Analyzing Machine Data with Datadog

Rajesh Kumar
www.DevOpsSchool.com

Basic Windows Administration

Basic Linux Administration

Windows Environment

Hadoop Sandbox

www.DevOpsSchool.com

Introducing Datadog

www.DevOpsSchool.com

Common Mistakes

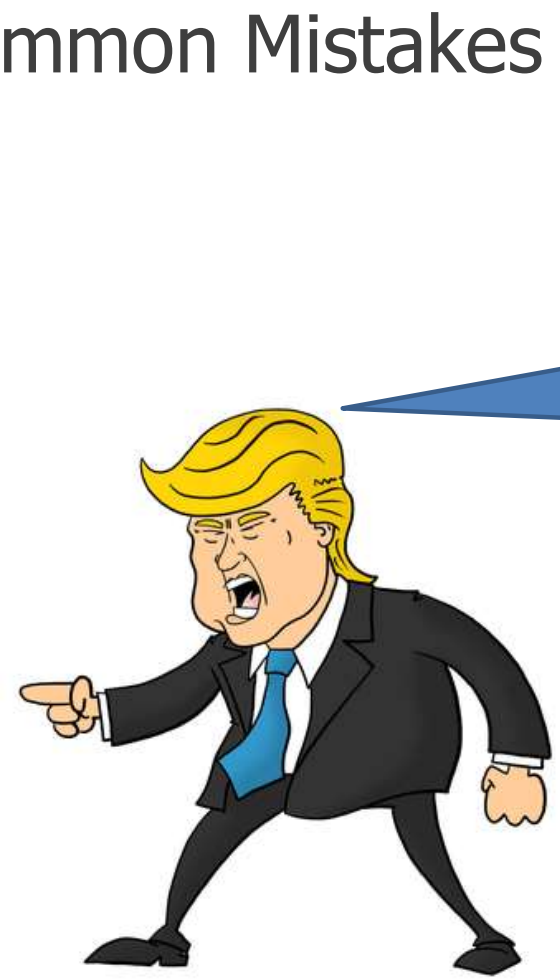

"My web host already offer online statistics. This is more than enough information."

Big mistake. Huge mistake!!!

Comparable to “well there’s usually enough money in the bank, so who needs to plan and budget?”

“I DON’T HAVE TIME”

EVEN BIGGER MISTAKE!!!


SYSTEM FAILURE
www.DevOpsSchool.com


www.DevOpsSchool.com


HELP


I'VE MADE A
HUGE MISTAKE
www.DevOpsSchool.com


WHAT HAVE YOU GOTTEN US INTO?
www.DevOpsSchool.com

Log

Log = record related to whatever activities occurring on an information system

Also: alert, "event", alarm, message, record, etc

Log Data Sources

- IDS
- Firewalls/IPS
- Anti-maleware
- Proxies
- Network Infrastructure
- Servers
- Databases
- Application

Machine Data

Data generated by machines, computer processing, applications and sensor data.

Machine data is everywhere.
In fact you are generating it right now!

Server & Workstation Logs

- Linux/Windows
- Log files
- Access
- File system

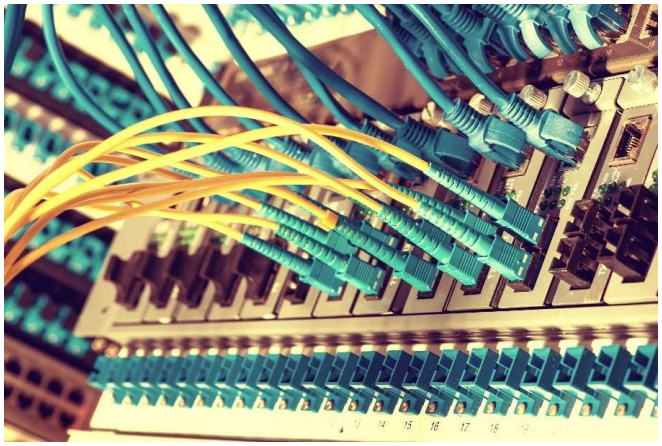

Networks

- Firewall
- Warnings
- Alerts
- IP addresses

Database

- Audit logs
- Configurations
- Schemas
- Tables
- Queries

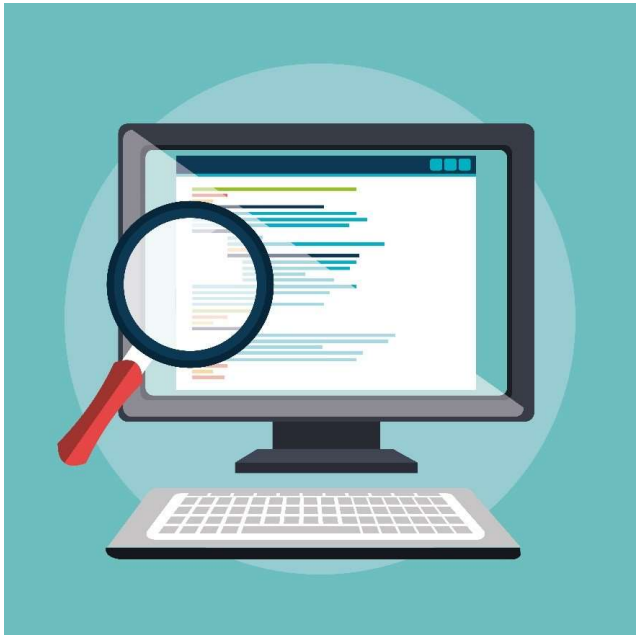

Web

- Transactions
- Click-stream
- Location
- Browser
- Time

DevOps

- Test logs
- log4j alerts
- Event logs
- Code check-in

IOT

- GPS
- RFID
- Biometric
- Temperature
- Limitless

Analyzing Log Data: Why It's Important

- Production Monitoring and Debugging
- Resource Usage
- HTTP Errors
- Slow Queries
- Rogue Automated Robots
- Security

Analyzing Log Data: Why It's Important

- Tracking Your Site's/Platform's Visitors
- Situational awareness and new threat discovery
- Getting more value out of network and security infrastructures
- Extracting what is really actionable automatically
- Measuring security (metrics, trends, etc)
- Compliance an regulations
- Incident response (last, but not least)

Best Practices For Log Analysis

Pattern detection and recognition: to filter messages based on a pattern book.
Normalization: to convert different log elements such as dates to the same format.
Tagging and classification: to tag log elements with keywords and categorize them into a number of classes so you can filter and adjust the way you display your data.
Correlation analysis: to collate logs from different sources and systems and sort meaningful messages that pertain to a particular event.
Artificial ignorance: a machine learning process to identify and "ignore" log entries that are not useful and detect anomalies.

What to look for in Logs?

- Password changes
- Unauthorized logins
- Login failures
- New login events
- Malware detection
- Malware attacks seen by IDS or other evidence
- Scans on your firewalls open and closed ports
- Denial of service attacks
- Errors on network devices
- File name changes
- File integrity changes
- Data exported
- New processes started or running processes stopped
- Shared access events
- Disconnected events
- New service installation
- File auditing
- New user accounts
- Modified registry values

Common Fields in Log

- ✓Time
- ✓Source
- ✓Destination
- ✓Protocol
- ✓Port(s)
- ✓User name
- ✓Event/Attach type
- ✓Bytes exchanged

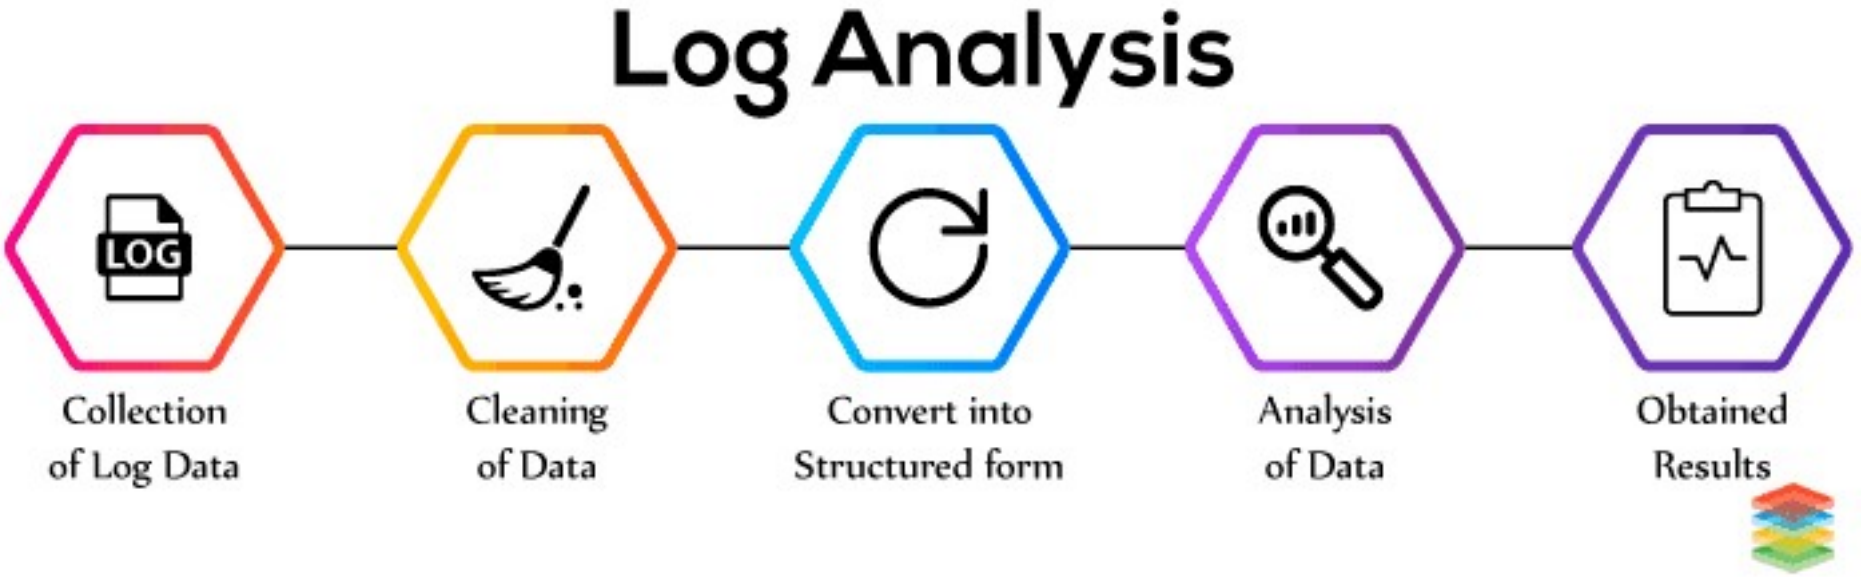
Log Analysis
LOG
Collection of Log Data
Cleaning of Data
Convert into Structured form
Analysis of Data
Obtained Results

Log Analysis Process

- Generate
- Collect
- Aggregate
- Normalize
- Alert
- Store
- Summarize, Baseline
- Make conclusions
- Act on them


logz.io


graylog


DATADOG


splunk>


logstash


sumologic


DATADOG

Analyzing Log Files

Ingest

Store

Visualize

Analyzing Log Files

Ingest Store Visualize

Who Is Datadog For?

DevOps

Data Analyst

Anyone

Security is one of the fastest growing sectors in IT

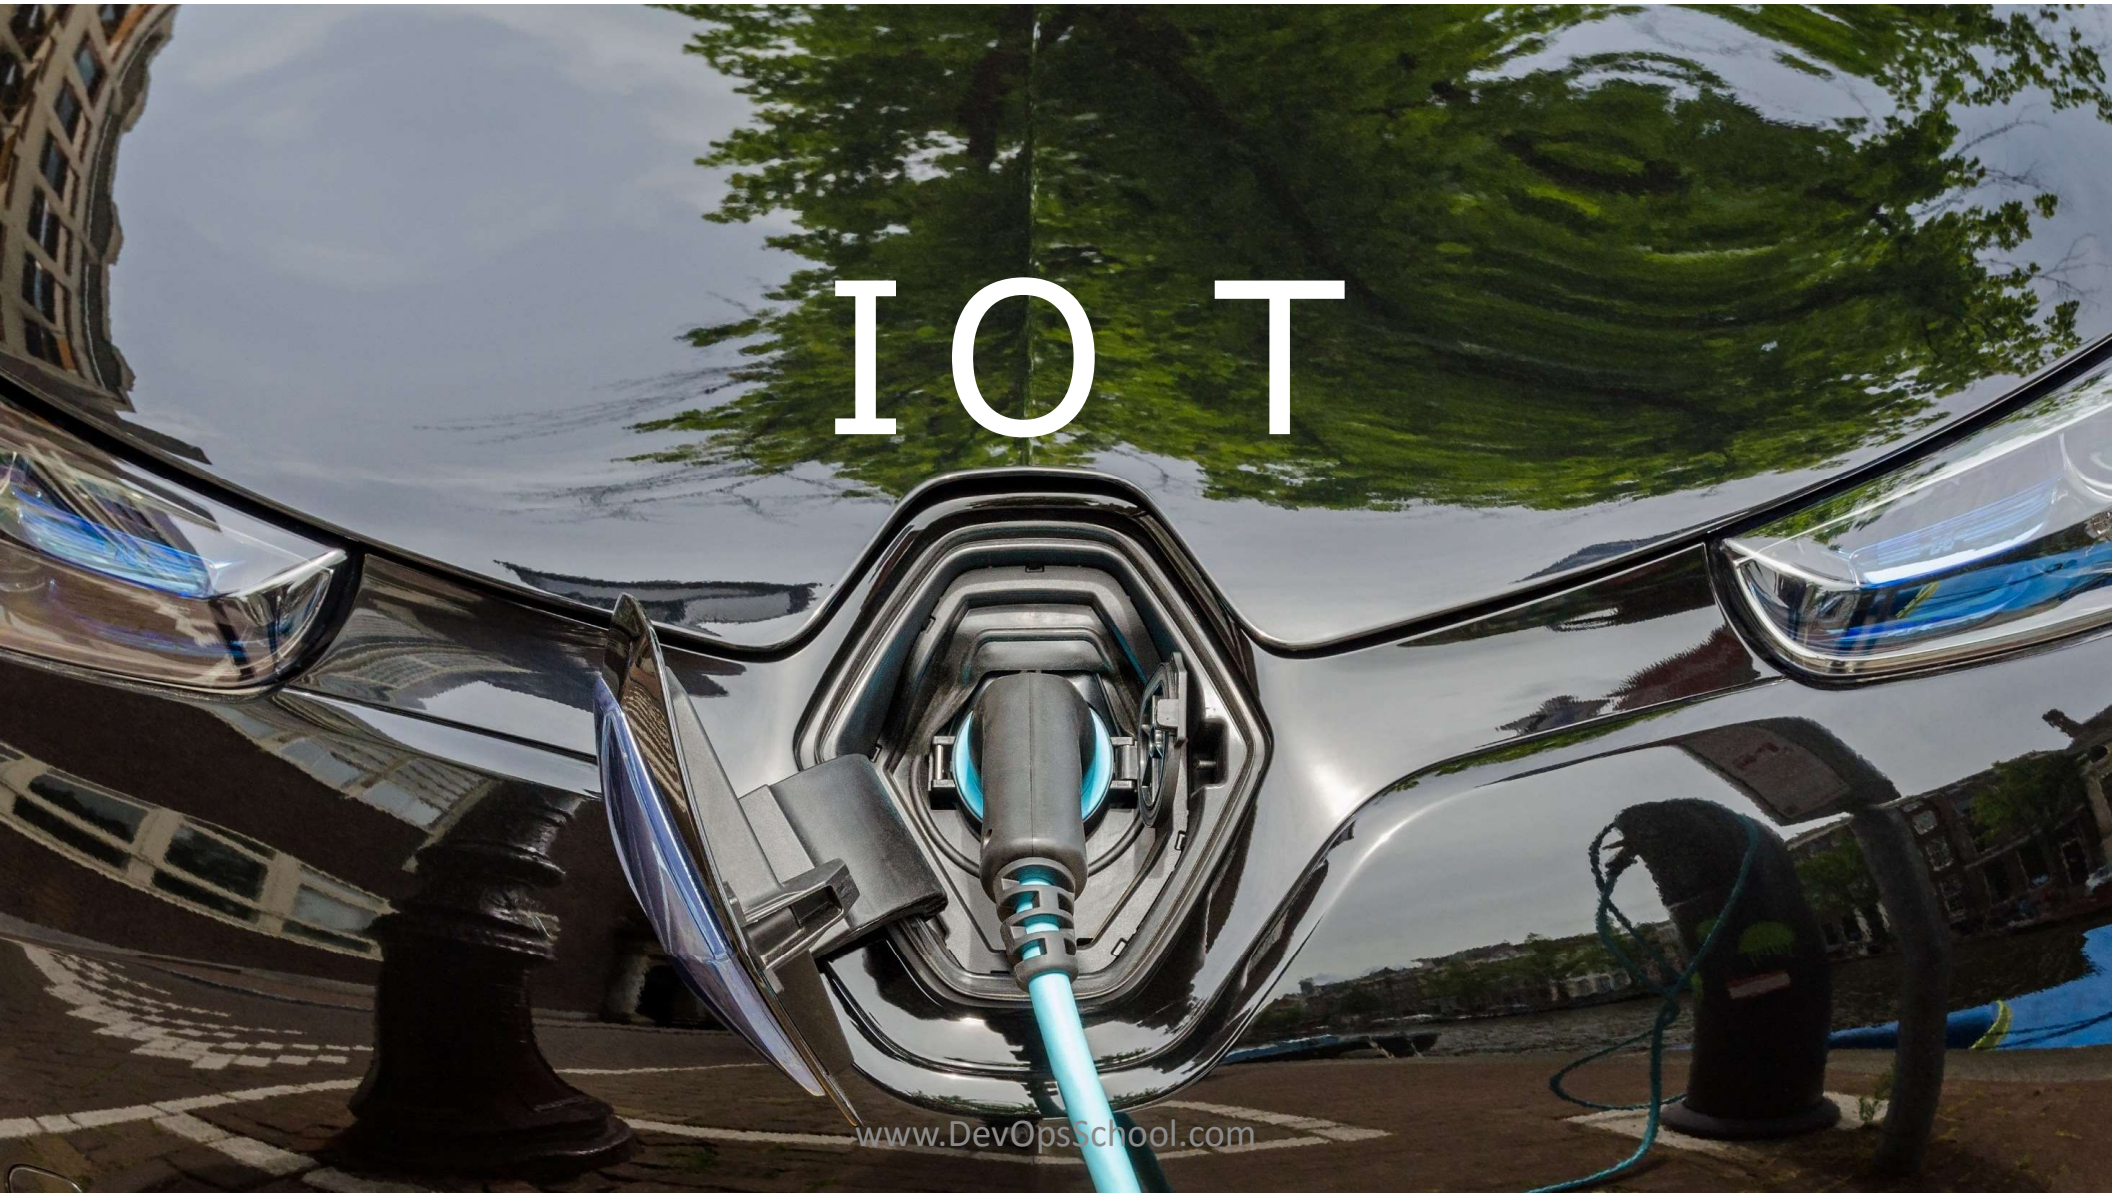
I O T
www.DevOpsSchool.com

What Data?

Index **ANY** data from **ANY** source

- Computers
- Network devices
- Virtual machines
- Internet devices
- Communication devices
- Sensors
- Databases

- Logs
- Configurations
- Messages
- Call detail records
- Clickstream
- Alerts
- Metrics
- Scripts
- Changes
- Tickets

Get Started!

Where are your logs coming from ?

Server
With the Datadog Agent

Container
Docker, Kubernetes...

Cloud
AWS, GCP, Azure...

Client
JavaScript, Mobile...

Other
Rsyslog, FluentD...

1 Select a log source

Select one of your existing integrations or a new one. Refer to the full documentation for additional configuration and frameworks. The list of supported OSes for those integrations can be found here.

APACHE | cassandra | Consul | elasticsearch | HAPROXY | NGINX | mongoDB | Java | Journald
APACHE Tomcat | GO | Microsoft .NET | Ruby | node | PostgreSQL | VARNISH CACHE | python | MySQL
redis | Microsoft IIS | kafka | APACHE ZooKeeper | RabbitMQ | php | | Memcached | MARATHON
gunicorn | Custom Files | Other ...

Full documentation | Logs FAQ | Contact Datadog Support

Where are your logs coming from ?

Server
With the Datadog Agent

Container
Docker, Kubernetes...

Cloud
AWS, GCP, Azure...

Client
JavaScript, Mobile...

Other
Rsyslog, FluentD...

1 Select a log source

Select one of your existing integrations, or a new one. Refer to the full Documentation for additional configuration and frameworks.

docker | kubernetes | Amazon ECS | Amazon EKS | MESOS | AWS Fargate

Datadog Log Management doesn't yet support the following integrations. Register below for updates and we'll let you know when something is ready.

CoreOS ☐ Request

RED HAT OPENSHIFT ☐ Request

Other integrations: [ ] Submit

Full documentation | Logs FAQ | Contact Datadog Support

Where are your logs coming from ?

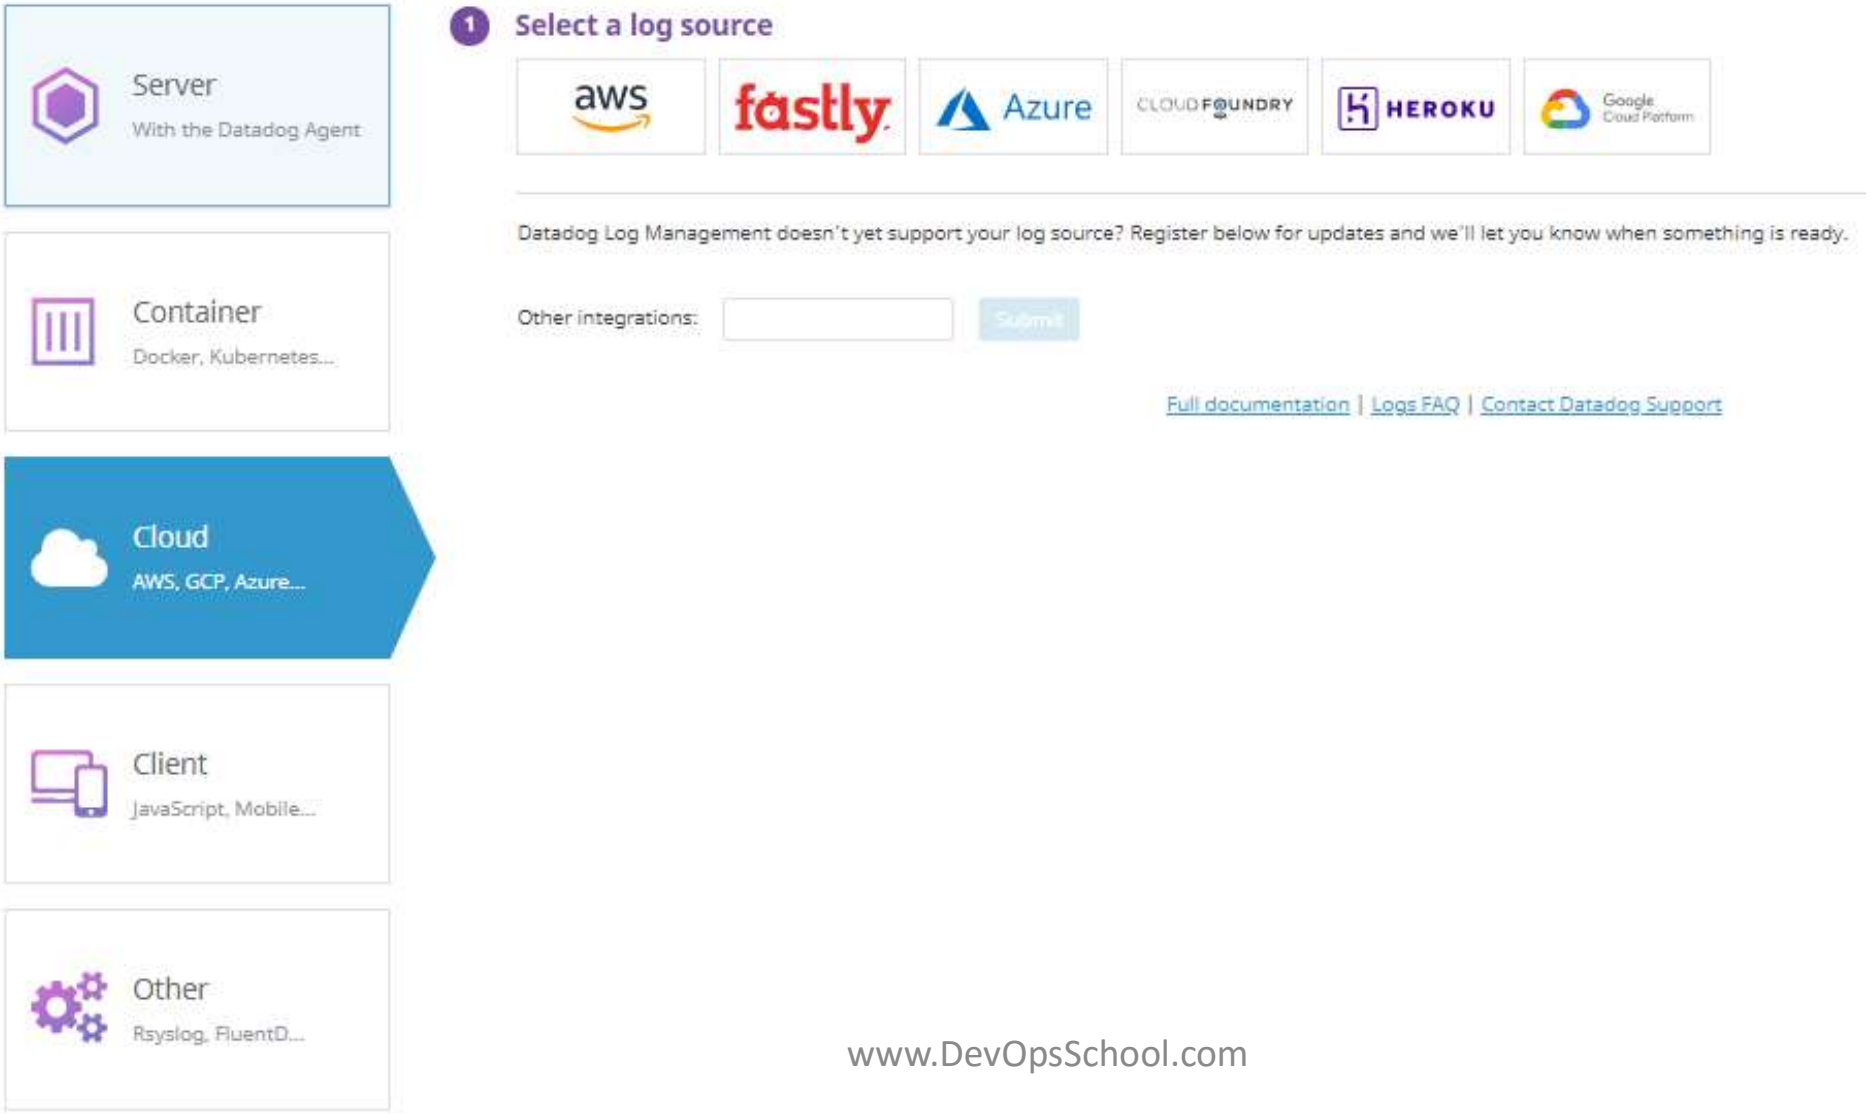

Where are your logs coming from ?

Server
With the Datadog Agent

Container
Docker, Kubernetes...

Cloud
AWS, GCP, Azure...

Client
JavaScript, Mobile...

Other
Rsyslog, FluentD...

1 Select a log source

Select one of your exisiting integrations or a new one. Refer to the full documentation for additional configuration and frameworks.

JS android

Datadog Log Management doesn't yet support the following integrations. Register below for updates and we'll let you know when something is ready.

iOS

Request

Other integrations: Submit

Full documentation | Logs FAQ | Contact Datadog Support

Where are your logs coming from ?

Server
With the Datadog Agent

Container
Docker, Kubernetes...

Cloud
AWS, GCP, Azure...

Client
JavaScript, Mobile...

Other
Rsyslog, FluentD...

1 Select a log source

rsyslog fluentd logstash syslog-ng NXLog fluentbit

Datadog Log Management doesn't yet support your log source? Register below for updates and we'll let you know when something is ready.

Other integrations: Submit

Full documentation | Logs FAQ | Contact Datadog Support

Datadog Log Workflow

Log Collection & Integrations: Ingest all your logs from your hosts, containers, and cloud providers.
Processing: Process and enrich all of your logs with pipelines and processors.
Live Tail: See your ingested logs in real time across all your environments.
Generate Metrics: Generate Metrics from Ingested Logs.
Archives: Archive all enriched logs into S3 buckets.
Index: Dynamically decide what to include or exclude from your indexes to control your costs.

After indexing your logs, explore them in the Log Explorer:

Log Explorer: Discover the Log Explorer view, how to add Facets and Measures.
Search: Search through all of your indexed logs.
Analytics: Perform Log Analytics over your indexed logs.
Patterns: Spot Log Patterns by clustering your indexed logs together.
Saved Views: Use Saved Views to automatically configure your Log Explorer.

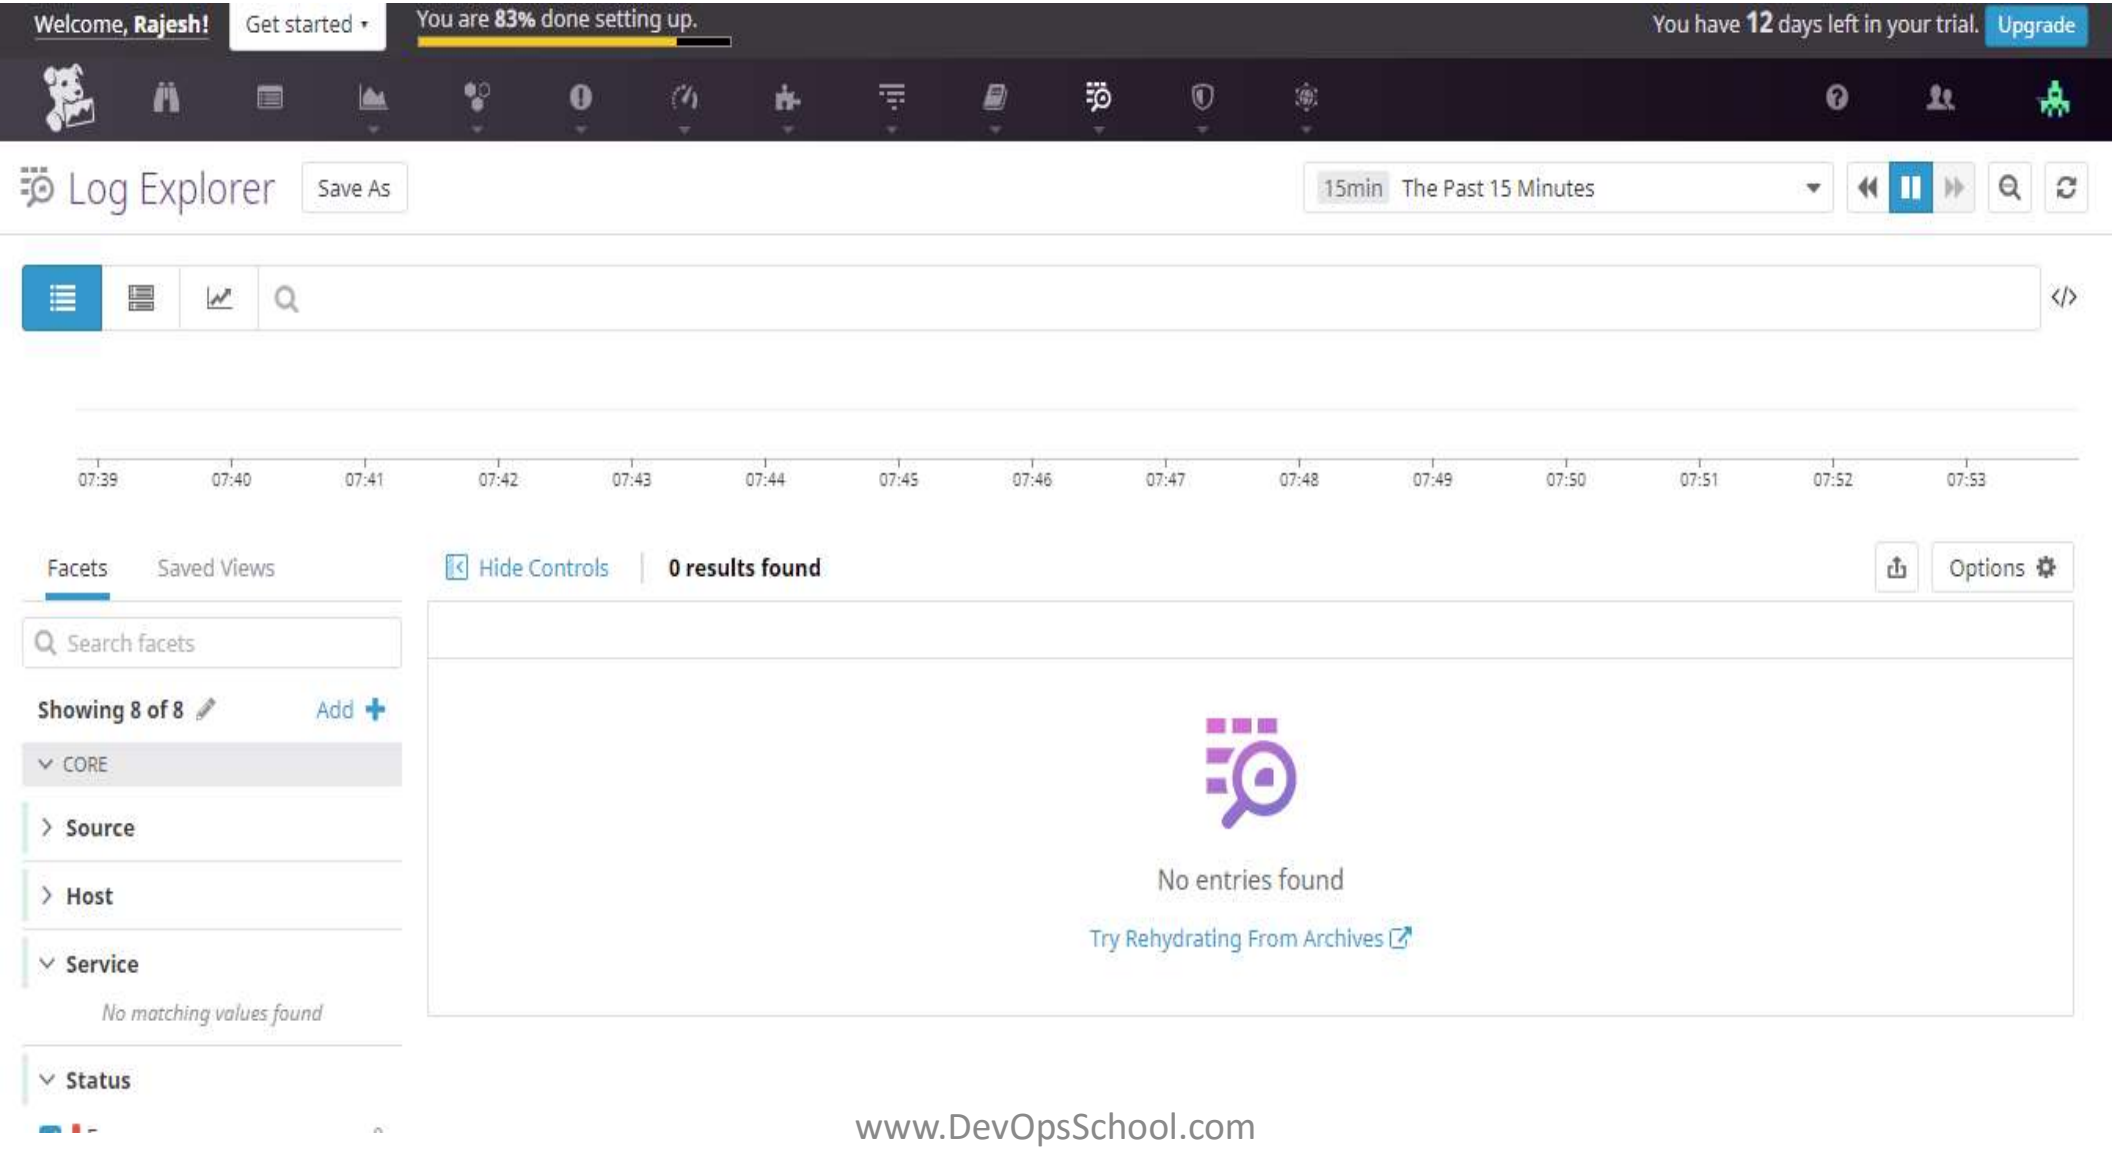

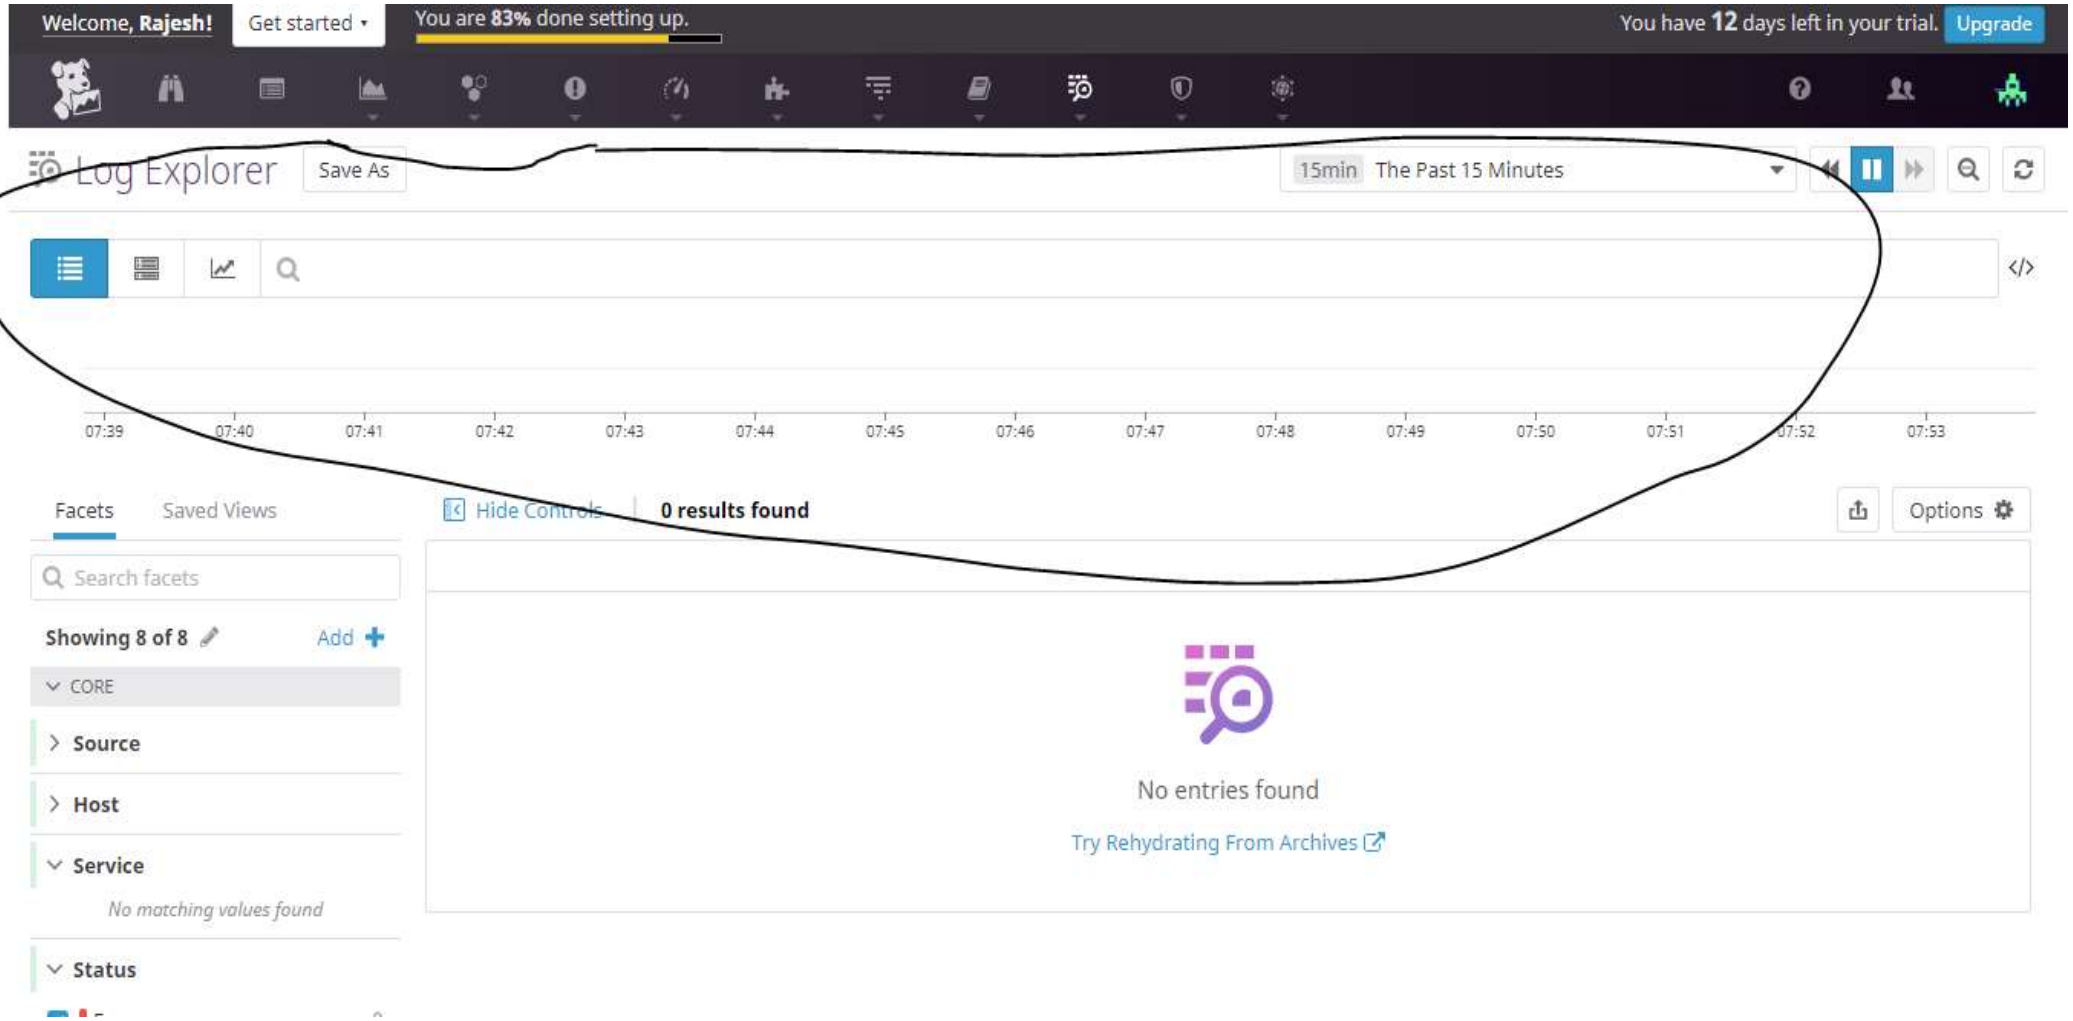
Welcome, Rajesh!
Get started
You are 83% done setting up.
You have 12 days left in your trial.
Upgrade
Log Explorer
Save As
15min
The Past 15 Minutes
07:39
07:40
07:41
07:42
07:43
07:44
07:45
07:46
07:47
07:48
07:49
07:50
07:51
07:52
07:53
Facets
Saved Views
Hide Controls
0 results found
Options
Search facets
Showing 8 of 8
Add
CORE
Source
Host
Service
No matching values found
Status
No entries found
Try Rehydrating From Archives