffic to any of the endpoints in that child profile. This argument only applies to Endpoints of type `nestedEndpoints` and defaults to `1`.
- `geo_mappings` - (Optional) A list of Geographic Regions used to distribute traffic, such as `WORLD`, `UK` or `DE`. The same location can't be specified in two endpoints. See the Geographic Hierarchies documentation for more information (<https://docs.microsoft.com/en-us/rest/api/trafficmanager/geographichierarchies/getdefault>).
- `custom_header` - (Optional) One or more `custom_header` blocks as defined below
- `subnet` - (Optional) One or more `subnet` blocks as defined below

A `custom_header` block supports the following:

- `name` - (Required) The name of the custom header.
- `value` - (Required) The value of custom header. Applicable for Http and Https protocol.

A `subnet` block supports the following:

- `first` - (Required) The First IP....
- `last` - (Optional) The Last IP...
- `scope` - (Optional) The Scope...

NOTE: One and only one of either `last` (in case of IP range) or `scope` (in case of CIDR) must be specified.

Attributes Reference

The following attributes are exported:

- `id` - The Traffic Manager Endpoint id.

Import

Traffic Manager Endpoints can be imported using the `resource id`, e.g.

```
terraform import azurerm_traffic_manager_endpoint.testEndpoints /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/trafficManagerProfiles/mytrafficmanagerprofile1/azureEndpoints/mytrafficmanagerendpoint
```

azurerm_traffic_manager_profile

Manages a Traffic Manager Profile to which multiple endpoints can be attached.

Example Usage

```
resource "random_id" "server" {
  keepers = {
    azi_id = 1
  }

  byte_length = 8
}

resource "azurerm_resource_group" "test" {
  name      = "trafficmanagerProfile"
  location  = "West US"
}

resource "azurerm_traffic_manager_profile" "test" {
  name                        = "${random_id.server.hex}"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  traffic_routing_method     = "Weighted"

  dns_config {
    relative_name = "${random_id.server.hex}"
    ttl           = 100
  }

  monitor_config {
    protocol           = "http"
    port               = 80
    path               = "/"
    interval_in_seconds = 30
    timeout_in_seconds  = 9
    tolerated_number_of_failures = 3
  }

  tags = {
    environment = "Production"
  }
}
```

Argument Reference

The following arguments are supported:

- **name** - (Required) The name of the virtual network. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) The name of the resource group in which to create the virtual network.

- `profile_status` - (Optional) The status of the profile, can be set to either `Enabled` or `Disabled`. Defaults to `Enabled`.
- `traffic_routing_method` - (Required) Specifies the algorithm used to route traffic, possible values are:
 - `Geographic` - Traffic is routed based on Geographic regions specified in the Endpoint.
 - `MultiValue` - All healthy Endpoints are returned. MultiValue routing method works only if all the endpoints of type 'External' and are specified as IPv4 or IPv6 addresses.
 - `Performance` - Traffic is routed via the User's closest Endpoint
 - `Priority` - Traffic is routed to the Endpoint with the lowest `priority` value.
 - `Subnet` - Traffic is routed based on a mapping of sets of end-user IP address ranges to a specific Endpoint within a Traffic Manager profile.
 - `Weighted` - Traffic is spread across Endpoints proportional to their `weight` value.
- `dns_config` - (Required) This block specifies the DNS configuration of the Profile, it supports the fields documented below.
- `monitor_config` - (Required) This block specifies the Endpoint monitoring configuration for the Profile, it supports the fields documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `dns_config` block supports:

- `relative_name` - (Required) The relative domain name, this is combined with the domain name used by Traffic Manager to form the FQDN which is exported as documented below. Changing this forces a new resource to be created.
- `ttl` - (Required) The TTL value of the Profile used by Local DNS resolvers and clients.

The `monitor_config` block supports:

- `protocol` - (Required) The protocol used by the monitoring checks, supported values are `HTTP`, `HTTPS` and `TCP`.
- `port` - (Required) The port number used by the monitoring checks.
- `path` - (Optional) The path used by the monitoring checks. Required when `protocol` is set to `HTTP` or `HTTPS` - cannot be set when `protocol` is set to `TCP`.
- `interval_in_seconds` - (Optional) The interval used to check the endpoint health from a Traffic Manager probing agent. You can specify two values here: `30` (normal probing) and `10` (fast probing). The default value is `30`.
- `timeout_in_seconds` - (Optional) The amount of time the Traffic Manager probing agent should wait before considering that check a failure when a health check probe is sent to the endpoint. If `interval_in_seconds` is set to `30`, then `timeout_in_seconds` can be between `5` and `10`. The default value is `10`. If `interval_in_seconds` is set to `10`, then valid values are between `5` and `9` and `timeout_in_seconds` is required.
- `tolerated_number_of_failures` - (Optional) The number of failures a Traffic Manager probing agent tolerates before marking that endpoint as unhealthy. Valid values are between `0` and `9`. The default value is `3`.

Attributes Reference

The following attributes are exported:

- `id` - The Traffic Manager Profile id.
- `fqdn` - The FQDN of the created Profile.

Notes

The Traffic Manager is created with the location `global`.

Import

Traffic Manager Profiles can be imported using the `resource id`, e.g.

```
terraform import azurerm_traffic_manager_profile.testProfile /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/trafficManagerProfiles/mytrafficmanagerprofile1
```

azurerm_user_assigned_identity

Manages a user assigned identity.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "eastus"
}

resource "azurerm_user_assigned_identity" "testIdentity" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  location             = "${azurerm_resource_group.test.location}"

  name = "search-api"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the user assigned identity. Changing this forces a new identity to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the user assigned identity.
- `location` - (Required) The location/region where the user assigned identity is created.
- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The user assigned identity ID.
- `principal_id` - Service Principal ID associated with the user assigned identity.
- `client_id` - Client ID associated with the user assigned identity.

Import

User Assigned Identities can be imported using the `resource id`, e.g.

```
terraform import azurerm_user_assigned_identity.testIdentity /subscriptions/00000000-0000-0000-0000-000000000000/resourcegroups/acceptanceTestResourceGroup1/providers/Microsoft.ManagedIdentity/userAssignedIdentities/testIdentity
```

azurerm_virtual_machine_data_disk_attachment

Manages attaching a Disk to a Virtual Machine.

NOTE: Data Disks can be attached either directly on the `azurerm_virtual_machine` resource, or using the `azurerm_virtual_machine_data_disk_attachment` resource - but the two cannot be used together. If both are used against the same Virtual Machine, spurious changes will occur.

Please Note: only Managed Disks are supported via this separate resource, Unmanaged Disks can be attached using the `storage_data_disk` block in the `azurerm_virtual_machine` resource.

Example Usage

```
variable "prefix" {
  default = "example"
}

locals {
  vm_name = "${var.prefix}-vm"
}

resource "azurerm_resource_group" "main" {
  name     = "${var.prefix}-resources"
  location = "West Europe"
}

resource "azurerm_virtual_network" "main" {
  name            = "${var.prefix}-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"
}

resource "azurerm_subnet" "internal" {
  name                 = "internal"
  resource_group_name = "${azurerm_resource_group.main.name}"
  virtual_network_name = "${azurerm_virtual_network.main.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "main" {
  name                = "${var.prefix}-nic"
  location            = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"

  ip_configuration {
    name                          = "internal"
    subnet_id                    = "${azurerm_subnet.internal.id}"
    private_ip_address_allocation = "Dynamic"
  }
}
```

```

resource "azurerm_virtual_machine" "test" {
  name                = "${local.vm_name}"
  location            = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"
  network_interface_ids = ["${azurerm_network_interface.main.id}"]
  vm_size             = "Standard_F2"

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }

  storage_os_disk {
    name                = "myosdisk1"
    caching             = "ReadWrite"
    create_option       = "FromImage"
    managed_disk_type   = "Standard_LRS"
  }

  os_profile {
    computer_name  = "${local.vm_name}"
    admin_username = "testadmin"
    admin_password = "Password1234!"
  }

  os_profile_linux_config {
    disable_password_authentication = false
  }
}

resource "azurerm_managed_disk" "test" {
  name                = "${local.vm_name}-disk1"
  location            = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"
  storage_account_type = "Standard_LRS"
  create_option       = "Empty"
  disk_size_gb        = 10
}

resource "azurerm_virtual_machine_data_disk_attachment" "test" {
  managed_disk_id    = "${azurerm_managed_disk.test.id}"
  virtual_machine_id = "${azurerm_virtual_machine.test.id}"
  lun                = "10"
  caching             = "ReadWrite"
}

```

Argument Reference

The following arguments are supported:

- `virtual_machine_id` - (Required) The ID of the Virtual Machine to which the Data Disk should be attached. Changing this forces a new resource to be created.

- `managed_disk_id` - (Required) The ID of an existing Managed Disk which should be attached. Changing this forces a new resource to be created.
- `lun` - (Required) The Logical Unit Number of the Data Disk, which needs to be unique within the Virtual Machine. Changing this forces a new resource to be created.
- `caching` - (Required) Specifies the caching requirements for this Data Disk. Possible values include `None` , `ReadOnly` and `ReadWrite` .
- `create_option` - (Optional) The Create Option of the Data Disk, such as `Empty` or `Attach` . Defaults to `Attach` . Changing this forces a new resource to be created.
- `write_accelerator_enabled` - (Optional) Specifies if Write Accelerator is enabled on the disk. This can only be enabled on `Premium_LRS` managed disks with no caching and M-Series VMs (<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/how-to-enable-write-accelerator>). Defaults to `false` .

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual Machine Data Disk attachment.

Import

Virtual Machines Data Disk Attachments can be imported using the `resource id` , e.g.

```
terraform import azurerm_virtual_machine_data_disk_attachment.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.compute/virtualMachines/machine1/dataDisks/disk1
```

Please Note: This is a Terraform Unique ID matching the format: `{virtualMachineID}/dataDisks/{diskName}`

azurerm_virtual_machine_extension

Manages a Virtual Machine Extension to provide post deployment configuration and run automated tasks.

NOTE: Custom Script Extensions for Linux & Windows require that the `commandToExecute` returns a `0` exit code to be classified as successfully deployed. You can achieve this by appending `exit 0` to the end of your `commandToExecute`.

NOTE: Custom Script Extensions require that the Azure Virtual Machine Guest Agent is running on the Virtual Machine.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "acctestRG"
  location  = "West US"
}

resource "azurerm_virtual_network" "test" {
  name            = "acctvn"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "acctsub"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "test" {
  name                = "acctni"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  ip_configuration {
    name                       = "testconfiguration1"
    subnet_id                 = "${azurerm_subnet.test.id}"
    private_ip_address_allocation = "Dynamic"
  }
}

resource "azurerm_storage_account" "test" {
  name                = "accsa"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"
  account_tier        = "Standard"
  account_replication_type = "LRS"

  tags = {
    environment = "staging"
  }
}
```

```

}

resource "azurerm_storage_container" "test" {
  name                = "vhds"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
  container_access_type = "private"
}

resource "azurerm_virtual_machine" "test" {
  name                = "acctvm"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  network_interface_ids = [ "${azurerm_network_interface.test.id}" ]
  vm_size             = "Standard_F2"

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }

  storage_os_disk {
    name          = "myosdisk1"
    vhd_uri       = "${azurerm_storage_account.test.primary_blob_endpoint}${azurerm_storage_container.test.name}/myosdisk1.vhd"
    caching       = "ReadWrite"
    create_option = "FromImage"
  }

  os_profile {
    computer_name  = "hostname"
    admin_username = "testadmin"
    admin_password = "Password1234!"
  }

  os_profile_linux_config {
    disable_password_authentication = false
  }

  tags = {
    environment = "staging"
  }
}

resource "azurerm_virtual_machine_extension" "test" {
  name                = "hostname"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_machine_name = "${azurerm_virtual_machine.test.name}"
  publisher           = "Microsoft.Azure.Extensions"
  type                = "CustomScript"
  type_handler_version = "2.0"

  settings = <<SETTINGS
  {
    "commandToExecute": "hostname && uptime"
  }
}

```

SETTINGS

```
tags = {  
  environment = "Production"  
}  
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the virtual machine extension peering. Changing this forces a new resource to be created.
- `location` - (Required) The location where the extension is created. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the virtual network. Changing this forces a new resource to be created.
- `virtual_machine_name` - (Required) The name of the virtual machine. Changing this forces a new resource to be created.
- `publisher` - (Required) The publisher of the extension, available publishers can be found by using the Azure CLI.
- `type` - (Required) The type of extension, available types for a publisher can be found using the Azure CLI.

Note: The Publisher and Type of Virtual Machine Extensions can be found using the Azure CLI, via: `shell $ az vm extension image list --location westus -o table`

- `type_handler_version` - (Required) Specifies the version of the extension to use, available versions can be found using the Azure CLI.
- `auto_upgrade_minor_version` - (Optional) Specifies if the platform deploys the latest minor version update to the `type_handler_version` specified.
- `settings` - (Required) The settings passed to the extension, these are specified as a JSON object in a string.

Please Note: Certain VM Extensions require that the keys in the `settings` block are case sensitive. If you're seeing unhelpful errors, please ensure the keys are consistent with how Azure is expecting them (for instance, for the `JsonAddDomainExtension` extension, the keys are expected to be in `TitleCase`.)

- `protected_settings` - (Optional) The `protected_settings` passed to the extension, like `settings`, these are specified as a JSON object in a string.

Please Note: Certain VM Extensions require that the keys in the `protected_settings` block are case sensitive. If you're seeing unhelpful errors, please ensure the keys are consistent with how Azure is expecting them (for instance, for the `JsonAddDomainExtension` extension, the keys are expected to be in `TitleCase`.)

- `tags` - (Optional) A mapping of tags to assign to the resource.

Attributes Reference

The following attributes are exported:

- `id` - The Virtual Machine Extension ID.

Import

Virtual Machine Extensions can be imported using the `resource id`, e.g.

```
terraform import azurerm_virtual_machine_extension.test /subscriptions/00000000-0000-0000-0000-000000000000/00/resourceGroups/mygroup1/providers/Microsoft.Compute/virtualMachines/myVM/extensions/hostname
```

azurerm_virtual_machine

Manages a Virtual Machine.

NOTE: Data Disks can be attached either directly on the `azurerm_virtual_machine` resource, or using the `azurerm_virtual_machine_data_disk_attachment` resource - but the two cannot be used together. If both are used against the same Virtual Machine, spurious changes will occur.

NOTE: The `azurerm_virtual_machine` resource will be superseded by two new resources in the next major version of the Azure Provider (2.0) - you can find out more about these changes here (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/2807>).

Example Usage (from an Azure Platform Image)

This example provisions a Virtual Machine with Managed Disks. Other examples of the `azurerm_virtual_machine` resource can be found in the `./examples/virtual-machines` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/virtual-machines>)

```
variable "prefix" {
  default = "tfvmex"
}

resource "azurerm_resource_group" "main" {
  name       = "${var.prefix}-resources"
  location   = "West US 2"
}

resource "azurerm_virtual_network" "main" {
  name            = "${var.prefix}-network"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"
}

resource "azurerm_subnet" "internal" {
  name                 = "internal"
  resource_group_name = "${azurerm_resource_group.main.name}"
  virtual_network_name = "${azurerm_virtual_network.main.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_network_interface" "main" {
  name                = "${var.prefix}-nic"
  location            = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"

  ip_configuration {
    name                          = "testconfiguration1"
    subnet_id                    = "${azurerm_subnet.internal.id}"
    private_ip_address_allocation = "Dynamic"
  }
}
```

```

}

resource "azurerm_virtual_machine" "main" {
  name                = "${var.prefix}-vm"
  location            = "${azurerm_resource_group.main.location}"
  resource_group_name = "${azurerm_resource_group.main.name}"
  network_interface_ids = ["${azurerm_network_interface.main.id}"]
  vm_size             = "Standard_DS1_v2"

  # Uncomment this line to delete the OS disk automatically when deleting the VM
  # delete_os_disk_on_termination = true

  # Uncomment this line to delete the data disks automatically when deleting the VM
  # delete_data_disks_on_termination = true

  storage_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }
  storage_os_disk {
    name          = "myosdisk1"
    caching       = "ReadWrite"
    create_option = "FromImage"
    managed_disk_type = "Standard_LRS"
  }
  os_profile {
    computer_name  = "hostname"
    admin_username = "testadmin"
    admin_password = "Password1234!"
  }
  os_profile_linux_config {
    disable_password_authentication = false
  }
  tags = {
    environment = "staging"
  }
}

```

Argument Reference

The following arguments are supported:

- **name** - (Required) Specifies the name of the Virtual Machine. Changing this forces a new resource to be created.
- **resource_group_name** - (Required) Specifies the name of the Resource Group in which the Virtual Machine should exist. Changing this forces a new resource to be created.
- **location** - (Required) Specifies the Azure Region where the Virtual Machine exists. Changing this forces a new resource to be created.
- **network_interface_ids** - (Required) A list of Network Interface ID's which should be associated with the Virtual Machine.

- `os_profile_linux_config` - (Required, when a Linux machine) A `os_profile_linux_config` block.
- `os_profile_windows_config` - (Required, when a Windows machine) A `os_profile_windows_config` block.
- `vm_size` - (Required) Specifies the size of the Virtual Machine (<https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-size-specs/>).

- `availability_set_id` - (Optional) The ID of the Availability Set in which the Virtual Machine should exist. Changing this forces a new resource to be created.
- `boot_diagnostics` - (Optional) A `boot_diagnostics` block.
- `additional_capabilities` - (Optional) A `additional_capabilities` block.
- `delete_os_disk_on_termination` - (Optional) Should the OS Disk (either the Managed Disk / VHD Blob) be deleted when the Virtual Machine is destroyed? Defaults to `false`.

Note: This setting works when instance is deleted via Terraform only and don't forget to delete disks manually if you deleted VM manually. It can increase spending.

- `delete_data_disks_on_termination` - (Optional) Should the Data Disks (either the Managed Disks / VHD Blobs) be deleted when the Virtual Machine is destroyed? Defaults to `false`.

Note: This setting works when instance is deleted via Terraform only and don't forget to delete disks manually if you deleted VM manually. It can increase spending.

- `identity` - (Optional) A `identity` block.
- `license_type` - (Optional) Specifies the BYOL Type for this Virtual Machine. This is only applicable to Windows Virtual Machines. Possible values are `Windows_Client` and `Windows_Server`.
- `os_profile` - (Optional) An `os_profile` block. Required when `create_option` in the `storage_os_disk` block is set to `FromImage`.
- `os_profile_secrets` - (Optional) One or more `os_profile_secrets` blocks.
- `plan` - (Optional) A `plan` block.
- `primary_network_interface_id` - (Optional) The ID of the Network Interface (which must be attached to the Virtual Machine) which should be the Primary Network Interface for this Virtual Machine.
- `proximity_placement_group_id` - (Optional) The ID of the Proximity Placement Group to which this Virtual Machine should be assigned. Changing this forces a new resource to be created
- `storage_data_disk` - (Optional) One or more `storage_data_disk` blocks.

Please Note: Data Disks can also be attached either using this block or the `azurerm_virtual_machine_data_disk_attachment` resource ([/docs/providers/azurerm/r/virtual_machine_data_disk_attachment.html](https://docs.providers.azurearm.com/r/virtual_machine_data_disk_attachment.html)) - but not both.

- `storage_image_reference` - (Optional) A `storage_image_reference` block.

- `storage_os_disk` - (Required) A `storage_os_disk` block.
- `tags` - (Optional) A mapping of tags to assign to the Virtual Machine.
- `zones` - (Optional) A list of a single item of the Availability Zone which the Virtual Machine should be allocated in.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

For more information on the different example configurations, please check out the [Azure documentation](https://docs.microsoft.com/en-gb/rest/api/compute/virtualmachines/createorupdate#examples) (<https://docs.microsoft.com/en-gb/rest/api/compute/virtualmachines/createorupdate#examples>)

A `additional_unattend_config` block supports the following:

- `pass` - (Required) Specifies the name of the pass that the content applies to. The only allowable value is `oobeSystem`.
- `component` - (Required) Specifies the name of the component to configure with the added content. The only allowable value is `Microsoft-Windows-Shell-Setup`.
- `setting_name` - (Required) Specifies the name of the setting to which the content applies. Possible values are: `FirstLogonCommands` and `AutoLogon`.
- `content` - (Optional) Specifies the base-64 encoded XML formatted content that is added to the `unattend.xml` file for the specified path and component.

A `boot_diagnostics` block supports the following:

- `enabled` - (Required) Should Boot Diagnostics be enabled for this Virtual Machine?
- `storage_uri` - (Required) The Storage Account's Blob Endpoint which should hold the virtual machine's diagnostic files.

NOTE: This needs to be the root of a Storage Account and not a Storage Container.

A `additional_capabilities` block supports the following:

- `ultra_ssd_enabled` - (Required) Should Ultra SSD disk be enabled for this Virtual Machine?

Note: Azure Ultra Disk Storage is currently in preview and are not available to subscriptions that have not requested (<https://aka.ms/UltraSSDPreviewSignUp>) onboarding to Azure Ultra Disk Storage preview. Azure Ultra Disk Storage is only available in East US 2, North Europe, and Southeast Asia regions. For more information see the Azure Ultra Disk Storage product documentation (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-enable-ultra-ssd>), product blog (<https://azure.microsoft.com/en-us/blog/announcing-the-general-availability-of-azure-ultra-disk-storage/>) and FAQ (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/faq-for-disks#ultra-disks>).

A `identity` block supports the following:

- `type` - (Required) The Managed Service Identity Type of this Virtual Machine. Possible values are `SystemAssigned`

(where Azure will generate a Service Principal for you), `UserAssigned` (where you can specify the Service Principal ID's) to be used by this Virtual Machine using the `identity_ids` field, and `SystemAssigned`, `UserAssigned` which assigns both a system managed identity as well as the specified user assigned identities.

NOTE: Managed Service Identity previously required the installation of a VM Extension, but this information is now available via the Azure Instance Metadata Service (<https://docs.microsoft.com/en-us/azure/active-directory/managed-service-identity/overview#how-does-it-work>).

NOTE: When `type` is set to `SystemAssigned`, identity the Principal ID can be retrieved after the virtual machine has been created. More details are available below. See documentation (<https://docs.microsoft.com/en-us/azure/active-directory/managed-service-identity/overview>) for additional information.

- `identity_ids` - (Optional) Specifies a list of user managed identity ids to be assigned to the VM. Required if `type` is `UserAssigned`.

A `os_profile` block supports the following:

- `computer_name` - (Required) Specifies the name of the Virtual Machine.
- `admin_username` - (Required) Specifies the name of the local administrator account.
- `admin_password` - (Required for Windows, Optional for Linux) The password associated with the local administrator account.

NOTE: If using Linux, it may be preferable to use SSH Key authentication (available in the `os_profile_linux_config` block) instead of password authentication.

NOTE: `admin_password` must be between 6-72 characters long and must satisfy at least 3 of password complexity requirements from the following: 1. Contains an uppercase character 2. Contains a lowercase character 3. Contains a numeric digit 4. Contains a special character

- `custom_data` - (Optional) Specifies custom data to supply to the machine. On Linux-based systems, this can be used as a cloud-init script. On other systems, this will be copied as a file on disk. Internally, Terraform will base64 encode this value before sending it to the API. The maximum length of the binary array is 65535 bytes.

A `os_profile_linux_config` block supports the following:

- `disable_password_authentication` - (Required) Specifies whether password authentication should be disabled. If set to `false`, an `admin_password` must be specified.
- `ssh_keys` - (Optional) One or more `ssh_keys` blocks. This field is required if `disable_password_authentication` is set to `true`.

A `os_profile_secrets` block supports the following:

- `source_vault_id` - (Required) Specifies the ID of the Key Vault to use.
- `vault_certificates` - (Required) One or more `vault_certificates` blocks.

A `os_profile_windows_config` block supports the following:

- `provision_vm_agent` - (Optional) Should the Azure Virtual Machine Guest Agent be installed on this Virtual Machine? Defaults to `false`.

NOTE: This is different from the Default value used for this field within Azure.

- `enable_automatic_upgrades` - (Optional) Are automatic updates enabled on this Virtual Machine? Defaults to `false`.
- `timezone` - (Optional) Specifies the time zone of the virtual machine, the possible values are defined here (<http://jackstromberg.com/2017/01/list-of-time-zones-consumed-by-azure/>).
- `winrm` - (Optional) One or more `winrm` block.
- `additional_unattend_config` - (Optional) A `additional_unattend_config` block.

A `plan` block supports the following:

- `name` - (Required) Specifies the name of the image from the marketplace.
- `publisher` - (Required) Specifies the publisher of the image.
- `product` - (Required) Specifies the product of the image from the marketplace.

A `ssh_keys` block supports the following:

- `key_data` - (Required) The Public SSH Key which should be written to the `path` defined above.

NOTE: Rather than defining this in-line you can source this from a local file using the `file` function (<https://www.terraform.io/docs/configuration/functions/file.html>) - for example `key_data = file("~/ssh/id_rsa.pub")`.

- `path` - (Required) The path of the destination file on the virtual machine

NOTE: Due to a limitation in the Azure VM Agent the only allowed `path` is `/home/{username}/.ssh/authorized_keys`.

A `storage_image_reference` block supports the following:

This block provisions the Virtual Machine from one of two sources: an Azure Platform Image (e.g. Ubuntu/Windows Server) or a Custom Image.

To provision from an Azure Platform Image, the following fields are applicable:

- `publisher` - (Required) Specifies the publisher of the image used to create the virtual machine. Changing this forces a new resource to be created.

- `offer` - (Required) Specifies the offer of the image used to create the virtual machine. Changing this forces a new resource to be created.
- `sku` - (Required) Specifies the SKU of the image used to create the virtual machine. Changing this forces a new resource to be created.
- `version` - (Optional) Specifies the version of the image used to create the virtual machine. Changing this forces a new resource to be created.

To provision a Custom Image, the following fields are applicable:

- `id` - (Required) Specifies the ID of the Custom Image which the Virtual Machine should be created from. Changing this forces a new resource to be created.

NOTE: An example of how to use this is available within the `./examples/virtual-machines/managed-disks/from-custom-image` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/virtual-machines/managed-disks/from-custom-image>)

A `storage_data_disk` block supports the following:

NOTE: Data Disks can also be attached either using this block or the `azurerm_virtual_machine_data_disk_attachment` resource ([/docs/providers/azurerm/r/virtual_machine_data_disk_attachment.html](/docs/providers/azurerm/r/virtual_machine_data_disk_attachment.html)) - but not both.

- `name` - (Required) The name of the Data Disk.
- `caching` - (Optional) Specifies the caching requirements for the Data Disk. Possible values include `None`, `ReadOnly` and `ReadWrite`.
- `create_option` - (Required) Specifies how the data disk should be created. Possible values are `Attach`, `FromImage` and `Empty`.

NOTE: If using an image that does not have data to be written to the Data Disk, use `Empty` as the create option in order to create the desired disk without any data.

- `disk_size_gb` - (Optional) Specifies the size of the data disk in gigabytes.
- `lun` - (Required) Specifies the logical unit number of the data disk. This needs to be unique within all the Data Disks on the Virtual Machine.
- `write_accelerator_enabled` - (Optional) Specifies if Write Accelerator is enabled on the disk. This can only be enabled on `Premium_LRS` managed disks with no caching and M-Series VMs (<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/how-to-enable-write-accelerator>). Defaults to `false`.

The following properties apply when using Managed Disks:

- `managed_disk_type` - (Optional) Specifies the type of managed disk to create. Possible values are either `Standard_LRS`, `StandardSSD_LRS`, `Premium_LRS` or `UltraSSD_LRS`.

Note: `managed_disk_type` of type `UltraSSD_LRS` is currently in preview and are not available to subscriptions that have not requested (<https://aka.ms/UltraSSDPreviewSignUp>) onboarding to Azure Ultra Disk Storage preview. Azure Ultra Disk Storage is only available in East US 2, North Europe, and Southeast Asia regions. For more information see the Azure Ultra Disk Storage product documentation (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-enable-ultra-ssd>), product blog (<https://azure.microsoft.com/en-us/blog/announcing-the-general-availability-of-azure-ultra-disk-storage/>) and FAQ (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/faq-for-disks#ultra-disks>). You must also set `additional_capabilities.ultra_ssd_enabled` to `true`.

- `managed_disk_id` - (Optional) Specifies the ID of an Existing Managed Disk which should be attached to this Virtual Machine. When this field is set `create_option` must be set to `Attach`.

The following properties apply when using Unmanaged Disks:

- `vhd_uri` - (Optional) Specifies the URI of the VHD file backing this Unmanaged Data Disk. Changing this forces a new resource to be created.

A `storage_os_disk` block supports the following:

- `name` - (Required) Specifies the name of the OS Disk.
- `create_option` - (Required) Specifies how the OS Disk should be created. Possible values are `Attach` (managed disks only) and `FromImage`.
- `caching` - (Optional) Specifies the caching requirements for the OS Disk. Possible values include `None`, `ReadOnly` and `ReadWrite`.
- `disk_size_gb` - (Optional) Specifies the size of the OS Disk in gigabytes.
- `image_uri` - (Optional) Specifies the Image URI in the format `publisherName:offer:skus:version`. This field can also specify the VHD uri (<https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-cli-deploy-templates/#create-a-custom-vm-image>) of a custom VM image to clone. When cloning a Custom (Unmanaged) Disk Image the `os_type` field must be set.
- `os_type` - (Optional) Specifies the Operating System on the OS Disk. Possible values are `Linux` and `Windows`.
- `write_accelerator_enabled` - (Optional) Specifies if Write Accelerator is enabled on the disk. This can only be enabled on `Premium_LRS` managed disks with no caching and M-Series VMs (<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/how-to-enable-write-accelerator>). Defaults to `false`.

The following properties apply when using Managed Disks:

- `managed_disk_id` - (Optional) Specifies the ID of an existing Managed Disk which should be attached as the OS Disk of this Virtual Machine. If this is set then the `create_option` must be set to `Attach`.
- `managed_disk_type` - (Optional) Specifies the type of Managed Disk which should be created. Possible values are `Standard_LRS`, `StandardSSD_LRS` or `Premium_LRS`.

The following properties apply when using Unmanaged Disks:

- `vhd_uri` - (Optional) Specifies the URI of the VHD file backing this Unmanaged OS Disk. Changing this forces a new resource to be created.
-

A `vault_certificates` block supports the following:

- `certificate_url` - (Required) The ID of the Key Vault Secret. Stored secret is the Base64 encoding of a JSON Object that which is encoded in UTF-8 of which the contents need to be:

```
{
  "data": "<Base64-encoded-certificate>",
  "dataType": "pfx",
  "password": "<pfx-file-password>"
}
```

NOTE: If your certificate is stored in Azure Key Vault - this can be sourced from the `secret_id` property on the `azurerm_key_vault_certificate` resource.

- `certificate_store` - (Required, on windows machines) Specifies the certificate store on the Virtual Machine where the certificate should be added to, such as `My` .

A `winrm` block supports the following:

- `protocol` - (Required) Specifies the protocol of listener. Possible values are `HTTP` or `HTTPS` .
- `certificate_url` - (Optional) The ID of the Key Vault Secret which contains the encrypted Certificate which should be installed on the Virtual Machine. This certificate must also be specified in the `vault_certificates` block within the `os_profile_secrets` block.

NOTE: This can be sourced from the `secret_id` field on the `azurerm_key_vault_certificate` resource.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual Machine.
- `identity` - An `identity` block as defined below, which contains the Managed Service Identity information for this Virtual Machine.

A `identity` block exports the following:

- `principal_id` - The Principal ID for the Service Principal associated with the Managed Service Identity of this Virtual Machine.

You can access the Principal ID via `${azurerm_virtual_machine.test.identity.0.principal_id}`

Import

Virtual Machines can be imported using the `resource_id`, e.g.

```
terraform import azurerm_virtual_machine.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/microsoft.compute/virtualMachines/machine1
```

azurerm_virtual_machine_scale_set

Manages a virtual machine scale set.

NOTE: All arguments including the administrator login and password will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

NOTE: The `azurerm_virtual_machine_scale_set` resource will be superseded by two new resources in the next major version of the Azure Provider (2.0) - you can find out more about these changes here (<https://github.com/terraform-providers/terraform-provider-azurerm/issues/2807>).

Example Usage with Managed Disks (Recommended)

```
resource "azurerm_resource_group" "test" {
  name     = "acctestRG"
  location = "West US 2"
}

resource "azurerm_virtual_network" "test" {
  name            = "acctvn"
  address_space   = ["10.0.0.0/16"]
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                 = "acctsub"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_public_ip" "test" {
  name                = "test"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method   = "Static"
  domain_name_label   = "${azurerm_resource_group.test.name}"

  tags = {
    environment = "staging"
  }
}

resource "azurerm_lb" "test" {
  name                = "test"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  frontend_ip_configuration {
    name          = "PublicIPAddress"
    public_ip_address = "${azurerm_public_ip.test.id}"
  }
}
```



```

    public_ip_address_id = "${azurerm_public_ip.test.id}"
  }
}

resource "azurerm_lb_backend_address_pool" "bpepool" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id      = "${azurerm_lb.test.id}"
  name                 = "BackEndAddressPool"
}

resource "azurerm_lb_nat_pool" "lbnatpool" {
  resource_group_name      = "${azurerm_resource_group.test.name}"
  name                     = "ssh"
  loadbalancer_id          = "${azurerm_lb.test.id}"
  protocol                 = "Tcp"
  frontend_port_start      = 50000
  frontend_port_end        = 50119
  backend_port             = 22
  frontend_ip_configuration_name = "PublicIPAddress"
}

resource "azurerm_lb_probe" "test" {
  resource_group_name = "${azurerm_resource_group.test.name}"
  loadbalancer_id      = "${azurerm_lb.test.id}"
  name                 = "http-probe"
  protocol             = "Http"
  request_path         = "/health"
  port                 = 8080
}

resource "azurerm_virtual_machine_scale_set" "test" {
  name                = "mytestscaleaset-1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  # automatic rolling upgrade
  automatic_os_upgrade = true
  upgrade_policy_mode  = "Rolling"

  rolling_upgrade_policy {
    max_batch_instance_percent      = 20
    max_unhealthy_instance_percent = 20
    max_unhealthy_upgraded_instance_percent = 5
    pause_time_between_batches     = "PT0S"
  }

  # required when using rolling upgrade policy
  health_probe_id = "${azurerm_lb_probe.test.id}"

  sku {
    name      = "Standard_F2"
    tier       = "Standard"
    capacity  = 2
  }

  storage_profile_image_reference {
    publisher = "Canonical"
    offer     = "UbuntuServer"
    sku       = "16.04-LTS"
    version   = "latest"
  }
}

```

```

}

storage_profile_os_disk {
  name          = ""
  caching        = "ReadWrite"
  create_option  = "FromImage"
  managed_disk_type = "Standard_LRS"
}

storage_profile_data_disk {
  lun          = 0
  caching       = "ReadWrite"
  create_option = "Empty"
  disk_size_gb  = 10
}

os_profile {
  computer_name_prefix = "testvm"
  admin_username       = "myadmin"
}

os_profile_linux_config {
  disable_password_authentication = true

  ssh_keys {
    path      = "/home/myadmin/.ssh/authorized_keys"
    key_data = "${file("~/ssh/demo_key.pub")}"
  }
}

network_profile {
  name      = "terraformnetworkprofile"
  primary   = true

  ip_configuration {
    name                        = "TestIPConfiguration"
    primary                    = true
    subnet_id                  = "${azurerm_subnet.test.id}"
    load_balancer_backend_address_pool_ids = ["${azurerm_lb_backend_address_pool.bpepool.id}"]
    load_balancer_inbound_nat_rules_ids   = ["${azurerm_lb_nat_pool.lbnatpool.id}"]
  }
}

tags = {
  environment = "staging"
}
}

```

Example Usage with Unmanaged Disks

```

resource "azurerm_resource_group" "test" {
  name     = "acctestRG"
  location = "West US"
}

```

```

resource "azurerm_virtual_network" "test" {
  name                = "acctvn"
  address_space       = ["10.0.0.0/16"]
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_subnet" "test" {
  name                = "acctsub"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.2.0/24"
}

resource "azurerm_storage_account" "test" {
  name                = "accsa"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "westus"
  account_tier        = "Standard"
  account_replication_type = "LRS"

  tags = {
    environment = "staging"
  }
}

resource "azurerm_storage_container" "test" {
  name                = "vhds"
  resource_group_name = "${azurerm_resource_group.test.name}"
  storage_account_name = "${azurerm_storage_account.test.name}"
  container_access_type = "private"
}

resource "azurerm_virtual_machine_scale_set" "test" {
  name                = "mytestscaleaset-1"
  location            = "West US"
  resource_group_name = "${azurerm_resource_group.test.name}"
  upgrade_policy_mode = "Manual"

  sku {
    name      = "Standard_F2"
    tier       = "Standard"
    capacity = 2
  }

  os_profile {
    computer_name_prefix = "testvm"
    admin_username       = "myadmin"
  }

  os_profile_linux_config {
    disable_password_authentication = true

    ssh_keys {
      path      = "/home/myadmin/.ssh/authorized_keys"
      key_data = "${file("${~}/.ssh/demo_key.pub")}"
    }
  }

  network_profile {

```

```

name      = "TestNetworkProfile"
primary   = true

ip_configuration {
  name      = "TestIPConfiguration"
  primary   = true
  subnet_id = "${azurerm_subnet.test.id}"
}

storage_profile_os_disk {
  name          = "osDiskProfile"
  caching        = "ReadWrite"
  create_option = "FromImage"
  vhd_containers = ["${azurerm_storage_account.test.primary_blob_endpoint}${azurerm_storage_container.test.name}"]
}

storage_profile_image_reference {
  publisher = "Canonical"
  offer     = "UbuntuServer"
  sku       = "16.04-LTS"
  version   = "latest"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the virtual machine scale set resource. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the virtual machine scale set. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `network_profile` - (Required) A collection of network profile block as documented below.
- `os_profile` - (Required) A Virtual Machine OS Profile block as documented below.
- `os_profile_windows_config` - (Required, when a windows machine) A Windows config block as documented below.
- `os_profile_linux_config` - (Required, when a linux machine) A Linux config block as documented below.
- `proximity_placement_group_id` - (Optional) The ID of the Proximity Placement Group to which this Virtual Machine should be assigned. Changing this forces a new resource to be created
- `sku` - (Required) A sku block as documented below.
- `storage_profile_os_disk` - (Required) A storage profile os disk block as documented below

- `upgrade_policy_mode` - (Required) Specifies the mode of an upgrade to virtual machines in the scale set. Possible values, `Rolling`, `Manual`, or `Automatic`. When choosing `Rolling`, you will need to set a health probe.
- `automatic_os_upgrade` - (Optional) Automatic OS patches can be applied by Azure to your scaleset. This is particularly useful when `upgrade_policy_mode` is set to `Rolling`. Defaults to `false`.
- `boot_diagnostics` - (Optional) A boot diagnostics profile block as referenced below.
- `extension` - (Optional) Can be specified multiple times to add extension profiles to the scale set. Each `extension` block supports the fields documented below.
- `eviction_policy` - (Optional) Specifies the eviction policy for Virtual Machines in this Scale Set. Possible values are `Deallocate` and `Delete`.

NOTE: `eviction_policy` can only be set when `priority` is set to `Low`.

- `health_probe_id` - (Optional) Specifies the identifier for the load balancer health probe. Required when using `Rolling` as your `upgrade_policy_mode`.
- `license_type` - (Optional, when a Windows machine) Specifies the Windows OS license type. If supplied, the only allowed values are `Windows_Client` and `Windows_Server`.
- `os_profile_secrets` - (Optional) A collection of Secret blocks as documented below.
- `overprovision` - (Optional) Specifies whether the virtual machine scale set should be overprovisioned.
- `plan` - (Optional) A plan block as documented below.
- `priority` - (Optional) Specifies the priority for the Virtual Machines in the Scale Set. Defaults to `Regular`. Possible values are `Low` and `Regular`.
- `rolling_upgrade_policy` - (Optional) A `rolling_upgrade_policy` block as defined below. This is only applicable when the `upgrade_policy_mode` is `Rolling`.
- `single_placement_group` - (Optional) Specifies whether the scale set is limited to a single placement group with a maximum size of 100 virtual machines. If set to `false`, managed disks must be used. Default is `true`. Changing this forces a new resource to be created. See documentation (<http://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-placement-groups>) for more information.
- `storage_profile_data_disk` - (Optional) A storage profile data disk block as documented below
- `storage_profile_image_reference` - (Optional) A storage profile image reference block as documented below.
- `tags` - (Optional) A mapping of tags to assign to the resource.
- `zones` - (Optional) A collection of availability zones to spread the Virtual Machines over.

Please Note: Availability Zones are only supported in several regions at this time (<https://docs.microsoft.com/en-us/azure/availability-zones/az-overview>).

`sku` supports the following:

- `name` - (Required) Specifies the size of virtual machines in a scale set.

- `tier` - (Optional) Specifies the tier of virtual machines in a scale set. Possible values, `standard` or `basic`.
- `capacity` - (Required) Specifies the number of virtual machines in the scale set.

`rolling_upgrade_policy` supports the following:

- `max_batch_instance_percent` - (Optional) The maximum percent of total virtual machine instances that will be upgraded simultaneously by the rolling upgrade in one batch. As this is a maximum, unhealthy instances in previous or future batches can cause the percentage of instances in a batch to decrease to ensure higher reliability. Defaults to 20.
- `max_unhealthy_instance_percent` - (Optional) The maximum percentage of the total virtual machine instances in the scale set that can be simultaneously unhealthy, either as a result of being upgraded, or by being found in an unhealthy state by the virtual machine health checks before the rolling upgrade aborts. This constraint will be checked prior to starting any batch. Defaults to 20.
- `max_unhealthy_upgraded_instance_percent` - (Optional) The maximum percentage of upgraded virtual machine instances that can be found to be in an unhealthy state. This check will happen after each batch is upgraded. If this percentage is ever exceeded, the rolling update aborts. Defaults to 20.
- `pause_time_between_batches` - (Optional) The wait time between completing the update for all virtual machines in one batch and starting the next batch. The time duration should be specified in ISO 8601 format for duration (https://en.wikipedia.org/wiki/ISO_8601#Durations (https://en.wikipedia.org/wiki/ISO_8601#Durations)). Defaults to 0 seconds represented as `PT0S`.

`identity` supports the following:

- `type` - (Required) Specifies the identity type to be assigned to the scale set. Allowable values are `SystemAssigned`, `UserAssigned`, and `SystemAssigned, UserAssigned`. For the `SystemAssigned` identity the scale set's Service Principal ID (SPN) can be retrieved after the scale set has been created. See documentation (<https://docs.microsoft.com/en-us/azure/active-directory/managed-service-identity/overview>) for more information.
- `identity_ids` - (Optional) Specifies a list of user managed identity ids to be assigned to the VMSS. Required if `type` is `UserAssigned`.

```

resource "azurerm_virtual_machine_scale_set" "test" {
  name                = "vm-scaleset"
  resource_group_name = "${azurerm_resource_group.test.name}"
  location            = "${azurerm_resource_group.test.location}"

  sku {
    name      = "${var.vm_sku}"
    tier       = "Standard"
    capacity  = "${var.instance_count}"
  }

  identity {
    type = "systemAssigned"
  }

  extension {
    name                = "MSILinuxExtension"
    publisher           = "Microsoft.ManagedIdentity"
    type               = "ManagedIdentityExtensionForLinux"
    type_handler_version = "1.0"
    settings            = "{\"port\": 50342}"
  }

  # ...
}

output "principal_id" {
  value = "${lookup(azurerm_virtual_machine_scale_set.test.identity[0], "principal_id")}"
}

```

`os_profile` supports the following:

- `computer_name_prefix` - (Required) Specifies the computer name prefix for all of the virtual machines in the scale set. Computer name prefixes must be 1 to 9 characters long for windows images and 1 - 58 for linux. Changing this forces a new resource to be created.
- `admin_username` - (Required) Specifies the administrator account name to use for all the instances of virtual machines in the scale set.
- `admin_password` - (Required) Specifies the administrator password to use for all the instances of virtual machines in a scale set.
- `custom_data` - (Optional) Specifies custom data to supply to the machine. On linux-based systems, this can be used as a cloud-init script. On other systems, this will be copied as a file on disk. Internally, Terraform will base64 encode this value before sending it to the API. The maximum length of the binary array is 65535 bytes.

`os_profile_secrets` supports the following:

- `source_vault_id` - (Required) Specifies the key vault to use.
- `vault_certificates` - (Required, on windows machines) A collection of Vault Certificates as documented below

`vault_certificates` support the following:

- `certificate_url` - (Required) It is the Base64 encoding of a JSON Object that which is encoded in UTF-8 of which the contents need to be `data`, `data_type` and `password`.

- `certificate_store` - (Required, on windows machines) Specifies the certificate store on the Virtual Machine where the certificate should be added to.

`os_profile_windows_config` supports the following:

- `provision_vm_agent` - (Optional) Indicates whether virtual machine agent should be provisioned on the virtual machines in the scale set.
- `enable_automatic_upgrades` - (Optional) Indicates whether virtual machines in the scale set are enabled for automatic updates.
- `winrm` - (Optional) A collection of WinRM configuration blocks as documented below.
- `additional_unattend_config` - (Optional) An Additional Unattended Config block as documented below.

`winrm` supports the following:

- `protocol` - (Required) Specifies the protocol of listener
- `certificate_url` - (Optional) Specifies URL of the certificate with which new Virtual Machines is provisioned.

`additional_unattend_config` supports the following:

- `pass` - (Required) Specifies the name of the pass that the content applies to. The only allowable value is `oobeSystem`.
- `component` - (Required) Specifies the name of the component to configure with the added content. The only allowable value is `Microsoft-Windows-Shell-Setup`.
- `setting_name` - (Required) Specifies the name of the setting to which the content applies. Possible values are: `FirstLogonCommands` and `AutoLogon`.
- `content` - (Optional) Specifies the base-64 encoded XML formatted content that is added to the `unattend.xml` file for the specified path and component.

`os_profile_linux_config` supports the following:

- `disable_password_authentication` - (Optional) Specifies whether password authentication should be disabled. Defaults to `false`. Changing this forces a new resource to be created.
- `ssh_keys` - (Optional) Specifies a collection of `path` and `key_data` to be placed on the virtual machine.

Note: Please note that the only allowed `path` is `/home/<username>/.ssh/authorized_keys` due to a limitation of Azure

`network_profile` supports the following:

- `name` - (Required) Specifies the name of the network interface configuration.
- `primary` - (Required) Indicates whether network interfaces created from the network interface configuration will be the primary NIC of the VM.
- `ip_configuration` - (Required) An `ip_configuration` block as documented below.
- `accelerated_networking` - (Optional) Specifies whether to enable accelerated networking or not. Defaults to `false`.
- `dns_settings` - (Optional) A `dns_settings` block as documented below.
- `ip_forwarding` - (Optional) Whether IP forwarding is enabled on this NIC. Defaults to `false`.

- `network_security_group_id` - (Optional) Specifies the identifier for the network security group.

`dns_settings` supports the following:

- `dns_servers` - (Required) Specifies an array of dns servers.

`ip_configuration` supports the following:

- `name` - (Required) Specifies name of the IP configuration.
- `subnet_id` - (Required) Specifies the identifier of the subnet.
- `application_gateway_backend_address_pool_ids` - (Optional) Specifies an array of references to backend address pools of application gateways. A scale set can reference backend address pools of multiple application gateways. Multiple scale sets cannot use the same application gateway.
- `load_balancer_backend_address_pool_ids` - (Optional) Specifies an array of references to backend address pools of load balancers. A scale set can reference backend address pools of one public and one internal load balancer. Multiple scale sets cannot use the same load balancer.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `load_balancer_inbound_nat_rules_ids` - (Optional) Specifies an array of references to inbound NAT pools for load balancers. A scale set can reference inbound nat pools of one public and one internal load balancer. Multiple scale sets cannot use the same load balancer.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `primary` - (Required) Specifies if this `ip_configuration` is the primary one.
- `application_security_group_ids` - (Optional) Specifies up to 20 application security group IDs.
- `public_ip_address_configuration` - (Optional) Describes a virtual machines scale set IP Configuration's PublicIPAddress configuration. The `public_ip_address_configuration` is documented below.

`public_ip_address_configuration` supports the following:

- `name` - (Required) The name of the public ip address configuration
- `idle_timeout` - (Required) The idle timeout in minutes. This value must be between 4 and 30.
- `domain_name_label` - (Required) The domain name label for the dns settings.

`storage_profile_os_disk` supports the following:

- `name` - (Optional) Specifies the disk name. Must be specified when using unmanaged disk ('`managed_disk_type`' property not set).
- `vhd_containers` - (Optional) Specifies the vhd uri. Cannot be used when `image` or `managed_disk_type` is specified.
- `managed_disk_type` - (Optional) Specifies the type of managed disk to create. Value you must be either `Standard_LRS`, `StandardSSD_LRS` or `Premium_LRS`. Cannot be used when `vhd_containers` or `image` is specified.

- `create_option` - (Required) Specifies how the virtual machine should be created. The only possible option is `FromImage`.
- `caching` - (Optional) Specifies the caching requirements. Possible values include: `None` (default), `ReadOnly`, `ReadWrite`.
- `image` - (Optional) Specifies the blob uri for user image. A virtual machine scale set creates an os disk in the same container as the user image. Updating the osDisk image causes the existing disk to be deleted and a new one created with the new image. If the VM scale set is in Manual upgrade mode then the virtual machines are not updated until they have manualUpgrade applied to them. When setting this field `os_type` needs to be specified. Cannot be used when `vhd_containers`, `managed_disk_type` or `storage_profile_image_reference` are specified.
- `os_type` - (Optional) Specifies the operating system Type, valid values are windows, linux.

`storage_profile_data_disk` supports the following:

- `lun` - (Required) Specifies the Logical Unit Number of the disk in each virtual machine in the scale set.
- `create_option` - (Optional) Specifies how the data disk should be created. The only possible options are `FromImage` and `Empty`.
- `caching` - (Optional) Specifies the caching requirements. Possible values include: `None` (default), `ReadOnly`, `ReadWrite`.
- `disk_size_gb` - (Optional) Specifies the size of the disk in GB. This element is required when creating an empty disk.
- `managed_disk_type` - (Optional) Specifies the type of managed disk to create. Value must be either `Standard_LRS`, `StandardSSD_LRS` or `Premium_LRS`.

`storage_profile_image_reference` supports the following:

- `id` - (Optional) Specifies the ID of the (custom) image to use to create the virtual machine scale set, as in the example below.
- `publisher` - (Optional) Specifies the publisher of the image used to create the virtual machines.
- `offer` - (Optional) Specifies the offer of the image used to create the virtual machines.
- `sku` - (Optional) Specifies the SKU of the image used to create the virtual machines.
- `version` - (Optional) Specifies the version of the image used to create the virtual machines.

`boot_diagnostics` supports the following:

- `enabled` : (Required) Whether to enable boot diagnostics for the virtual machine.
- `storage_uri` : (Required) Blob endpoint for the storage account to hold the virtual machine's diagnostic files. This must be the root of a storage account, and not a storage container.

`extension` supports the following:

- `name` - (Required) Specifies the name of the extension.
- `publisher` - (Required) The publisher of the extension, available publishers can be found by using the Azure CLI.
- `type` - (Required) The type of extension, available types for a publisher can be found using the Azure CLI.
- `type_handler_version` - (Required) Specifies the version of the extension to use, available versions can be found

using the Azure CLI.

- `auto_upgrade_minor_version` - (Optional) Specifies whether or not to use the latest minor version available.
- `provision_after_extensions` - (Optional) Specifies a dependency array of extensions required to be executed before, the array stores the name of each extension.
- `settings` - (Required) The settings passed to the extension, these are specified as a JSON object in a string.
- `protected_settings` - (Optional) The `protected_settings` passed to the extension, like `settings`, these are specified as a JSON object in a string.

`plan` supports the following:

- `name` - (Required) Specifies the name of the image from the marketplace.
- `publisher` - (Required) Specifies the publisher of the image.
- `product` - (Required) Specifies the product of the image from the marketplace.

Example of `storage_profile_image_reference` with `id`

```
resource "azurerm_image" "test" {
  name = "test"

  # ...
}

resource "azurerm_virtual_machine_scale_set" "test" {
  name = "test"

  # ...

  storage_profile_image_reference {
    id = "${azurerm_image.test.id}"
  }

  # ...
}
```

Attributes Reference

The following attributes are exported:

- `id` - The virtual machine scale set ID.

Import

Virtual Machine Scale Sets can be imported using the `resource id`, e.g.

```
terraform import azurerm_virtual_machine_scale_set.scaleset1 /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Compute/virtualMachineScaleSets/scaleset1
```

azurerm_virtual_network_gateway_connection

Manages a connection in an existing Virtual Network Gateway.

Example Usage

Site-to-Site connection

The following example shows a connection between an Azure virtual network and an on-premises VPN device and network.

```
resource "azurerm_resource_group" "test" {
  name     = "test"
  location = "West US"
}

resource "azurerm_virtual_network" "test" {
  name            = "test"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  address_space   = ["10.0.0.0/16"]
}

resource "azurerm_subnet" "test" {
  name                 = "GatewaySubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_local_network_gateway" "onpremise" {
  name            = "onpremise"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  gateway_address = "168.62.225.23"
  address_space   = ["10.1.1.0/24"]
}

resource "azurerm_public_ip" "test" {
  name            = "test"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  allocation_method = "Dynamic"
}

resource "azurerm_virtual_network_gateway" "test" {
  name            = "test"
  location        = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  type      = "Vpn"
  vpn_type = "RouteBased"

  active_active = false
}
```

```

enable_bgp    = false
sku           = "Basic"

ip_configuration {
  public_ip_address_id      = "${azurerm_public_ip.test.id}"
  private_ip_address_allocation = "Dynamic"
  subnet_id                 = "${azurerm_subnet.test.id}"
}
}

resource "azurerm_virtual_network_gateway_connection" "onpremise" {
  name                = "onpremise"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  type                = "IPsec"
  virtual_network_gateway_id = "${azurerm_virtual_network_gateway.test.id}"
  local_network_gateway_id  = "${azurerm_local_network_gateway.onpremise.id}"

  shared_key = "4-v3ry-53cr37-1p53c-5h4r3d-k3y"
}

```

VNet-to-VNet connection

The following example shows a connection between two Azure virtual network in different locations/regions.

```

resource "azurerm_resource_group" "us" {
  name     = "us"
  location = "East US"
}

resource "azurerm_virtual_network" "us" {
  name                = "us"
  location             = "${azurerm_resource_group.us.location}"
  resource_group_name = "${azurerm_resource_group.us.name}"
  address_space       = ["10.0.0.0/16"]
}

resource "azurerm_subnet" "us_gateway" {
  name                = "GatewaySubnet"
  resource_group_name = "${azurerm_resource_group.us.name}"
  virtual_network_name = "${azurerm_virtual_network.us.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "us" {
  name                = "us"
  location             = "${azurerm_resource_group.us.location}"
  resource_group_name = "${azurerm_resource_group.us.name}"
  allocation_method   = "Dynamic"
}

resource "azurerm_virtual_network_gateway" "us" {
  name                = "us-gateway"
  location             = "${azurerm_resource_group.us.location}"
  resource_group_name = "${azurerm_resource_group.us.name}"
}

```

```

type      = "Vpn"
vpn_type  = "RouteBased"
sku       = "Basic"

ip_configuration {
  public_ip_address_id      = "${azurerm_public_ip.us.id}"
  private_ip_address_allocation = "Dynamic"
  subnet_id                 = "${azurerm_subnet.us_gateway.id}"
}
}

resource "azurerm_resource_group" "europe" {
  name     = "europe"
  location = "West Europe"
}

resource "azurerm_virtual_network" "europe" {
  name                = "europe"
  location             = "${azurerm_resource_group.europe.location}"
  resource_group_name = "${azurerm_resource_group.europe.name}"
  address_space       = ["10.1.0.0/16"]
}

resource "azurerm_subnet" "europe_gateway" {
  name                = "GatewaySubnet"
  resource_group_name = "${azurerm_resource_group.europe.name}"
  virtual_network_name = "${azurerm_virtual_network.europe.name}"
  address_prefix       = "10.1.1.0/24"
}

resource "azurerm_public_ip" "europe" {
  name                = "europe"
  location             = "${azurerm_resource_group.europe.location}"
  resource_group_name = "${azurerm_resource_group.europe.name}"
  allocation_method   = "Dynamic"
}

resource "azurerm_virtual_network_gateway" "europe" {
  name                = "europe-gateway"
  location             = "${azurerm_resource_group.europe.location}"
  resource_group_name = "${azurerm_resource_group.europe.name}"

  type      = "Vpn"
  vpn_type  = "RouteBased"
  sku       = "Basic"

  ip_configuration {
    public_ip_address_id      = "${azurerm_public_ip.europe.id}"
    private_ip_address_allocation = "Dynamic"
    subnet_id                 = "${azurerm_subnet.europe_gateway.id}"
  }
}

resource "azurerm_virtual_network_gateway_connection" "us_to_europe" {
  name                = "us-to-europe"
  location             = "${azurerm_resource_group.us.location}"
  resource_group_name = "${azurerm_resource_group.us.name}"

  type      = "Vnet2Vnet"

```

```

type                                = "Vnet2Vnet"
virtual_network_gateway_id          = "${azurerm_virtual_network_gateway.us.id}"
peer_virtual_network_gateway_id     = "${azurerm_virtual_network_gateway.europe.id}"

shared_key = "4-v3ry-53cr37-1p53c-5h4r3d-k3y"
}

resource "azurerm_virtual_network_gateway_connection" "europe_to_us" {
  name                = "europe-to-us"
  location            = "${azurerm_resource_group.europe.location}"
  resource_group_name = "${azurerm_resource_group.europe.name}"

  type                                = "Vnet2Vnet"
  virtual_network_gateway_id          = "${azurerm_virtual_network_gateway.europe.id}"
  peer_virtual_network_gateway_id     = "${azurerm_virtual_network_gateway.us.id}"

  shared_key = "4-v3ry-53cr37-1p53c-5h4r3d-k3y"
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the connection. Changing the name forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the connection. Changing the name forces a new resource to be created.
- `location` - (Required) The location/region where the connection is located. Changing this forces a new resource to be created.
- `type` - (Required) The type of connection. Valid options are `IPsec` (Site-to-Site), `ExpressRoute` (ExpressRoute), and `Vnet2Vnet` (VNet-to-VNet). Each connection type requires different mandatory arguments (refer to the examples above). Changing the connection type will force a new connection to be created.
- `virtual_network_gateway_id` - (Required) The ID of the Virtual Network Gateway in which the connection will be created. Changing the gateway forces a new resource to be created.
- `authorization_key` - (Optional) The authorization key associated with the Express Route Circuit. This field is required only if the type is an ExpressRoute connection.
- `express_route_circuit_id` - (Optional) The ID of the Express Route Circuit when creating an ExpressRoute connection (i.e. when `type` is `ExpressRoute`). The Express Route Circuit can be in the same or in a different subscription.
- `peer_virtual_network_gateway_id` - (Optional) The ID of the peer virtual network gateway when creating a VNet-to-VNet connection (i.e. when `type` is `Vnet2Vnet`). The peer Virtual Network Gateway can be in the same or in a different subscription.
- `local_network_gateway_id` - (Optional) The ID of the local network gateway when creating Site-to-Site connection (i.e. when `type` is `IPsec`).
- `routing_weight` - (Optional) The routing weight. Defaults to `10`.

- `shared_key` - (Optional) The shared IPsec key. A key must be provided if a Site-to-Site or VNet-to-VNet connection is created whereas ExpressRoute connections do not need a shared key.
- `enable_bgp` - (Optional) If `true`, BGP (Border Gateway Protocol) is enabled for this connection. Defaults to `false`.
- `express_route_gateway_bypass` - (Optional) If `true`, data packets will bypass ExpressRoute Gateway for data forwarding. This is only valid for ExpressRoute connections.
- `use_policy_based_traffic_selectors` - (Optional) If `true`, policy-based traffic selectors are enabled for this connection. Enabling policy-based traffic selectors requires an `ipsec_policy` block. Defaults to `false`.
- `ipsec_policy` (Optional) A `ipsec_policy` block which is documented below. Only a single policy can be defined for a connection. For details on custom policies refer to the relevant section in the Azure documentation (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-ipsecikepolicy-rm-powershell>).
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `ipsec_policy` block supports:

- `dh_group` - (Required) The DH group used in IKE phase 1 for initial SA. Valid options are `DHGroup1`, `DHGroup14`, `DHGroup2`, `DHGroup2048`, `DHGroup24`, `ECP256`, `ECP384`, or `None`.
- `ike_encryption` - (Required) The IKE encryption algorithm. Valid options are `AES128`, `AES192`, `AES256`, `DES`, or `DES3`.
- `ike_integrity` - (Required) The IKE integrity algorithm. Valid options are `MD5`, `SHA1`, `SHA256`, or `SHA384`.
- `ipsec_encryption` - (Required) The IPsec encryption algorithm. Valid options are `AES128`, `AES192`, `AES256`, `DES`, `DES3`, `GCMAES128`, `GCMAES192`, `GCMAES256`, or `None`.
- `ipsec_integrity` - (Required) The IPsec integrity algorithm. Valid options are `GCMAES128`, `GCMAES192`, `GCMAES256`, `MD5`, `SHA1`, or `SHA256`.
- `pfs_group` - (Required) The DH group used in IKE phase 2 for new child SA. Valid options are `ECP256`, `ECP384`, `PFS1`, `PFS2`, `PFS2048`, `PFS24`, or `None`.
- `sa_datasize` - (Optional) The IPsec SA payload size in KB. Must be at least `1024` KB. Defaults to `102400000` KB.
- `sa_lifetime` - (Optional) The IPsec SA lifetime in seconds. Must be at least `300` seconds. Defaults to `27000` seconds.

Attributes Reference

The following attributes are exported:

- `id` - The connection ID.

Import

Virtual Network Gateway Connections can be imported using their `resource_id`, e.g.

```
terraform import azurerm_virtual_network_gateway_connection.testConnection /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myGroup1/providers/Microsoft.Network/connections/myConnection1
```

azurerm_virtual_network_gateway

Manages a Virtual Network Gateway to establish secure, cross-premises connectivity.

Note: Please be aware that provisioning a Virtual Network Gateway takes a long time (between 30 minutes and 1 hour)

Example Usage

```
resource "azurerm_resource_group" "test" {
  name     = "test"
  location = "West US"
}

resource "azurerm_virtual_network" "test" {
  name                = "test"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  address_space       = ["10.0.0.0/16"]
}

resource "azurerm_subnet" "test" {
  name                 = "GatewaySubnet"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test.name}"
  address_prefix       = "10.0.1.0/24"
}

resource "azurerm_public_ip" "test" {
  name                = "test"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  allocation_method = "Dynamic"
}

resource "azurerm_virtual_network_gateway" "test" {
  name                = "test"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"

  type      = "Vpn"
  vpn_type  = "RouteBased"

  active_active = false
  enable_bgp    = false
  sku           = "Basic"

  ip_configuration {
    name                 = "vnetGatewayConfig"
    public_ip_address_id = "${azurerm_public_ip.test.id}"
    private_ip_address_allocation = "Dynamic"
    subnet_id            = "${azurerm_subnet.test.id}"
  }
}
```

```

vpn_client_configuration {
  address_space = ["10.2.0.0/24"]

  root_certificate {
    name = "DigiCert-Federated-ID-Root-CA"

    public_cert_data = <<EOF
MIIDuzCCAqOgAwIBAgIQCHTZWCM+IlFFIRXIvyKSrjANBgkqhkiG9w0BAQsFADBN
MQswCQYDVQQGEwJVUzEVMBMGA1UEChMMRGlnaUNlcnQgSW5jMRkwFwYDVQQLExB3
d3cuZGlnaWNLcnQuY29tMSYwJAYDVQQDEx1EaWdpQ2VydCBGZWRLcmF0ZWQgSUQg
Um9vdCBDQTAeFw0xMzAxMTUxMjAwMDBaFw0zMTUxMjAwMDBaMGcxChZAJBgNV
BAYTAlVTMRUwEwYDVQQKEwxEaWdpQ2VydCBJbmMxGTAXBgNVBAsTEHd3dy5kaWdp
Y2VydC5jb20xJjAkBgNVBAMTHURpZ2lDZXJ0IEZlZGVyYXRlZCBJRCB5b290IENB
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAvAEB4pcCqnNNOWE6Ur5j
QPUH+1y1F9KdHTRSza6k5iDlXq1kGS1qAkuKtw9JsiNRrjltmFnzMZRBbX8Tlfl8
zAhBmb6dDduDGED01kBsTKgywYPxXVTKec0WxYEEF0oMn4wSYNl0lt2eJAKHXjNf
GTwiibdP8CUR2ghSM2sUTI8Nt10mfc4SMHhGhYD64uJMbX98THQ/4LMGuYegou+d
GTiahfHtjn7AboSEknamAMJHCh5RLYZZ6B104QbKJ+34Q0eKgnI3X6Vc9u0zf6DH8
Dk+4zQDYRRtQtNvO3VT8jzqDlCRuNtq6YvryOWN74/dq8LQhUnXHvFyrSdMaE1X2
DwIDAQABo2MwYTAPBgNVHRMBAf8EBTADAQH/MA4GA1UdDwEB/wQEAwIBhjAdBgNV
HQ4EFgQUGRdkFnbGt1EWjKwbUne+50aZvRYwHwYDVR0jBBgwFoAUGRdkFnbGt1EW
jKwbUne+50aZvRYwDQYJKoZIhvcNAQELBQADggEBAHcqsHkrjpESqfuVTRiptJfP
9JbdtWqRTmOf6uJi2c8YVqI6XlKXsD8C1dUUaaHKLUIJzvKiazibVuBwMIT84AyqR
QELn3e0BtgEymEygMU569b01ZPxoFSnNXc7qDZBDef8WfqAV/sxkTi8L9BkmFYfL
uGLOhRJ0FprPdoDIUBB+tmCl3oDcBy3vnUe0Eioz8zAkprcb3GHwHAK+vHmfmfgcn
WsFMLH4JCLa/tRYL+Rw/N3ybCkDp00s0WUz+AoDywSl0Q/ZEnNY0MsFiw6LyIdbq
M/s/1JRt03bDSzD9TazRVzn2oBqzSa8VgIo5C1n0noAKJTlsClJKvIhnRlaLQqk=
EOF
  }

  revoked_certificate {
    name          = "Verizon-Global-Root-CA"
    thumbprint    = "912198EEF23DCAC40939312FEE97DD560BAE49B1"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Virtual Network Gateway. Changing the name forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Virtual Network Gateway. Changing the resource group name forces a new resource to be created.
- `location` - (Required) The location/region where the Virtual Network Gateway is located. Changing the location/region forces a new resource to be created.
- `type` - (Required) The type of the Virtual Network Gateway. Valid options are `Vpn` or `ExpressRoute`. Changing the type forces a new resource to be created.
- `vpn_type` - (Optional) The routing type of the Virtual Network Gateway. Valid options are `RouteBased` or `PolicyBased`. Defaults to `RouteBased`.

- `enable_bgp` - (Optional) If `true`, BGP (Border Gateway Protocol) will be enabled for this Virtual Network Gateway. Defaults to `false`.
- `active_active` - (Optional) If `true`, an active-active Virtual Network Gateway will be created. An active-active gateway requires a `HighPerformance` or an `UltraPerformance` sku. If `false`, an active-standby gateway will be created. Defaults to `false`.
- `default_local_network_gateway_id` - (Optional) The ID of the local network gateway through which outbound Internet traffic from the virtual network in which the gateway is created will be routed (*forced tunneling*). Refer to the Azure documentation on forced tunneling (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-forced-tunneling-rm>). If not specified, forced tunneling is disabled.
- `sku` - (Required) Configuration of the size and capacity of the virtual network gateway. Valid options are `Basic`, `Standard`, `HighPerformance`, `UltraPerformance`, `ErGw1AZ`, `ErGw2AZ`, `ErGw3AZ`, `VpnGw1`, `VpnGw2`, `VpnGw3`, `VpnGw1AZ`, `VpnGw2AZ`, and `VpnGw3AZ` and depend on the `type` and `vpn_type` arguments. A `PolicyBased` gateway only supports the `Basic` sku. Further, the `UltraPerformance` sku is only supported by an `ExpressRoute` gateway.
- `ip_configuration` (Required) One or two `ip_configuration` blocks documented below. An active-standby gateway requires exactly one `ip_configuration` block whereas an active-active gateway requires exactly two `ip_configuration` blocks.
- `vpn_client_configuration` (Optional) A `vpn_client_configuration` block which is documented below. In this block the Virtual Network Gateway can be configured to accept IPsec point-to-site connections.
- `tags` - (Optional) A mapping of tags to assign to the resource.

The `ip_configuration` block supports:

- `name` - (Optional) A user-defined name of the IP configuration. Defaults to `vnetGatewayConfig`.
- `private_ip_address_allocation` - (Optional) Defines how the private IP address of the gateways virtual interface is assigned. Valid options are `Static` or `Dynamic`. Defaults to `Dynamic`.
- `subnet_id` - (Required) The ID of the gateway subnet of a virtual network in which the virtual network gateway will be created. It is mandatory that the associated subnet is named `GatewaySubnet`. Therefore, each virtual network can contain at most a single Virtual Network Gateway.
- `public_ip_address_id` - (Optional) The ID of the public ip address to associate with the Virtual Network Gateway.

The `vpn_client_configuration` block supports:

- `address_space` - (Required) The address space out of which ip addresses for vpn clients will be taken. You can provide more than one address space, e.g. in CIDR notation.
- `root_certificate` - (Optional) One or more `root_certificate` blocks which are defined below. These root certificates are used to sign the client certificate used by the VPN clients to connect to the gateway. This setting is incompatible with the use of `radius_server_address` and `radius_server_secret`.
- `revoked_certificate` - (Optional) One or more `revoked_certificate` blocks which are defined below. This setting is incompatible with the use of `radius_server_address` and `radius_server_secret`.
- `radius_server_address` - (Optional) The address of the Radius server. This setting is incompatible with the use of `root_certificate` and `revoked_certificate`.
- `radius_server_secret` - (Optional) The secret used by the Radius server. This setting is incompatible with the use of

`root_certificate` and `revoked_certificate`.

- `vpn_client_protocols` - (Optional) List of the protocols supported by the vpn client. The supported values are SSTP, IKEV2 and OpenVPN.

The `bgp_settings` block supports:

- `asn` - (Optional) The Autonomous System Number (ASN) to use as part of the BGP.
- `peering_address` - (Optional) The BGP peer IP address of the virtual network gateway. This address is needed to configure the created gateway as a BGP Peer on the on-premises VPN devices. The IP address must be part of the subnet of the Virtual Network Gateway. Changing this forces a new resource to be created.
- `peer_weight` - (Optional) The weight added to routes which have been learned through BGP peering. Valid values can be between 0 and 100.

The `root_certificate` block supports:

- `name` - (Required) A user-defined name of the root certificate.
- `public_cert_data` - (Required) The public certificate of the root certificate authority. The certificate must be provided in Base-64 encoded X.509 format (PEM). In particular, this argument *must not* include the `-----BEGIN CERTIFICATE-----` or `-----END CERTIFICATE-----` markers.

The `root_revoked_certificate` block supports:

- `name` - (Required) A user-defined name of the revoked certificate.
- `public_cert_data` - (Required) The SHA1 thumbprint of the certificate to be revoked.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual Network Gateway.

Import

Virtual Network Gateways can be imported using the `resource id`, e.g.

```
terraform import azurerm_virtual_network_gateway.testGateway /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myGroup1/providers/Microsoft.Network/virtualNetworkGateways/myGateway1
```

azurerm_virtual_network

Manages a virtual network including any configured subnets. Each subnet can optionally be configured with a security group to be associated with the subnet.

NOTE on Virtual Networks and Subnet's: Terraform currently provides both a standalone Subnet resource (</docs/providers/azurerm/r/subnet.html>), and allows for Subnets to be defined in-line within the Virtual Network resource ([/docs/providers/azurerm/r/virtual_network.html](/docs/providers/azurerm/r/virtual_network.html)). At this time you cannot use a Virtual Network with in-line Subnets in conjunction with any Subnet resources. Doing so will cause a conflict of Subnet configurations and will overwrite Subnet's.

Example Usage

```

resource "azurerm_resource_group" "test" {
  name      = "acceptanceTestResourceGroup1"
  location  = "West US"
}

resource "azurerm_network_security_group" "test" {
  name                = "acceptanceTestSecurityGroup1"
  location             = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_ddos_protection_plan" "test" {
  name                = "ddospplan1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
}

resource "azurerm_virtual_network" "test" {
  name                = "virtualNetwork1"
  location            = "${azurerm_resource_group.test.location}"
  resource_group_name = "${azurerm_resource_group.test.name}"
  address_space       = ["10.0.0.0/16"]
  dns_servers         = ["10.0.0.4", "10.0.0.5"]

  ddos_protection_plan {
    id      = "${azurerm_ddos_protection_plan.test.id}"
    enable = true
  }

  subnet {
    name          = "subnet1"
    address_prefix = "10.0.1.0/24"
  }

  subnet {
    name          = "subnet2"
    address_prefix = "10.0.2.0/24"
  }

  subnet {
    name          = "subnet3"
    address_prefix = "10.0.3.0/24"
    security_group = "${azurerm_network_security_group.test.id}"
  }

  tags = {
    environment = "Production"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the virtual network. Changing this forces a new resource to be created.

- `resource_group_name` - (Required) The name of the resource group in which to create the virtual network.
 - `address_space` - (Required) The address space that is used the virtual network. You can supply more than one address space. Changing this forces a new resource to be created.
 - `location` - (Required) The location/region where the virtual network is created. Changing this forces a new resource to be created.
 - `ddos_protection_plan` - (Optional) A `ddos_protection_plan` block as documented below.
 - `dns_servers` - (Optional) List of IP addresses of DNS servers
 - `subnet` - (Optional) Can be specified multiple times to define multiple subnets. Each `subnet` block supports fields documented below.
 - `tags` - (Optional) A mapping of tags to assign to the resource.
-

A `ddos_protection_plan` block supports the following:

- `id` - (Required) The Resource ID of DDoS Protection Plan.
 - `enable` - (Required) Enable/disable DDoS Protection Plan on Virtual Network.
-

The `subnet` block supports:

- `name` - (Required) The name of the subnet.
- `address_prefix` - (Required) The address prefix to use for the subnet.
- `security_group` - (Optional) The Network Security Group to associate with the subnet. (Referenced by `id`, ie. `azurerm_network_security_group.test.id`)

Attributes Reference

The following attributes are exported:

- `id` - The virtual NetworkConfiguration ID.
 - `name` - The name of the virtual network.
 - `resource_group_name` - The name of the resource group in which to create the virtual network.
 - `location` - The location/region where the virtual network is created
 - `address_space` - The address space that is used the virtual network.
 - `subnet` - One or more `subnet` blocks as defined below.
-

The `subnet` block exports:

- `id` - The ID of this subnet.

Import

Virtual Networks can be imported using the `resource_id`, e.g.

```
terraform import azurerm_virtual_network.testNetwork /subscriptions/00000000-0000-0000-0000-000000000000/  
resourceGroups/mygroup1/providers/Microsoft.Network/virtualNetworks/myvnet1
```

azurerm_virtual_network_peering

Manages a virtual network peering which allows resources to access other resources in the linked virtual network.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "peeredvnets-rg"
  location  = "West US"
}

resource "azurerm_virtual_network" "test1" {
  name                = "peternetwork1"
  resource_group_name = "${azurerm_resource_group.test.name}"
  address_space       = ["10.0.1.0/24"]
  location            = "West US"
}

resource "azurerm_virtual_network" "test2" {
  name                = "peternetwork2"
  resource_group_name = "${azurerm_resource_group.test.name}"
  address_space       = ["10.0.2.0/24"]
  location            = "West US"
}

resource "azurerm_virtual_network_peering" "test1" {
  name                = "peer1to2"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test1.name}"
  remote_virtual_network_id = "${azurerm_virtual_network.test2.id}"
}

resource "azurerm_virtual_network_peering" "test2" {
  name                = "peer2to1"
  resource_group_name = "${azurerm_resource_group.test.name}"
  virtual_network_name = "${azurerm_virtual_network.test2.name}"
  remote_virtual_network_id = "${azurerm_virtual_network.test1.id}"
}
```

Example Usage (Global virtual network peering)

```

variable "location" {
  default = [
    "uksouth",
    "southeastasia",
  ]
}

variable "vnet_address_space" {
  default = [
    "10.0.0.0/16",
    "10.1.0.0/16",
  ]
}

resource "azurerm_resource_group" "vnet" {
  count    = "${length(var.location)}"
  name     = "rg-global-vnet-peering-${count.index}"
  location = "${element(var.location, count.index)}"
}

resource "azurerm_virtual_network" "vnet" {
  count            = "${length(var.location)}"
  name            = "vnet-${count.index}"
  resource_group_name = "${element(azurerm_resource_group.vnet.*, count.index)}"
  address_space    = ["${element(var.vnet_address_space, count.index)}"]
  location        = "${element(azurerm_resource_group.vnet.*, count.index)}"
}

resource "azurerm_subnet" "nva" {
  count            = "${length(var.location)}"
  name            = "nva"
  resource_group_name = "${element(azurerm_resource_group.vnet.*, count.index)}"
  virtual_network_name = "${element(azurerm_virtual_network.vnet.*, count.index)}"
  address_prefix    = "${cidrsubnet("${element(azurerm_virtual_network.vnet.*, address_space[count.index])", 13, 0)}" # /29
}

# enable global peering between the two virtual network
resource "azurerm_virtual_network_peering" "peering" {
  count            = "${length(var.location)}"
  name            = "peering-to-${element(azurerm_virtual_network.vnet.*, name, 1 - count.index)}"
  resource_group_name = "${element(azurerm_resource_group.vnet.*, count.index)}"
  virtual_network_name = "${element(azurerm_virtual_network.vnet.*, count.index)}"
  remote_virtual_network_id = "${element(azurerm_virtual_network.vnet.*, id, 1 - count.index)}"
  allow_virtual_network_access = true
  allow_forwarded_traffic      = true

  # `allow_gateway_transit` must be set to false for vnet Global Peering
  allow_gateway_transit = false
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the virtual network peering. Changing this forces a new resource to be created.
- `virtual_network_name` - (Required) The name of the virtual network. Changing this forces a new resource to be created.
- `remote_virtual_network_id` - (Required) The full Azure resource ID of the remote virtual network. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the virtual network. Changing this forces a new resource to be created.
- `allow_virtual_network_access` - (Optional) Controls if the VMs in the remote virtual network can access VMs in the local virtual network. Defaults to `false`.
- `allow_forwarded_traffic` - (Optional) Controls if forwarded traffic from VMs in the remote virtual network is allowed. Defaults to `false`.
- `allow_gateway_transit` - (Optional) Controls gatewayLinks can be used in the remote virtual network's link to the local virtual network.
- `use_remote_gateways` - (Optional) Controls if remote gateways can be used on the local virtual network. If the flag is set to `true`, and `allow_gateway_transit` on the remote peering is also `true`, virtual network will use gateways of remote virtual network for transit. Only one peering can have this flag set to `true`. This flag cannot be set if virtual network already has a gateway. Defaults to `false`.

NOTE: `use_remote_gateways` must be set to `false` if using Global Virtual Network Peerings.

Attributes Reference

The following attributes are exported:

- `id` - The Virtual Network Peering resource ID.

Note

Virtual Network peerings cannot be created, updated or deleted concurrently.

Import

Virtual Network Peerings can be imported using the `resource id`, e.g.

```
terraform import azurerm_virtual_network_peering.testPeering /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/providers/Microsoft.Network/virtualNetworks/myvnet1/virtualNetworkPeerings/myvnet1peering
```

azurerm_virtual_wan

Manages a Virtual WAN.

Example Usage

```
resource "azurerm_resource_group" "test" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_wan" "test" {
  name                        = "example-vwan"
  resource_group_name        = "${azurerm_resource_group.test.name}"
  location                   = "${azurerm_resource_group.test.location}"
}
```

Argument Reference

The following arguments are supported:

- `name` - (Required) Specifies the name of the Virtual WAN. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the resource group in which to create the Virtual WAN. Changing this forces a new resource to be created.
- `location` - (Required) Specifies the supported Azure location where the resource exists. Changing this forces a new resource to be created.
- `disable_vpn_encryption` - (Optional) Boolean flag to specify whether VPN encryption is disabled. Defaults to `false`.
- `security_provider_name` - (Optional) The name of the Security Provider.
- `allow_branch_to_branch_traffic` - (Optional) Boolean flag to specify whether branch to branch traffic is allowed. Defaults to `true`.
- `allow_vnet_to_vnet_traffic` - (Optional) Boolean flag to specify whether VNet to VNet traffic is allowed. Defaults to `false`.
- `office365_local_breakout_category` - (Optional) Specifies the Office365 local breakout category. Possible values include: `Optimize`, `OptimizeAndAllow`, `All`, `None`. Defaults to `None`.
- `tags` - (Optional) A mapping of tags to assign to the Virtual WAN.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the Virtual WAN.

Import

Virtual WAN can be imported using the `resource id`, e.g.

```
terraform import azurerm_virtual_wan.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/group1/providers/Microsoft.Network/virtualWans/testvwan
```

azurerm_web_application_firewall_policy

Manages a Azure Web Application Firewall Policy instance.

Example Usage

```

resource "azurerm_resource_group" "example" {
  name      = "example-rg"
  location  = "West US 2"
}

resource "azurerm_web_application_firewall_policy" "example" {
  name                        = "example-wafpolicy"
  resource_group_name        = "${azurerm_resource_group.example.name}"
  location                    = "${azurerm_resource_group.example.location}"

  custom_rules {
    name      = "Rule1"
    priority  = 1
    rule_type = "MatchRule"

    match_conditions {
      match_variables {
        variable_name = "RemoteAddr"
      }

      operator      = "IPMatch"
      negation_condition = false
      match_values  = ["192.168.1.0/24", "10.0.0.0/24"]
    }

    action = "Block"
  }

  custom_rules {
    name      = "Rule2"
    priority  = 2
    rule_type = "MatchRule"

    match_conditions {
      match_variables {
        variable_name = "RemoteAddr"
      }

      operator      = "IPMatch"
      negation_condition = false
      match_values  = ["192.168.1.0/24"]
    }

    match_conditions {
      match_variables {
        variable_name = "RequestHeaders"
        selector      = "UserAgent"
      }

      operator      = "Contains"
      negation_condition = false
      match_values  = ["Windows"]
    }

    action = "Block"
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the policy. Changing this forces a new resource to be created.
 - `resource_group_name` - (Required) The name of the resource group. Changing this forces a new resource to be created.
 - `location` - (Optional) Resource location. Changing this forces a new resource to be created.
 - `custom_rules` - (Optional) One or more `custom_rule` blocks as defined below.
 - `policy_settings` - (Optional) A `policy_setting` block as defined below.
 - `tags` - (Optional) A mapping of tags to assign to the Web Application Firewall Policy.
-

The `custom_rule` block supports the following:

- `name` - (Optional) Gets name of the resource that is unique within a policy. This name can be used to access the resource.
 - `priority` - (Required) Describes priority of the rule. Rules with a lower value will be evaluated before rules with a higher value
 - `rule_type` - (Required) Describes the type of rule
 - `match_conditions` - (Required) One or more `match_condition` block defined below.
 - `action` - (Required) Type of Actions
-

The `match_condition` block supports the following:

- `match_variables` - (Required) One or more `match_variable` block defined below.
 - `operator` - (Required) Describes operator to be matched
 - `negation_condition` - (Optional) Describes if this is negate condition or not
 - `match_values` - (Required) Match value
-

The `match_variable` block supports the following:

- `variable_name` - (Required) The name of the Match Variable
 - `selector` - (Optional) Describes field of the matchVariable collection
-

The `policy_setting` block supports the following:

- `enabled` - (Optional) Describes if the policy is in enabled state or disabled state Defaults to `Enabled` .
- `mode` - (Optional) Describes if it is in detection mode or prevention mode at the policy level Defaults to `Prevention` .

Attributes Reference

The following attributes are exported:

- `id` - Resource ID.

Import

Web Application Firewall Policy can be imported using the `resource id`, e.g.

```
$ terraform import azurerm_web_application_firewall_policy.example /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/example-rg/providers/Microsoft.Network/ApplicationGatewayWebApplicationFirewallPolicies/example-wafpolicy
```

azurerm_windows_virtual_machine_scale_set

NOTE: This resource is in Beta and as such the Schema can change in Minor versions of the Provider.

NOTE: All arguments including the administrator login and password will be stored in the raw state as plain-text. Read more about sensitive data in state (</docs/state/sensitive-data.html>).

Manages a Windows Virtual Machine Scale Set.

NOTE: This resource does not support Unmanaged Disks. If you need to use Unmanaged Disks you can continue to use the `azurerm_virtual_machine_scale_set` resource ([/docs/providers/azurerm/r/virtual_machine_scale_set.html](/docs/providers/azurerm/r/virtual_machine_scale_set.html)) instead

NOTE: Terraform will automatically update & reimage the nodes in the Scale Set if Required during an Update (for example, when changing Sku) - you can opt out of this by setting the `terraform_should_roll_instances_when_required` field to `false`.

Example Usage

This example provisions a basic Windows Virtual Machine Scale Set on an internal network. Additional examples of how to use the `azurerm_windows_virtual_machine_scale_set` resource can be found in the `./examples/vm-scale-set/windows` directory within the Github Repository (<https://github.com/terraform-providers/terraform-provider-azurerm/tree/master/examples/vm-scale-set/windows>).

```

resource "azurerm_resource_group" "example" {
  name      = "example-resources"
  location  = "West Europe"
}

resource "azurerm_virtual_network" "test" {
  name                = "example-network"
  resource_group_name = azurerm_resource_group.example.name
  location            = azurerm_resource_group.example.location
  address_space       = ["10.0.0.0/16"]
}

resource "azurerm_subnet" "internal" {
  name                = "internal"
  resource_group_name = azurerm_resource_group.example.name
  virtual_network_name = azurerm_virtual_network.example.name
  address_prefix      = "10.0.2.0/24"
}

resource "azurerm_windows_virtual_machine_scale_set" "example" {
  name                = "example-vmss"
  resource_group_name = azurerm_resource_group.example.name
  location            = azurerm_resource_group.example.location
  sku                 = "Standard_F2"
  instances           = 1
  admin_password      = "P@55w0rd1234!"
  admin_username      = "adminuser"

  source_image_reference {
    publisher = "MicrosoftWindowsServer"
    offer     = "WindowsServer"
    sku       = "2016-Datacenter-Server-Core"
    version   = "latest"
  }

  os_disk {
    storage_account_type = "Standard_LRS"
    caching               = "ReadWrite"
  }

  network_interface {
    name      = "example"
    primary   = true

    ip_configuration {
      name          = "internal"
      primary       = true
      subnet_id     = azurerm_subnet.internal.id
    }
  }
}

```

Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the Windows Virtual Machine Scale Set. Changing this forces a new resource to be created.
- `location` - (Required) The Azure location where the Windows Virtual Machine Scale Set should exist. Changing this forces a new resource to be created.
- `resource_group_name` - (Required) The name of the Resource Group in which the Windows Virtual Machine Scale Set should be exist. Changing this forces a new resource to be created.
- `admin_password` - (Required) The Password which should be used for the local-administrator on this Virtual Machine. Changing this forces a new resource to be created.
- `admin_username` - (Required) The username of the local administrator on each Virtual Machine Scale Set instance. Changing this forces a new resource to be created.
- `instances` - (Required) The number of Virtual Machines in the Scale Set.

NOTE: If you're using AutoScaling, you may wish to use Terraform's `ignore_changes` functionality (https://www.terraform.io/docs/configuration/resources.html#ignore_changes) to ignore changes to this field.

- `sku` - (Required) The Virtual Machine SKU for the Scale Set, such as `Standard_F2`.
- `network_interface` - (Required) One or more `network_interface` blocks as defined below.
- `os_disk` - (Required) An `os_disk` block as defined below.

-
- `additional_capabilities` - (Optional) A `additional_capabilities` block as defined below.
 - `additional_unattend_config` - (Optional) One or more `additional_unattend_config` blocks as defined below.
 - `automatic_os_upgrade_policy` - (Optional) A `automatic_os_upgrade_policy` block as defined below. This is Required and can only be specified when `upgrade_mode` is set to `Automatic`.
 - `boot_diagnostics` - (Optional) A `boot_diagnostics` block as defined below.
 - `computer_name_prefix` - (Optional) The prefix which should be used for the name of the Virtual Machines in this Scale Set. If unspecified this defaults to the value for the `name` field.
 - `custom_data` - (Optional) The Base64-Encoded Custom Data which should be used for this Virtual Machine Scale Set.

NOTE: When Custom Data has been configured, it's not possible to remove it without tainting the Virtual Machine Scale Set, due to a limitation of the Azure API.

- `data_disk` - (Optional) One or more `data_disk` blocks as defined below.
- `do_not_run_extensions_on_overprovisioned_machines` - (Optional) Should Virtual Machine Extensions be run on Overprovisioned Virtual Machines in the Scale Set? Defaults to `false`.
- `enable_automatic_updates` - (Optional) Are automatic updates enabled for this Virtual Machine? Defaults to `true`.
- `eviction_policy` - (Optional) The Policy which should be used Virtual Machines are Evicted from the Scale Set. Changing this forces a new resource to be created.

NOTE: This can only be configured when `priority` is set to `Low`.

- `health_probe_id` - (Optional) The ID of a Load Balancer Probe which should be used to determine the health of an instance. Changing this forces a new resource to be created. This is Required and can only be specified when `upgrade_mode` is set to `Automatic` or `Rolling`.
- `identity` - (Optional) A `identity` block as defined below.
- `license_type` - (Optional) Specifies the type of on-premise license (also known as Azure Hybrid Use Benefit (<https://docs.microsoft.com/azure/virtual-machines/virtual-machines-windows-hybrid-use-benefit-licensing>)) which should be used for this Virtual Machine Scale Set. Possible values are `Windows_Client` and `Windows_Server`. Changing this forces a new resource to be created.
- `max_bid_price` - (Optional) The maximum price you're willing to pay for a low-priority VM Scale Set, in US Dollars; which must be greater than the current low-priority price. If this bid price falls below the current low-priority price the Virtual Machines in the Scale Set will be evicted using the `eviction_policy`. Defaults to `-1`, which means that this VM Scale Set should not be evicted for price reasons.

NOTE: This can only be configured when `priority` is set to `Low`.

- `overprovision` - (Optional) Should Azure over-provision Virtual Machines in this Scale Set? This means that multiple Virtual Machines will be provisioned and Azure will keep the instances which become available first - which improves provisioning success rates and improves deployment time. You're not billed for these over-provisioned VM's and they don't count towards the Subscription Quota. Defaults to `false`.
- `priority` - (Optional) The Priority of this Virtual Machine Scale Set. Possible values are `Regular` and `Low`. Defaults to `Regular`. Changing this value forces a new resource.

NOTE: When `priority` is set to `Low` an `eviction_policy` must be specified.

- `provision_vm_agent` - (Optional) Should the Azure VM Agent be provisioned on each Virtual Machine in the Scale Set? Defaults to `true`. Changing this value forces a new resource to be created.
- `proximity_placement_group_id` - (Optional) The ID of the Proximity Placement Group in which the Virtual Machine Scale Set should be assigned to. Changing this forces a new resource to be created.
- `rolling_upgrade_policy` - (Optional) A `rolling_upgrade_policy` block as defined below. This is Required and can only be specified when `upgrade_mode` is set to `Automatic` or `Rolling`.
- `secret` - (Optional) One or more `secret` blocks as defined below.
- `single_placement_group` - (Optional) Should this Virtual Machine Scale Set be limited to a Single Placement Group, which means the number of instances will be capped at 100 Virtual Machines. Defaults to `true`.
- `source_image_id` - (Optional) The ID of an Image which each Virtual Machine in this Scale Set should be based on.

NOTE: One of either `source_image_id` or `source_image_reference` must be set.

- `source_image_reference` - (Optional) A `source_image_reference` block as defined below.

NOTE: One of either `source_image_id` or `source_image_reference` must be set.

- `tags` - (Optional) A mapping of tags which should be assigned to this Virtual Machine Scale Set.
- `terraform_should_roll_instances_when_required` - (Optional) Should Terraform automatically roll instances within the Virtual Machine Scale Set when required? This happens when the `data_disk`, `os_disk`, `sku`, `source_image_id`, or `source_image_reference` fields change. This field defaults to `true`.

NOTE: This field is specific to Terraform, when required Terraform will automatically roll the instances in a Scale Set one at a time.

- `timezone` - (Optional) Specifies the time zone of the virtual machine, the possible values are defined here (<https://jackstromberg.com/2017/01/list-of-time-zones-consumed-by-azure/>).
- `upgrade_mode` - (Optional) Specifies how Upgrades (e.g. changing the Image/SKU) should be performed to Virtual Machine Instances. Possible values are `Automatic`, `Manual` and `Rolling`. Defaults to `Manual`.
- `winrm_listener` - (Optional) One or more `winrm_listener` blocks as defined below.
- `zone_balance` - (Optional) Should the Virtual Machines in this Scale Set be strictly evenly distributed across Availability Zones? Defaults to `false`. Changing this forces a new resource to be created.

NOTE: This can only be set to `true` when one or more `zones` are configured.

- `zones` - (Optional) A list of Availability Zones in which the Virtual Machines in this Scale Set should be created in. Changing this forces a new resource to be created.

A `additional_capabilities` block supports the following:

- `ultra_ssd_enabled` - (Optional) Should the capacity to enable Data Disks of the `UltraSSD_LRS` storage account type be supported on this Virtual Machine Scale Set? Defaults to `false`. Changing this forces a new resource to be created.

A `additional_unattend_config` block supports the following:

- `content` - (Required) The XML formatted content that is added to the `unattend.xml` file for the specified path and component.
- `setting` - (Required) The name of the setting to which the content applies. Possible values are `AutoLogon` and `FirstLogonCommands`.

A `automatic_os_upgrade_policy` block supports the following:

- `disable_automatic_rollback` - (Required) Should automatic rollbacks be disabled? Changing this forces a new resource to be created.
- `enable_automatic_os_upgrade` - (Required) Should OS Upgrades automatically be applied to Scale Set instances in a rolling fashion when a newer version of the OS Image becomes available? Changing this forces a new resource to be created.

A `boot_diagnostics` block supports the following:

- `storage_account_uri` - (Required) The Primary/Secondary Endpoint for the Azure Storage Account which should be used to store Boot Diagnostics, including Console Output and Screenshots from the Hypervisor.

A `certificate` block supports the following:

- `store` - (Required) The certificate store on the Virtual Machine where the certificate should be added.
- `url` - (Required) The Secret URL of a Key Vault Certificate.

NOTE: This can be sourced from the `secret_id` field within the `azurerm_key_vault_certificate` Resource.

A `data_disk` block supports the following:

- `caching` - (Required) The type of Caching which should be used for this Data Disk. Possible values are `None` , `ReadOnly` and `ReadWrite` .
- `disk_size_gb` - (Required) The size of the Data Disk which should be created.
- `lun` - (Required) The Logical Unit Number of the Data Disk, which must be unique within the Virtual Machine.
- `storage_account_type` - (Required) The Type of Storage Account which should back this Data Disk. Possible values include `Standard_LRS` , `StandardSSD_LRS` , `Premium_LRS` and `UltraSSD_LRS` .

NOTE: `UltraSSD_LRS` is only supported when `ultra_ssd_enabled` within the `additional_capabilities` block is enabled.

- `write_accelerator_enabled` - (Optional) Should Write Accelerator be enabled for this Data Disk? Defaults to `false` .

NOTE: This requires that the `storage_account_type` is set to `Premium_LRS` and that `caching` is set to `None` .

A `diff_disk_settings` block supports the following:

`option` - (Required) Specifies the Ephemeral Disk Settings for the OS Disk. At this time the only possible value is `Local` . Changing this forces a new resource to be created.

A `identity` block supports the following:

- `type` - (Required) The type of Managed Identity which should be assigned to the Windows Virtual Machine Scale Set. Possible values are `SystemAssigned` , `UserAssigned` and `SystemAssigned, UserAssigned` .
- `identity_ids` - (Optional) A list of User Managed Identity ID's which should be assigned to the Windows Virtual Machine Scale Set.

NOTE: This is required when `type` is set to `UserAssigned` .

A `ip_configuration` block supports the following:

- `name` - (Required) The Name which should be used for this IP Configuration.
- `application_gateway_backend_address_pool_ids` - (Optional) A list of Backend Address Pools ID's from a Application Gateway which this Virtual Machine Scale Set should be connected to.
- `application_security_group_ids` - (Optional) A list of Application Security Group ID's which this Virtual Machine Scale Set should be connected to.
- `load_balancer_backend_address_pool_ids` - (Optional) A list of Backend Address Pools ID's from a Load Balancer which this Virtual Machine Scale Set should be connected to.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `load_balancer_inbound_nat_rules_ids` - (Optional) A list of NAT Rule ID's from a Load Balancer which this Virtual Machine Scale Set should be connected to.

NOTE: When using this field you'll also need to configure a Rule for the Load Balancer, and use a `depends_on` between this resource and the Load Balancer Rule.

- `primary` - (Optional) Is this the Primary IP Configuration for this Network Interface? Defaults to `false`.

NOTE: One `ip_configuration` block must be marked as Primary for each Network Interface.

- `public_ip_address` - (Optional) A `public_ip_address` block as defined below.
- `subnet_id` - (Optional) The ID of the Subnet which this IP Configuration should be connected to.

`subnet_id` is required if `version` is set to `IPv4`.

- `version` - (Optional) The Internet Protocol Version which should be used for this IP Configuration. Possible values are `IPv4` and `IPv6`. Defaults to `IPv4`.

A `ip_tag` block supports the following:

- `tag` - The IP Tag associated with the Public IP, such as `SQL` or `Storage`.
- `type` - The Type of IP Tag, such as `FirstPartyUsage`.

A `network_interface` block supports the following:

- `name` - (Required) The Name which should be used for this Network Interface. Changing this forces a new resource to be created.
- `ip_configuration` - (Required) One or more `ip_configuration` blocks as defined above.
- `dns_servers` - (Optional) A list of IP Addresses of DNS Servers which should be assigned to the Network Interface.

- `enable_accelerated_networking` - (Optional) Does this Network Interface support Accelerated Networking? Defaults to `false`.
- `enable_ip_forwarding` - (Optional) Does this Network Interface support IP Forwarding? Defaults to `false`.
- `network_security_group_id` - (Optional) The ID of a Network Security Group which should be assigned to this Network Interface.
- `primary` - (Optional) Is this the Primary IP Configuration?

NOTE: If multiple `network_interface` blocks are specified, one must be set to `primary`.

A `os_disk` block supports the following:

- `caching` - (Required) The Type of Caching which should be used for the Internal OS Disk. Possible values are `None`, `ReadOnly` and `ReadWrite`.
- `storage_account_type` - (Required) The Type of Storage Account which should back this the Internal OS Disk. Possible values include `Standard_LRS`, `StandardSSD_LRS` and `Premium_LRS`.
- `diff_disk_settings` - (Optional) A `diff_disk_settings` block as defined above. Changing this forces a new resource to be created.
- `disk_size_gb` - (Optional) The Size of the Internal OS Disk in GB, if you wish to vary from the size used in the image this Virtual Machine Scale Set is sourced from.

NOTE: If specified this must be equal to or larger than the size of the Image the VM Scale Set is based on. When creating a larger disk than exists in the image you'll need to repartition the disk to use the remaining space.

- `write_accelerator_enabled` - (Optional) Should Write Accelerator be Enabled for this OS Disk? Defaults to `false`.

NOTE: This requires that the `storage_account_type` is set to `Premium_LRS` and that `caching` is set to `None`.

A `public_ip_address` block supports the following:

- `name` - (Required) The Name of the Public IP Address Configuration.
- `domain_name_label` - (Optional) The Prefix which should be used for the Domain Name Label for each Virtual Machine Instance. Azure concatenates the Domain Name Label and Virtual Machine Index to create a unique Domain Name Label for each Virtual Machine.
- `idle_timeout_in_minutes` - (Optional) The Idle Timeout in Minutes for the Public IP Address. Possible values are in the range 4 to 32.
- `ip_tag` - (Optional) One or more `ip_tag` blocks as defined above.
- `public_ip_prefix_id` - (Optional) The ID of the Public IP Address Prefix from where Public IP Addresses should be allocated. Changing this forces a new resource to be created.

NOTE: This functionality is in Preview and must be opted into via `az feature register --namespace Microsoft.Network --name AllowBringYourOwnPublicIpAddress` and then `az provider register -n Microsoft.Network`.

A `rolling_upgrade_policy` block supports the following:

- `max_batch_instance_percent` - (Required) The maximum percent of total virtual machine instances that will be upgraded simultaneously by the rolling upgrade in one batch. As this is a maximum, unhealthy instances in previous or future batches can cause the percentage of instances in a batch to decrease to ensure higher reliability. Changing this forces a new resource to be created.
- `max_unhealthy_instance_percent` - (Required) The maximum percentage of the total virtual machine instances in the scale set that can be simultaneously unhealthy, either as a result of being upgraded, or by being found in an unhealthy state by the virtual machine health checks before the rolling upgrade aborts. This constraint will be checked prior to starting any batch. Changing this forces a new resource to be created.
- `max_unhealthy_upgraded_instance_percent` - (Required) The maximum percentage of upgraded virtual machine instances that can be found to be in an unhealthy state. This check will happen after each batch is upgraded. If this percentage is ever exceeded, the rolling update aborts. Changing this forces a new resource to be created.
- `pause_time_between_batches` - (Required) The wait time between completing the update for all virtual machines in one batch and starting the next batch. The time duration should be specified in ISO 8601 format. Changing this forces a new resource to be created.

A `secret` block supports the following:

- `certificate` - (Required) One or more `certificate` blocks as defined above.
- `key_vault_id` - (Required) The ID of the Key Vault from which all Secrets should be sourced.

A `winrm_listener` block supports the following:

- `certificate_url` - (Optional) The Secret URL of a Key Vault Certificate, which must be specified when `protocol` is set to `Https`.

NOTE: This can be sourced from the `secret_id` field within the `azurerm_key_vault_certificate` Resource.

- `protocol` - (Required) The Protocol of the WinRM Listener. Possible values are `Http` and `Https`.

Attributes Reference

In addition to all arguments above, the following attributes are exported:

- `id` - The ID of the Windows Virtual Machine Scale Set.
- `identity` - An `identity` block as defined below.
- `unique_id` - The Unique ID for this Windows Virtual Machine Scale Set.

An `identity` block exports the following:

- `principal_id` - The ID of the System Managed Service Principal.

Timeouts

The `timeouts` block allows you to specify timeouts

(<https://www.terraform.io/docs/configuration/resources.html#timeouts>) for certain actions:

- `create` - (Defaults to 30 minutes) Used when creating the Windows Virtual Machine Scale Set.
- `update` - (Defaults to 60 minutes) Used when updating (and rolling the instances of) the Windows Virtual Machine Scale Set (e.g. when changing SKU).
- `delete` - (Defaults to 30 minutes) Used when deleting the Windows Virtual Machine Scale Set.

Import

Windows Virtual Machine Scale Sets can be imported using the `resource_id`, e.g.

```
terraform import azurerm_windows_virtual_machine_scale_set.test /subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/mygroup1/Microsoft.Compute/virtualMachineScaleSets/scaleset1
```