

Redis at DATADOG

Ilan Rabinovitch
Director, Community
Datadog

RedisConf 2016
San Francisco, CA

\$ finger ilan@datadog

[datadoghq.com]

Name: Ilan Rabinovitch

Role: Director, Technical Community

Interests:

- * Open Source
- * Web Operations & Infra Automation
- * Monitoring and Metrics
- * FL/OSS Community Events

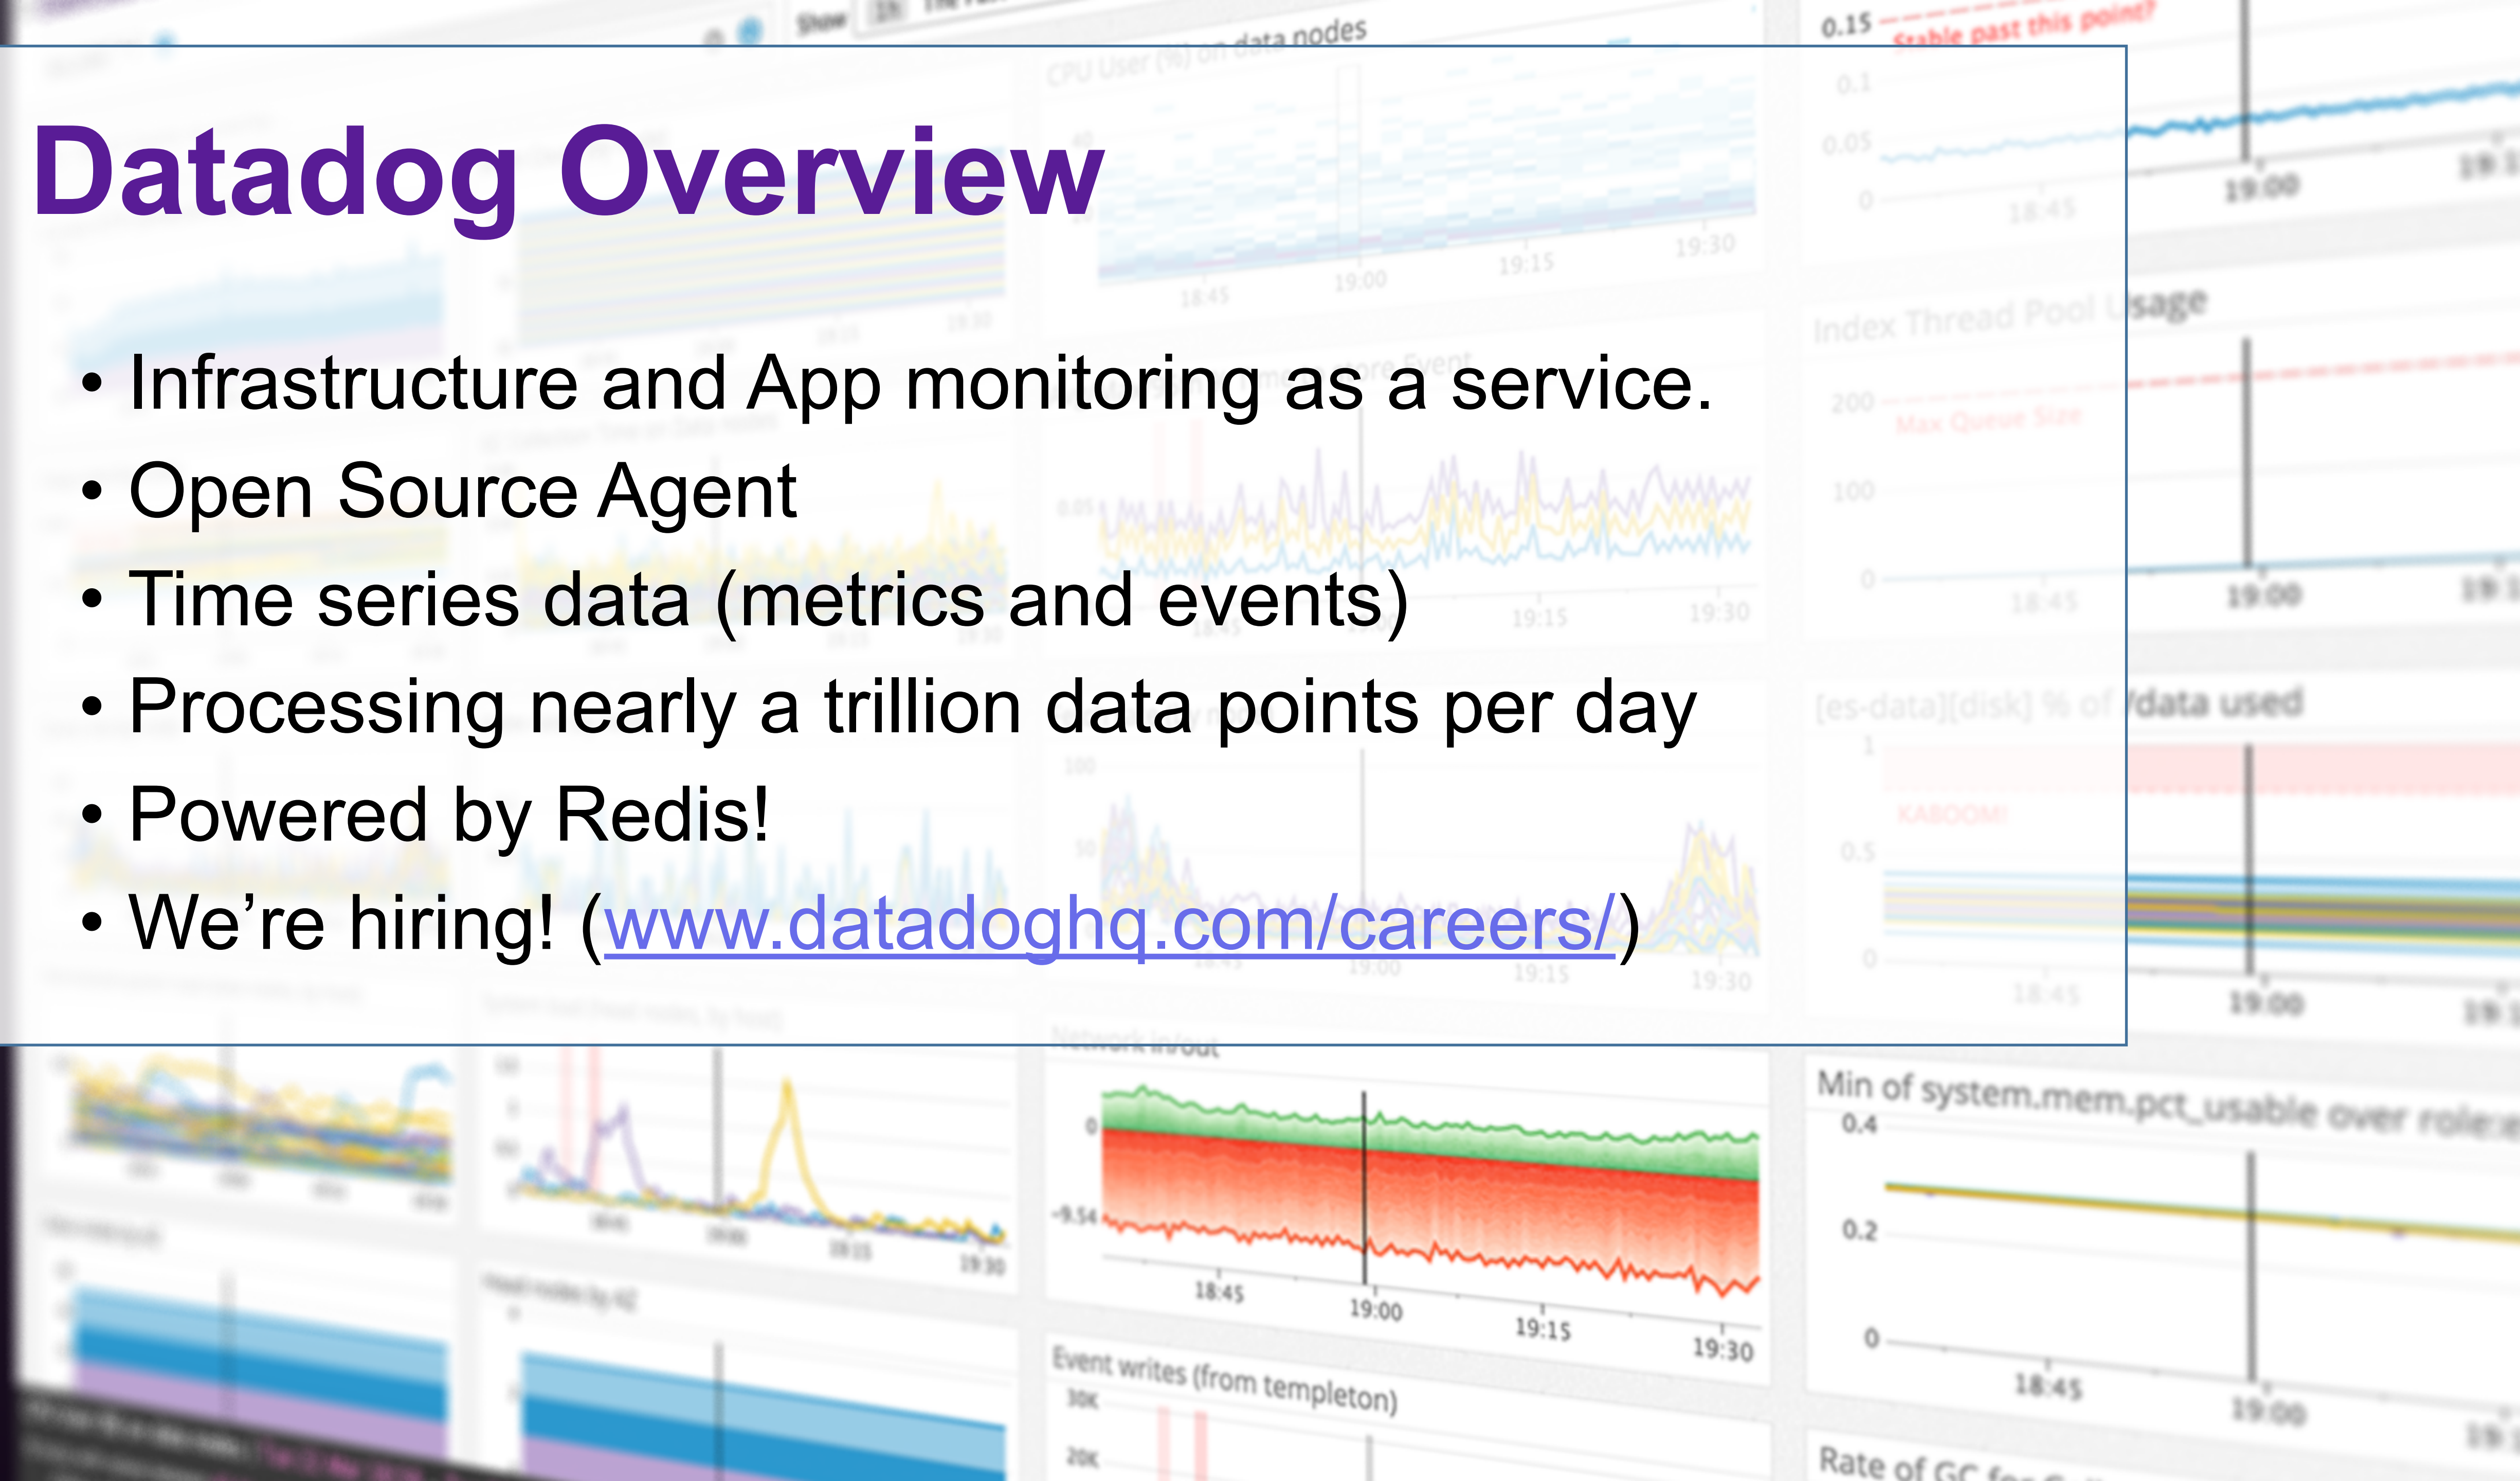


```
$ cat ~/.plan
```

- 1. An Overview of Datadog*
- 2. Monitoring 101*
- 3. How Datadog uses Redis*
- 4. Key Metrics and Examples*

Datadog Overview

- Infrastructure and App monitoring as a service.
- Open Source Agent
- Time series data (metrics and events)
- Processing nearly a trillion data points per day
- Powered by Redis!
- We're hiring! (www.datadoghq.com/careers/)



Monitor Everything



AWS



Docker



CoreOS



Chef



Puppet



Github



Pagerduty



Nagios



Go



Postgres



Java



VMware



Redis



MySQL



Apache



Tomcat



MongoDB




New Relic

Operating Systems, Cloud Providers (AWS), Containers, Web Servers, Datastores, Caches, Queues and more...



Top Technologies Running on Docker



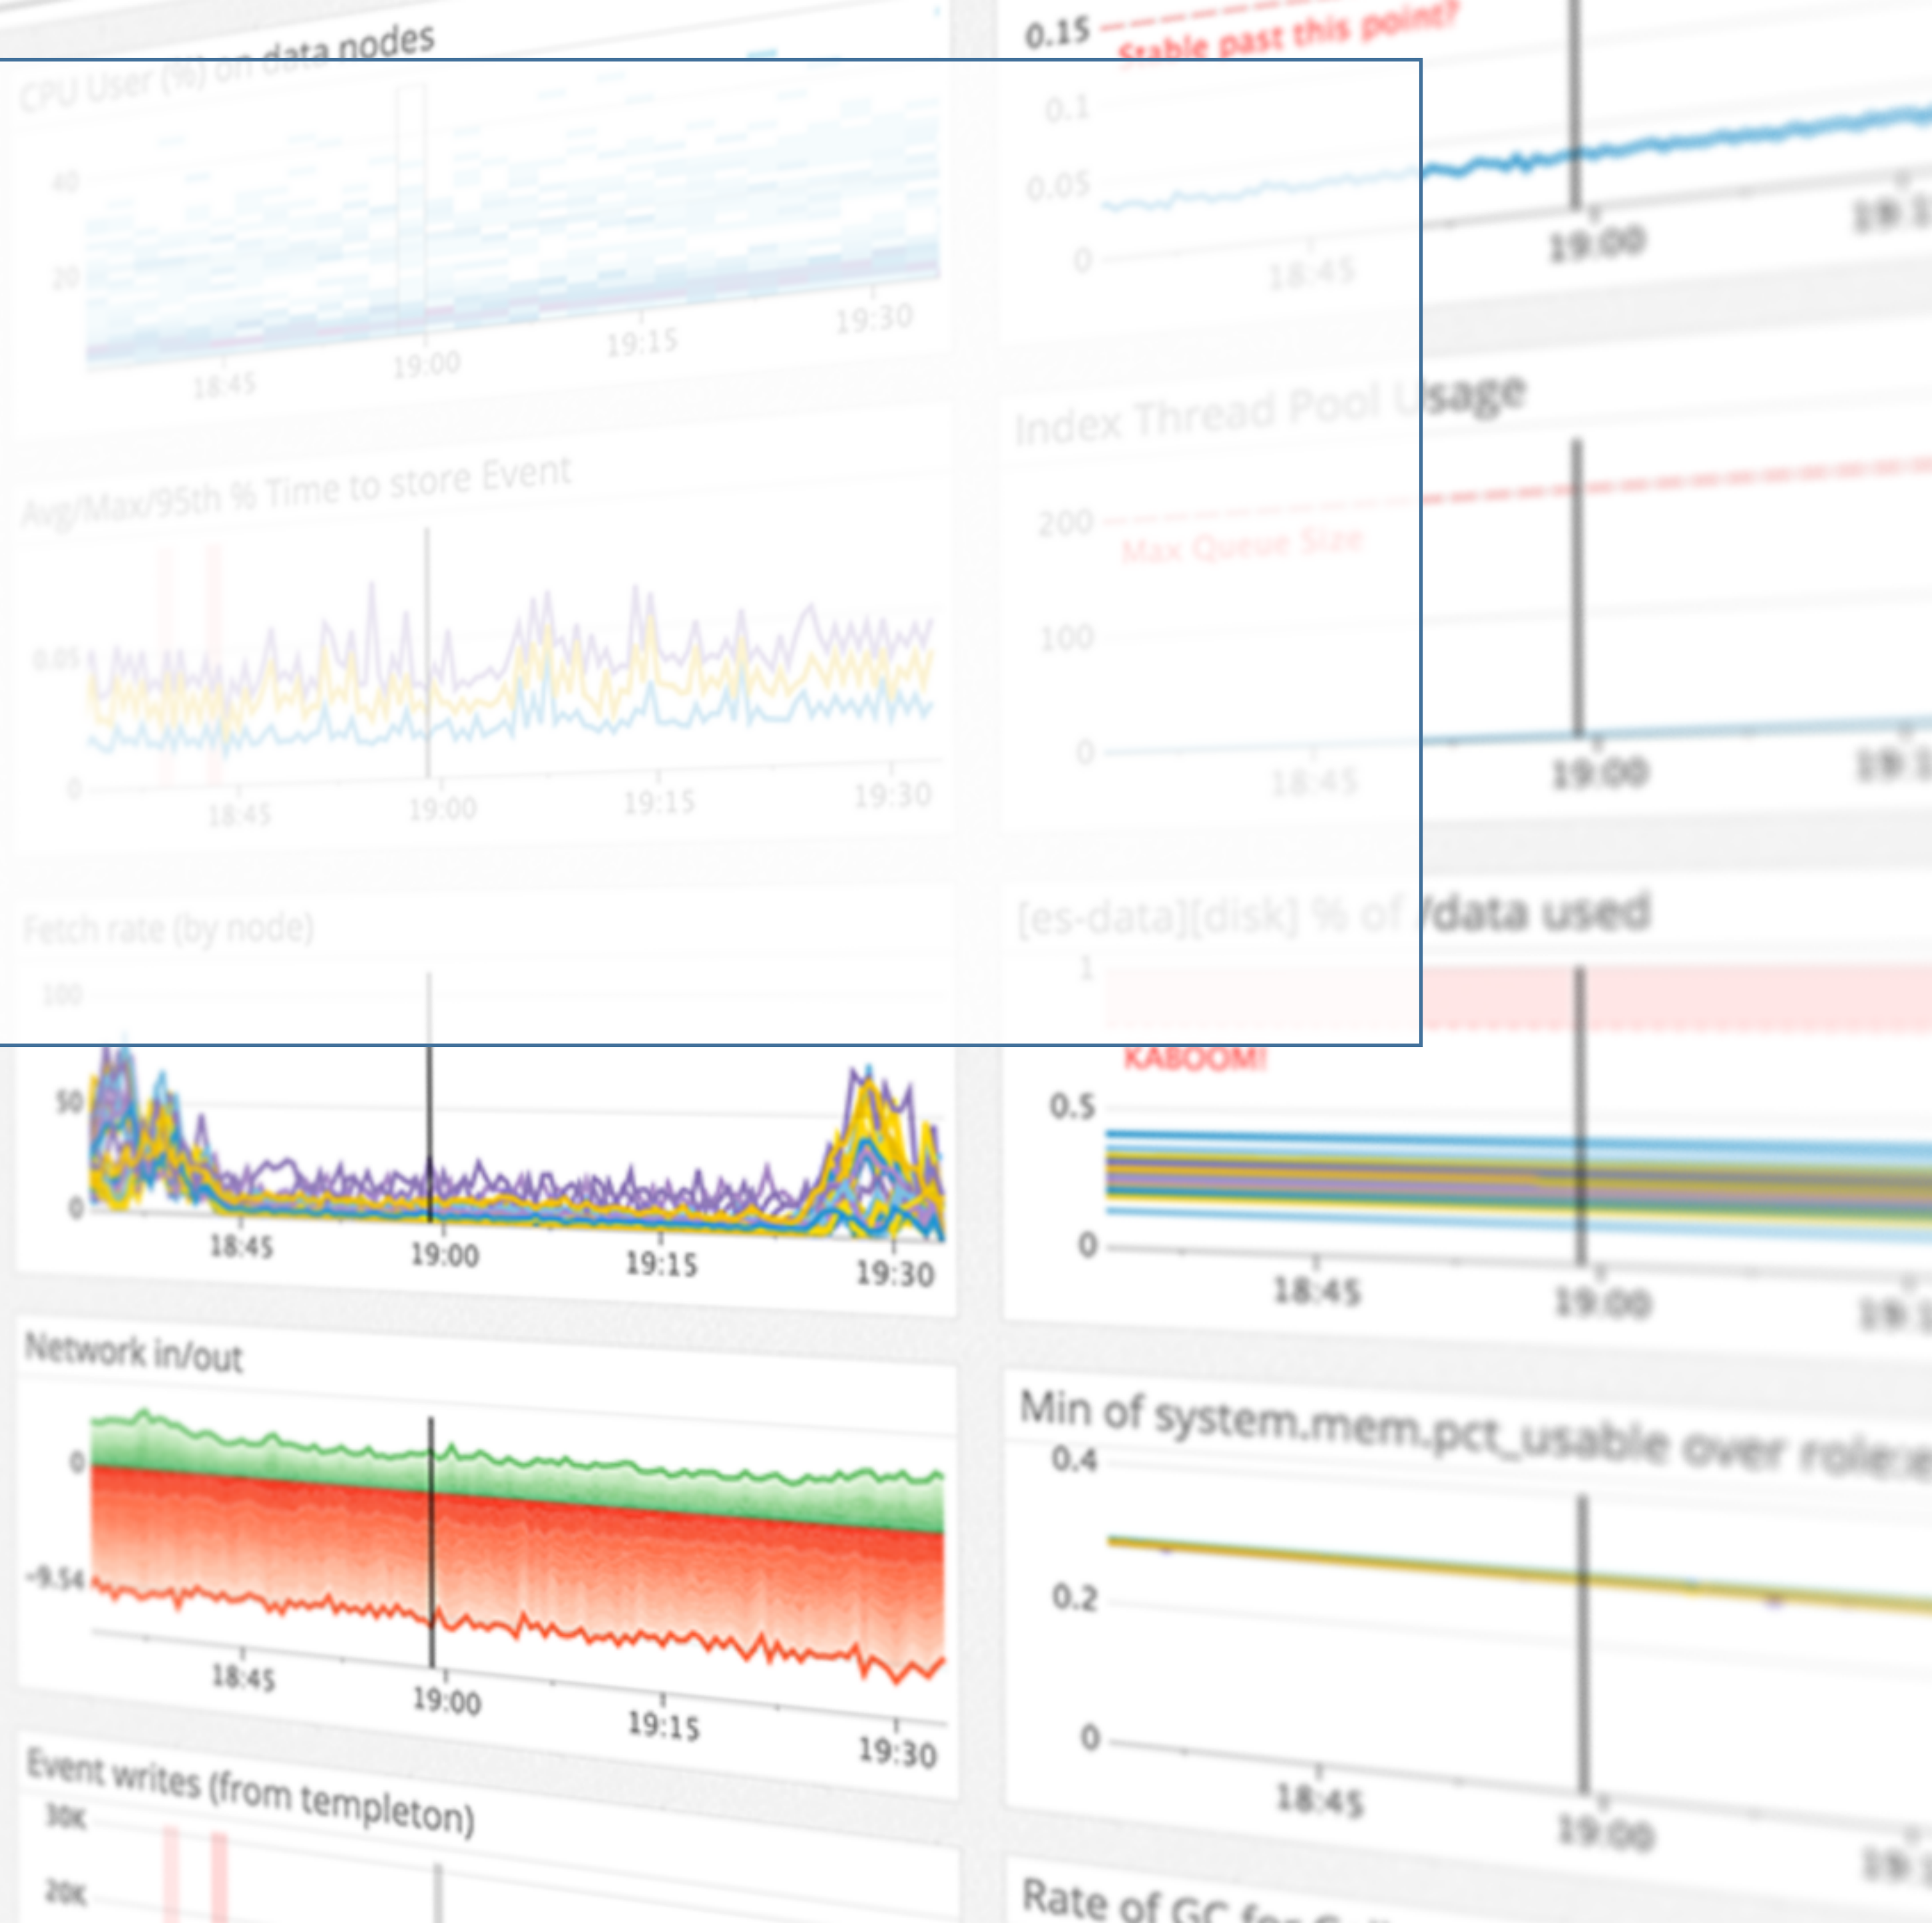
Source: Datadog

Redis Use Cases


Queueing

Caching

Session State



Monitoring 101





Honest Status Page @honest_update · 18 Sep 2015

We have no idea what's wrong so we're just gonna undo whatever we did last and whoever did it is the incident manager.

← 66 57 ...



Honest Status Page @honest_update · 20 Sep 2015

If there ain't an alert for it, it ain't broke!

← 33 28 ...




Honest Status Page @honest_update · 10 Mar 2015

Two words: unbounded queues.


← 61 75 ...

Follow @honest_update on Twitter



Pagerduty Incident Trends

Week by week



259 alerts
in the week of Nov 17

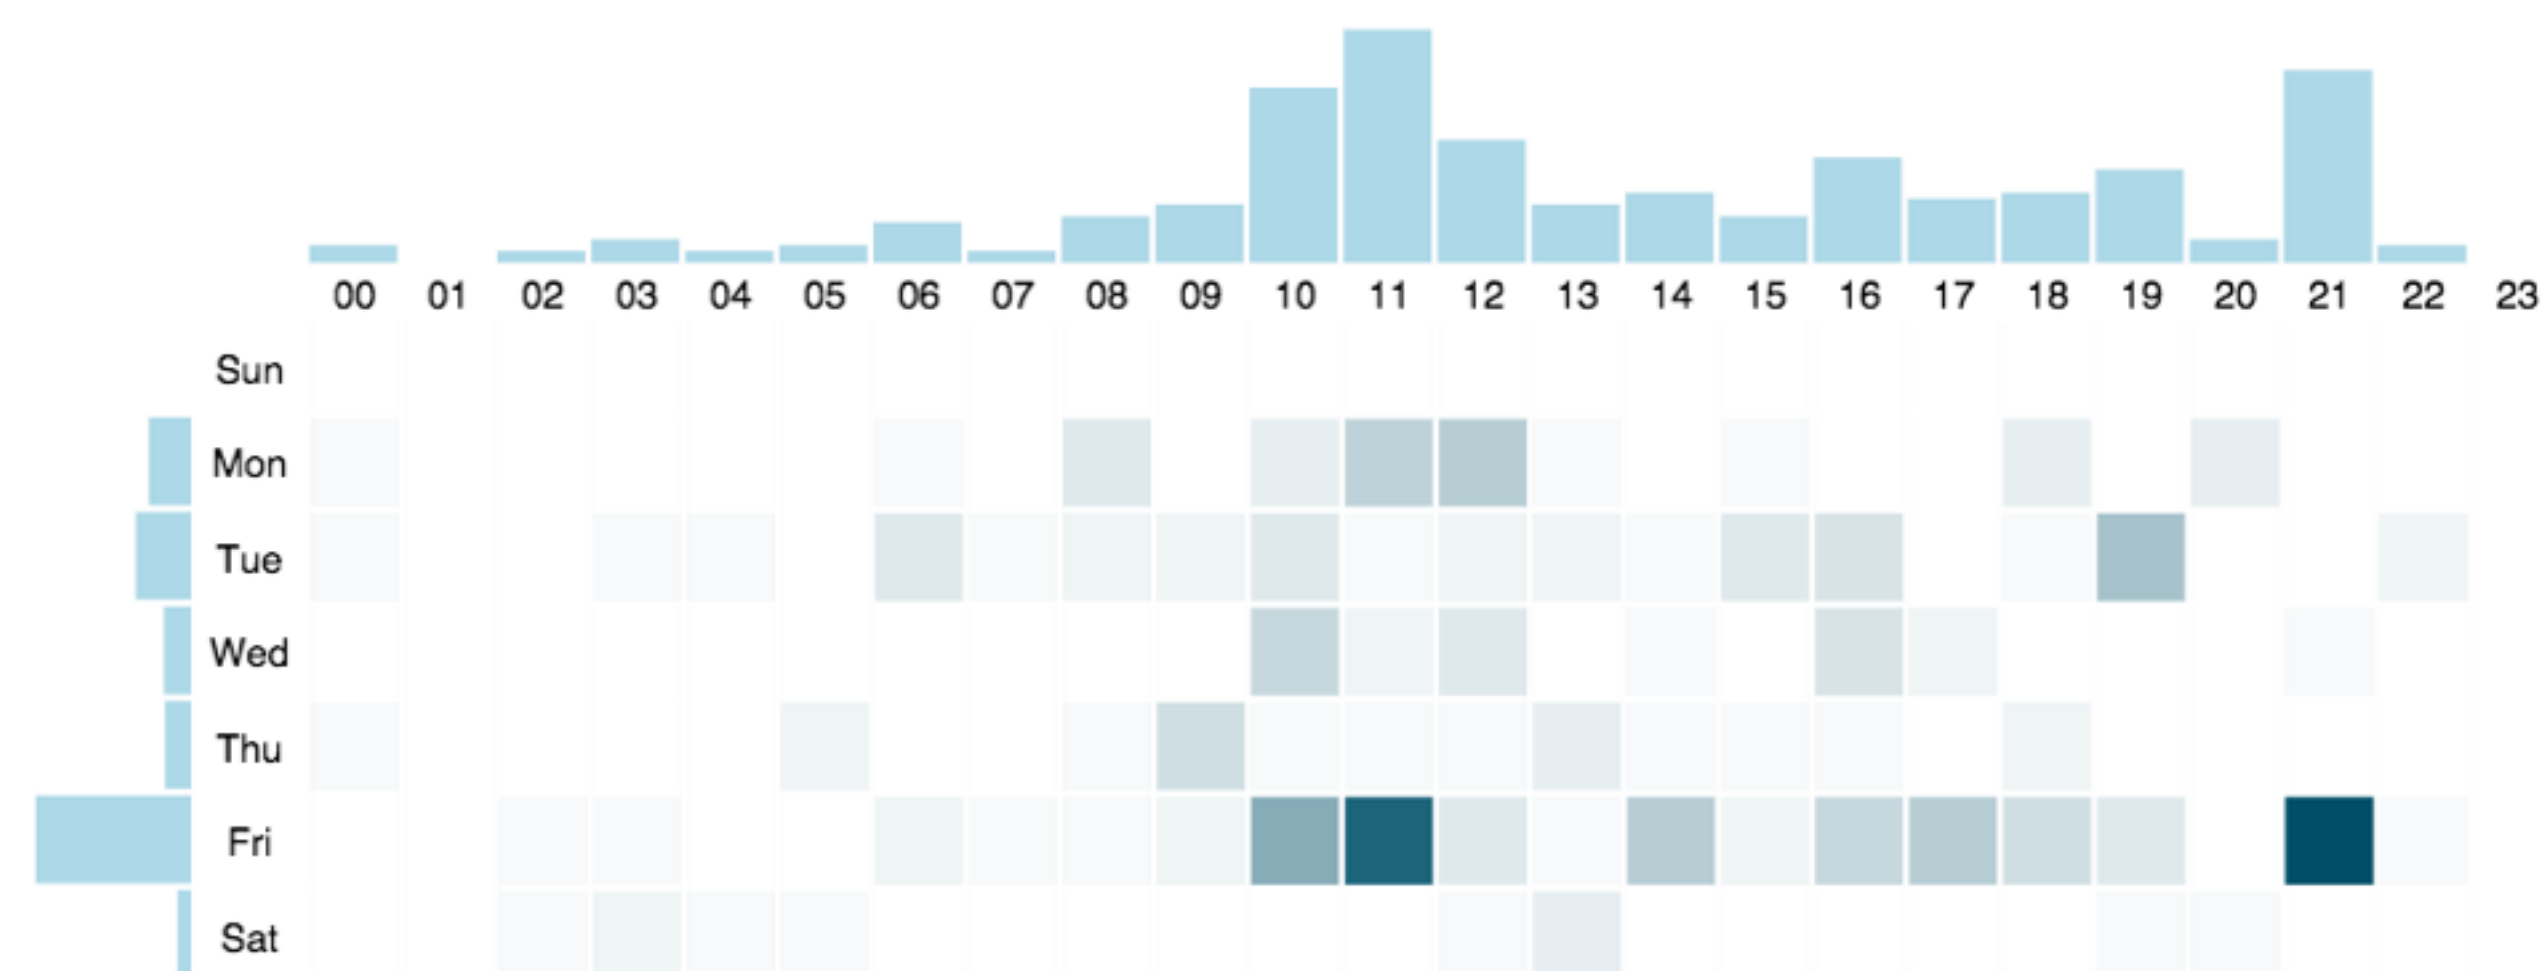
-65% from the previous week

Hadoop has alerted more than any other service for the last 3 weeks.

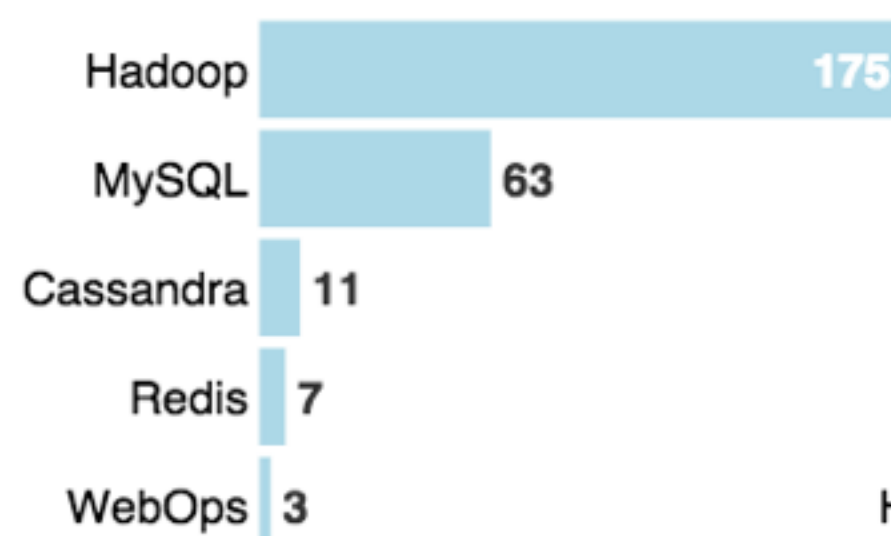
MySQL has alerted for the last 25 consecutive weeks.

MySQL alerted 63 times, **+1650%** from the previous week.


By day and hour



Top alerting services



Top alerting incidents



**Collecting data is cheap;
not having it when you
need it can be expensive**

Instrument all the things!

Operational Complexity Increases with..

- Number of things to measure
- Velocity of change

How much we measure?

1 instance


- 10 metrics from CloudWatch

1 operating system (e.g., Linux)

- 100 metrics

50~ metrics per redis instance

HOSTS RUN
FOUR
CONTAINERS
AT A TIME



Source: Datadog

Operational Complexity

100

instances



400

containers

Operational Complexity: Scale

160

metrics per host



640

metrics per host

Assuming 4 containers per host

Operational Complexity: Scale

100

instances



64,000

metrics

Assuming 4 containers per host

How much we measure?

1 instance

- 10 metrics from application

1 operating system (e.g. Linux)

- 100 metrics

50~100 metrics per application

N containers

- $150 * N$ metrics

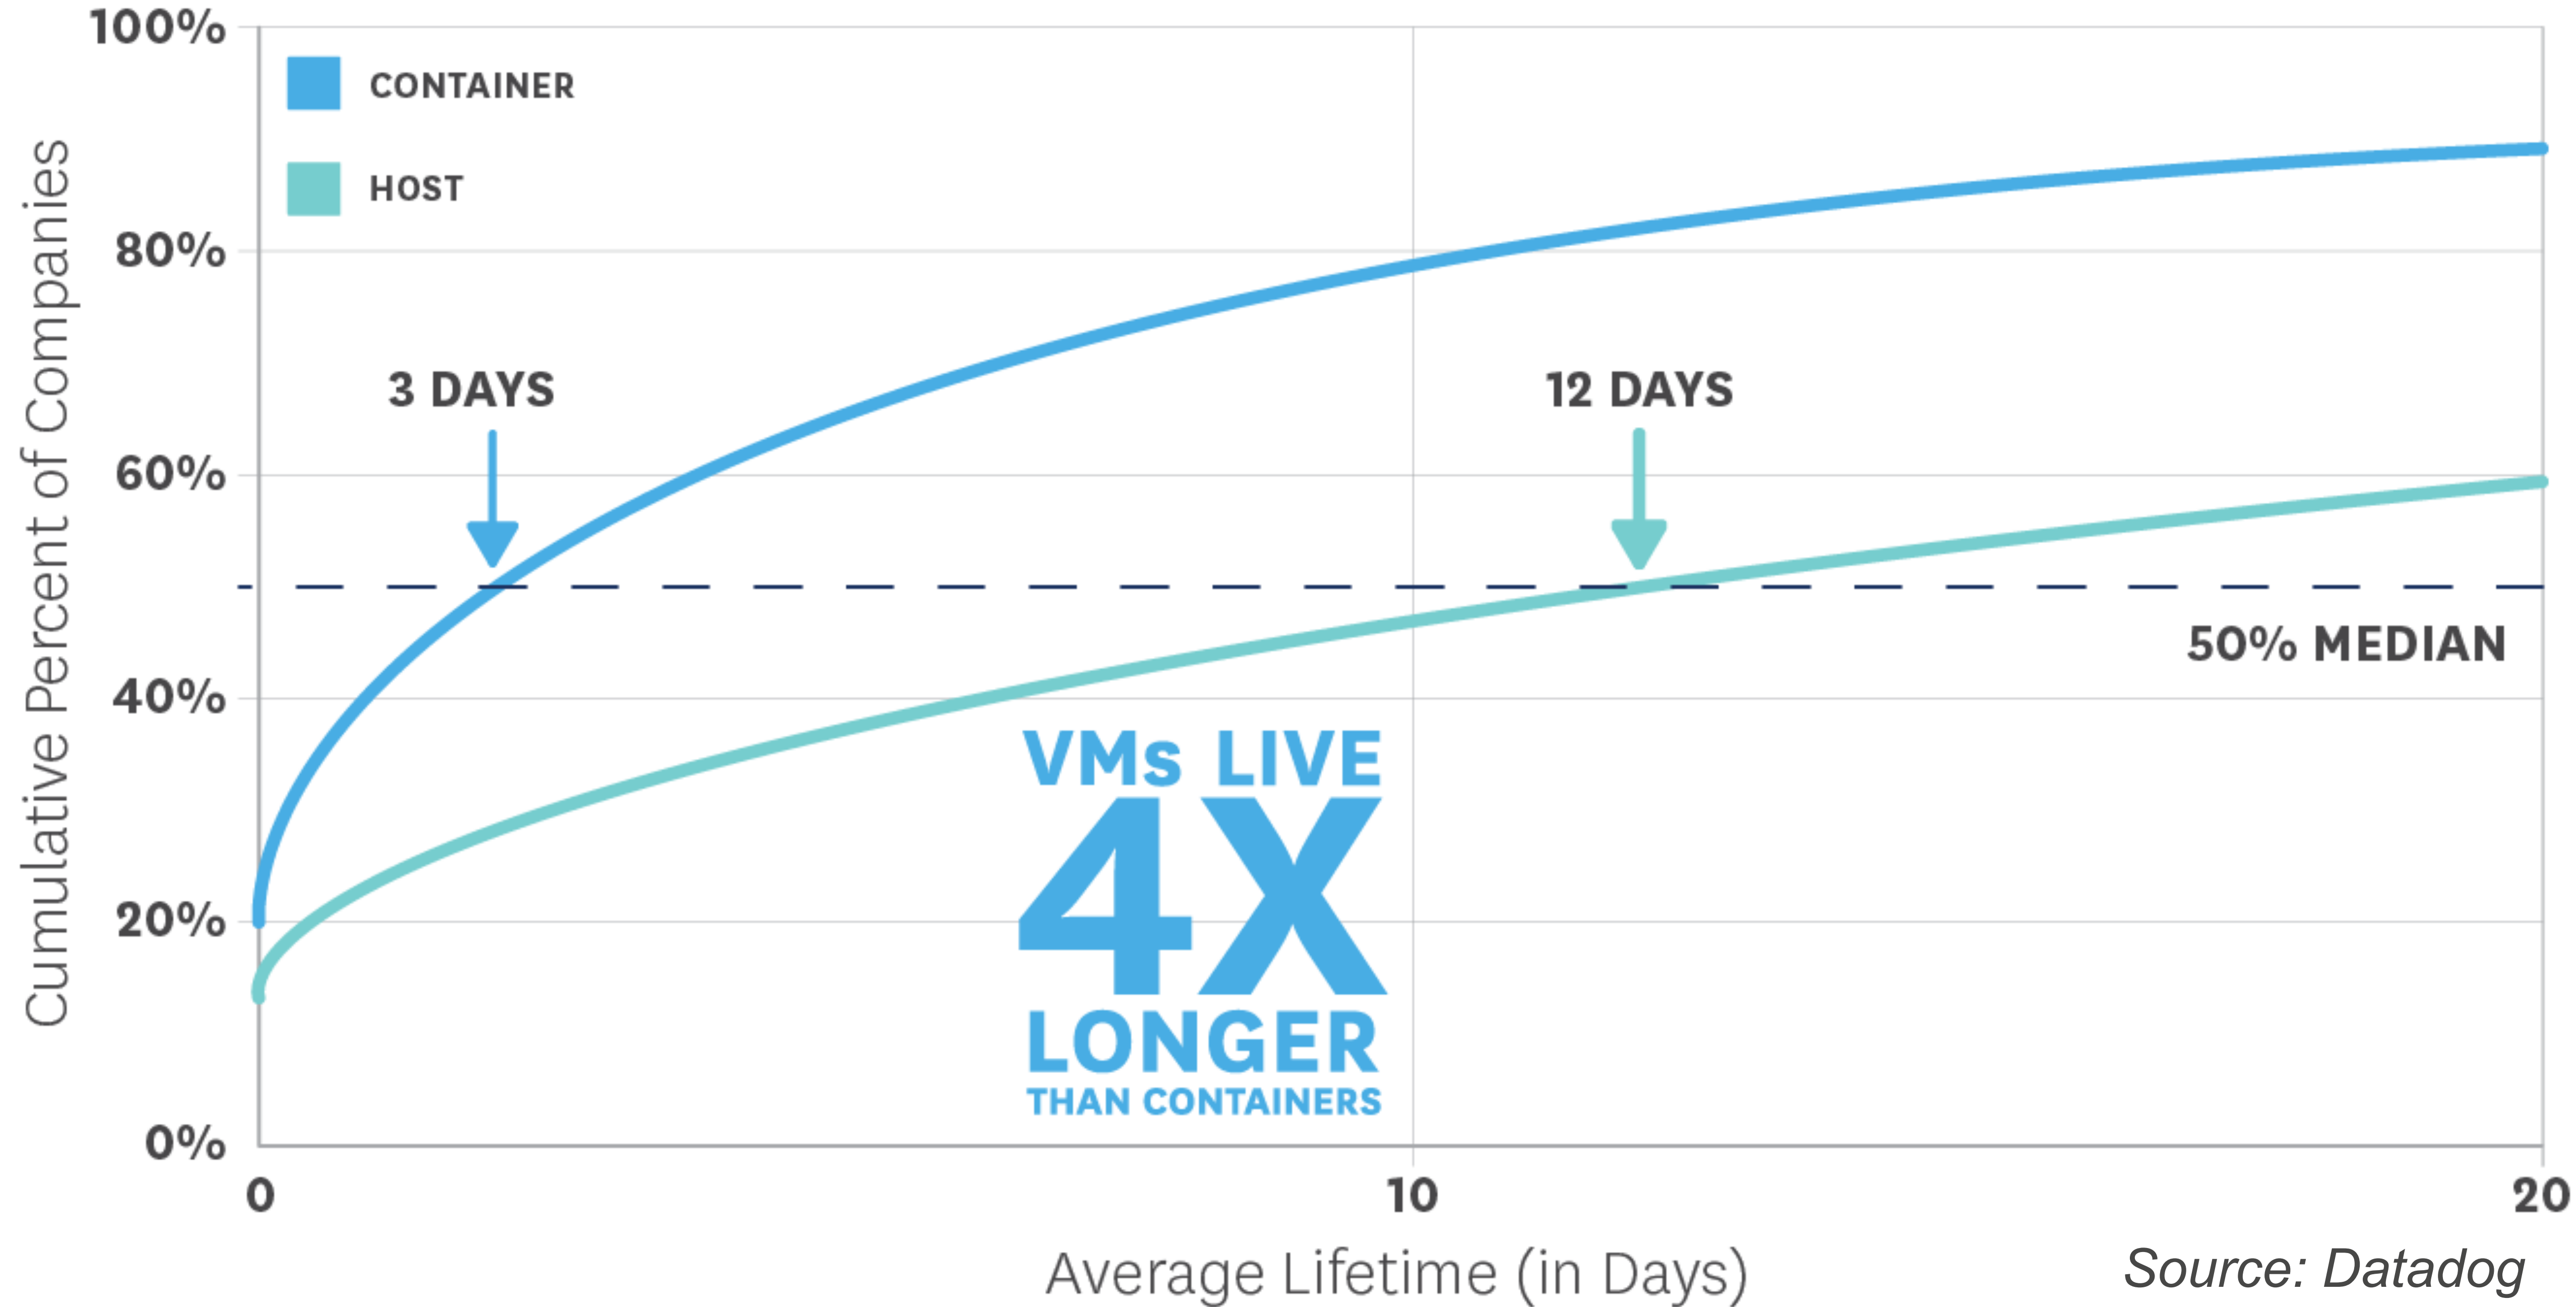
Metrics Overload!

Operational Complexity Increases with..

- Number of things to measure
- Velocity of change



Average Lifetimes of Hosts and Containers



Source: Datadog

Host half-life

hours,
days,
months




Container half-life


minutes,
hours,
days

Operational Complexity Increases with..

- Number of things to measure
- Velocity of change









Monitoring 101: Alerting on what matters

series / theory / alerting / monitoring / monitoring-101

More Details at: <http://www.datadoghq.com/blog/monitoring-101-alerting/>

Finding Signal - Categorizing Your Metrics

WORK METRICS

RESOURCE METRICS

EVENTS



THE COMMONWEALTH TRUST (INDIA) LTD

COMMONWEALTH TRUST

WORK METRICS

THROUGHPUT

SUCCESS

ERROR

PERFORMANCE



RESOURCE METRICS

UTILIZATION

SATURATION

ERROR

AVAILABILITY





Adrian Cole

@adrianfcoble



Following

Q: Are we losing money?

A: Can't answer that, but I can tell you what average CPU usage was 5ish mins ago..

EVENTS

CODE CHANGES

ALERTS

SCALING EVENTS

ETC



Examples: Web Application

Work Metrics:

- **Requests Per Second**
- **Dropped Connections**
- **Request Response Time**
- **Error Rates (4xx or 5xx)**
- **Success (2xx)**

Resource Metrics:

- **Disk I/O**
- **Network Bandwidth**
- **Memory**
- **CPU**
- **Queue Length**

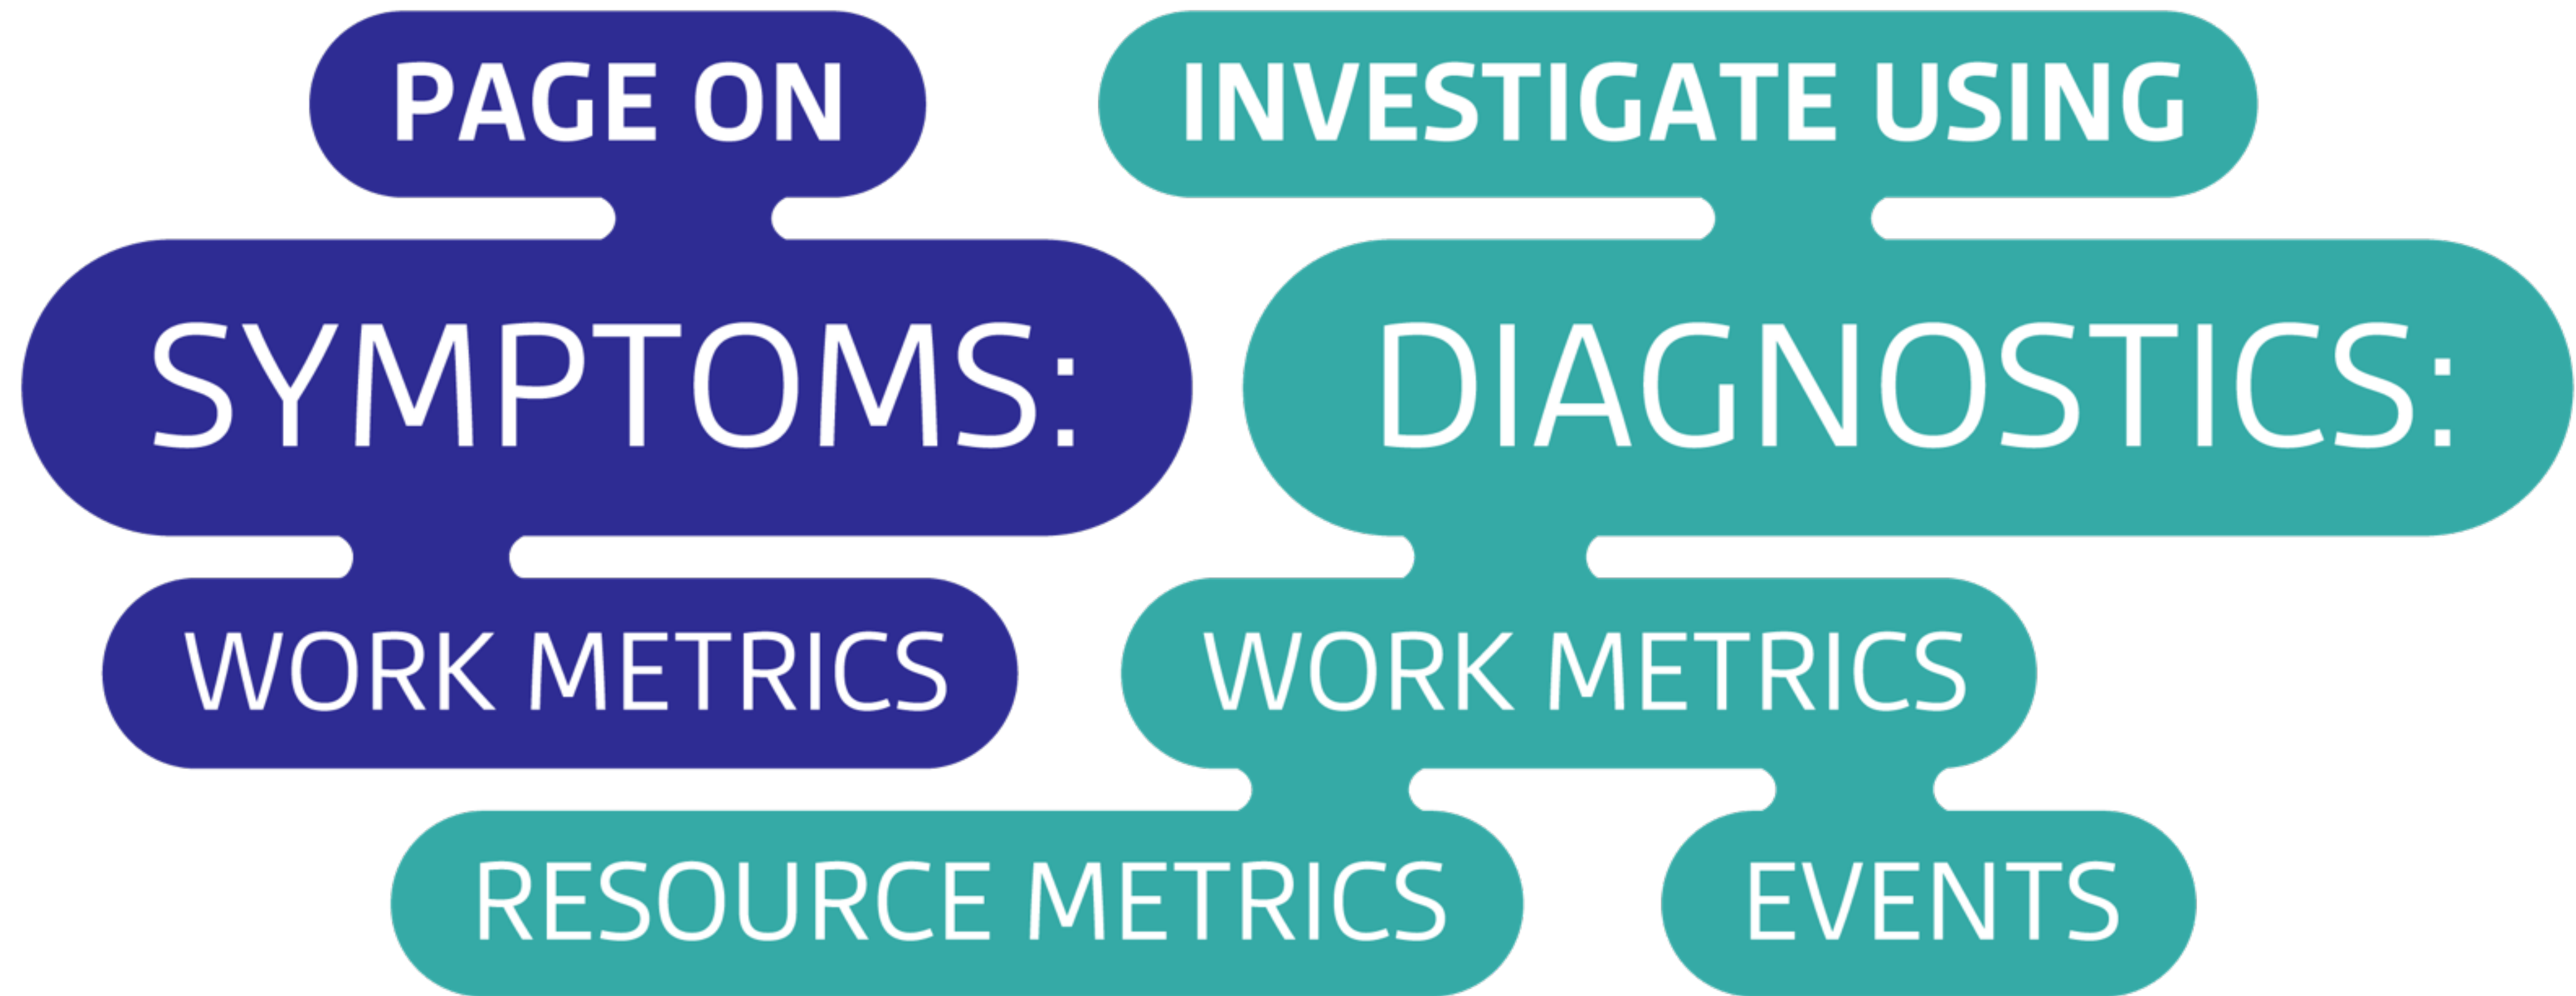
Examples: Web Application - Events

Work Metrics:

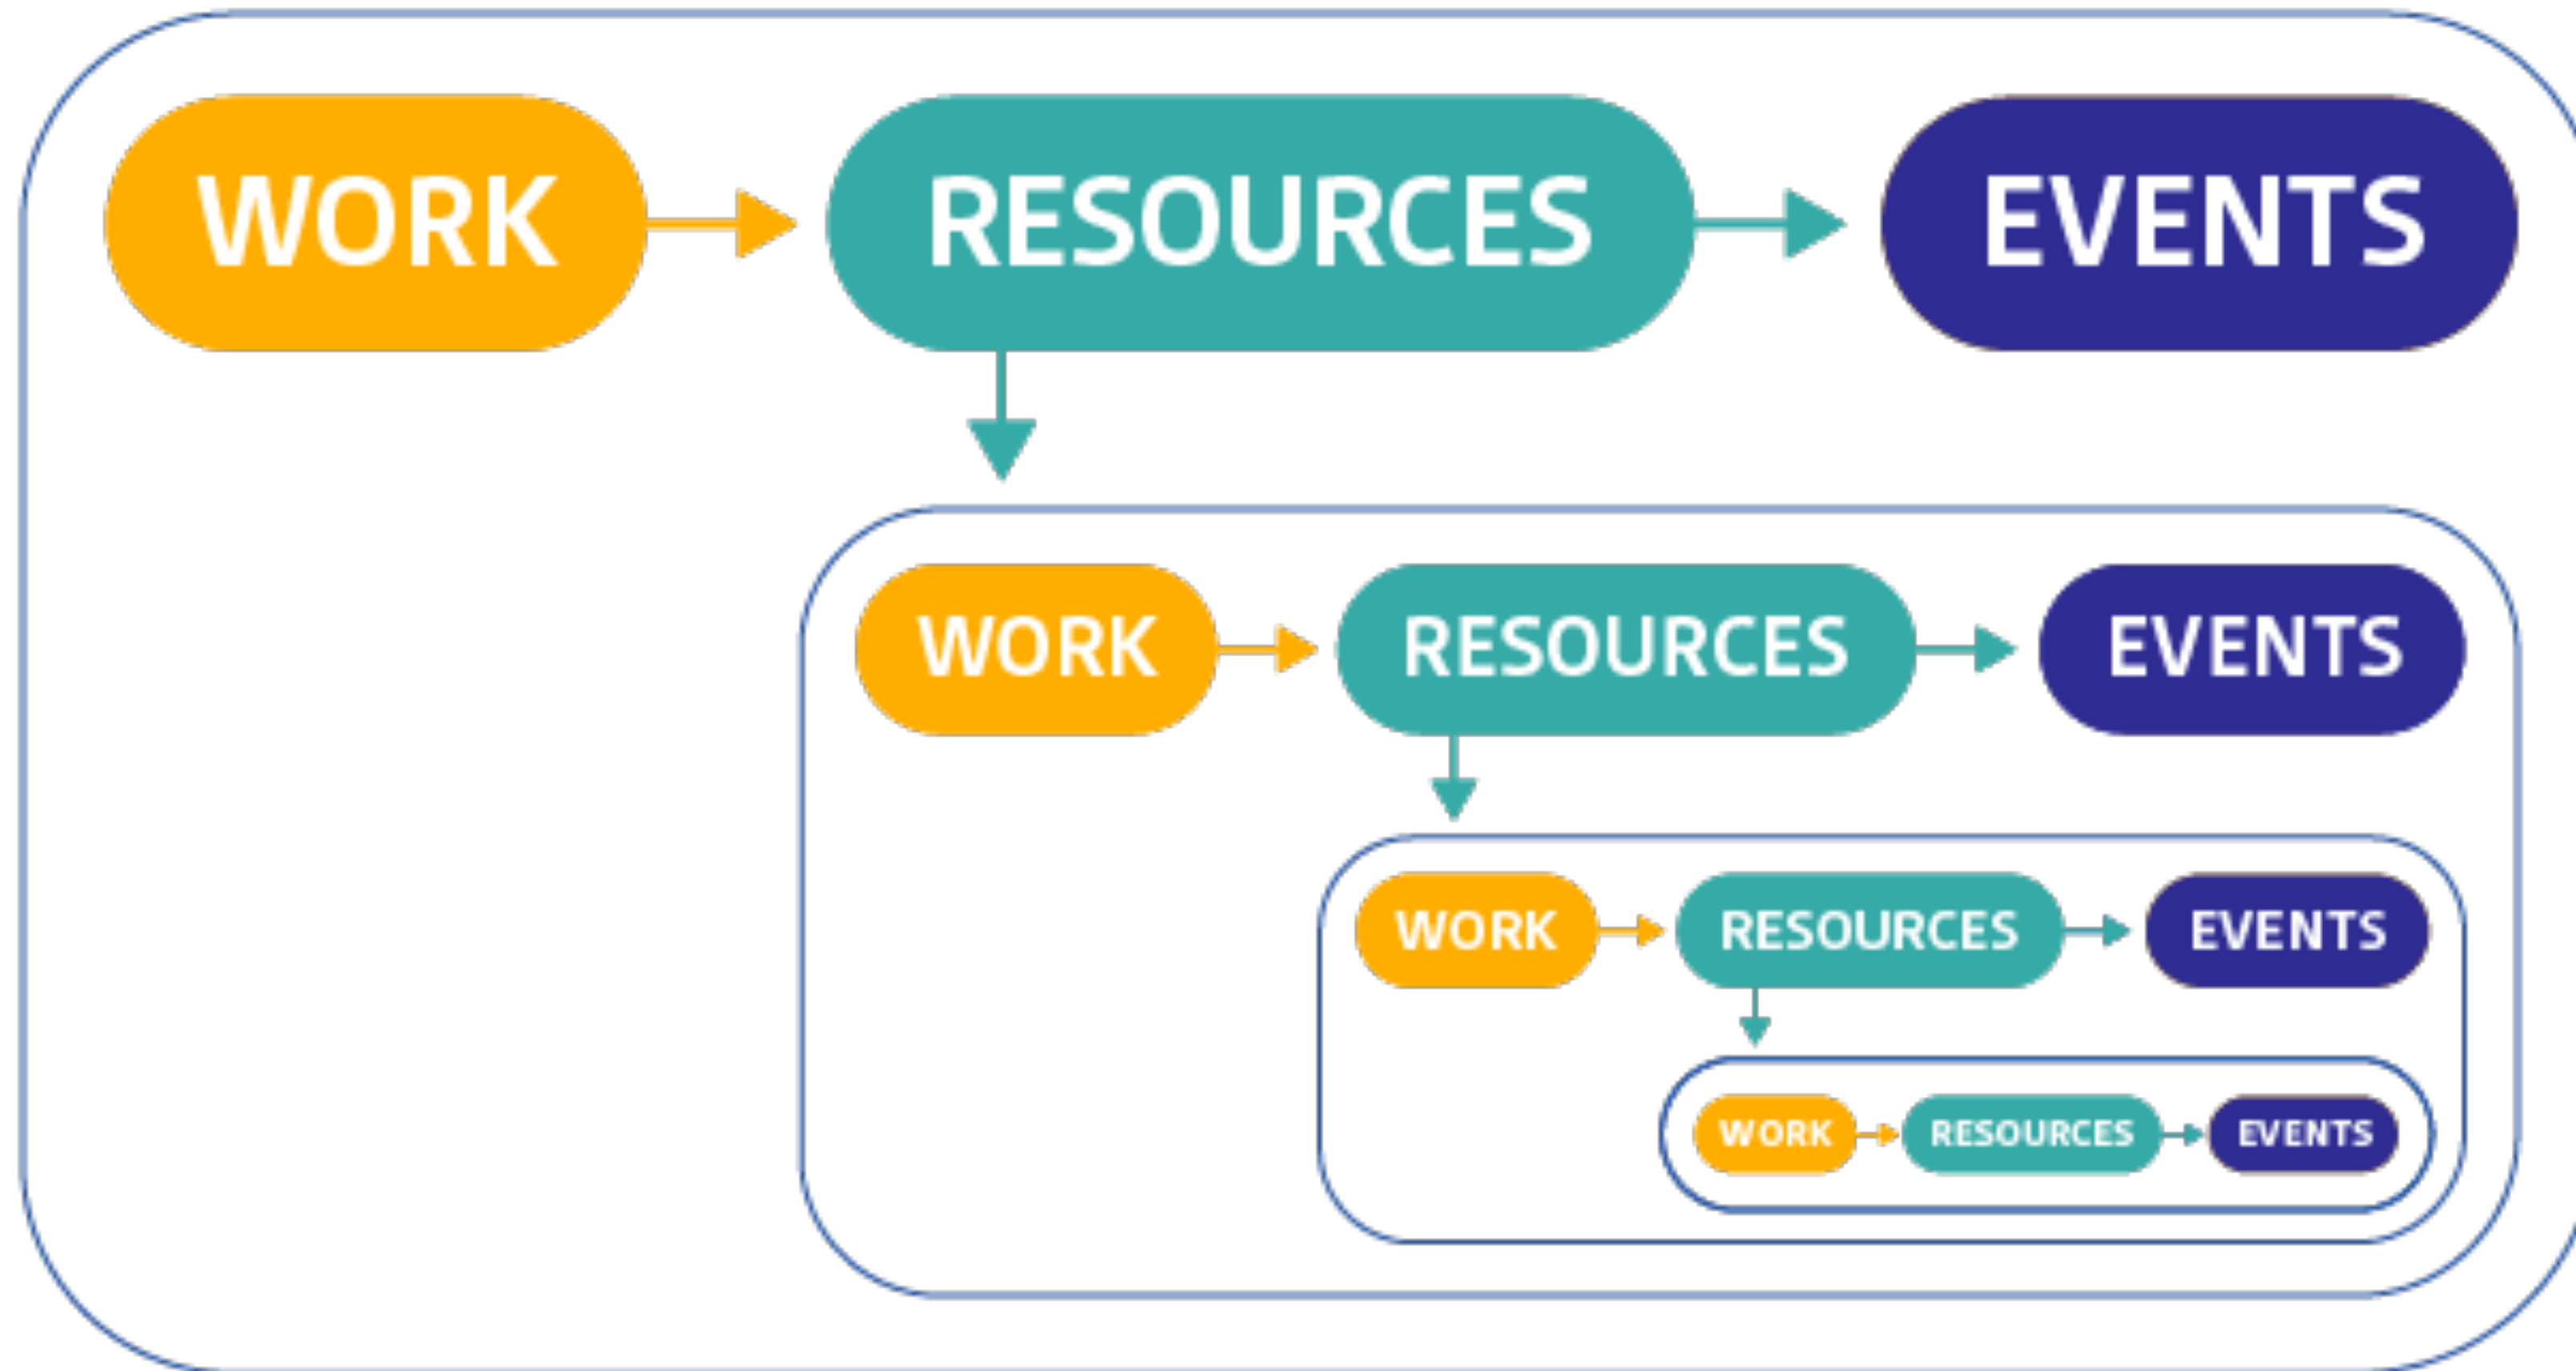
- **Configuration Change**
- **Code Deployment / Release**
- **Add / Remove Nodes**
- **Cache Purge**

**When to let a sleeping
engineer lie?**

When to alert?



Recurse until you find root cause



The Life of a Metric

Monitor Everything



AWS



Docker



CoreOS



Chef



Puppet



Github



Pagerduty



Nagios



Go



Postgres



Java



VMware



Redis



MySQL



Apache



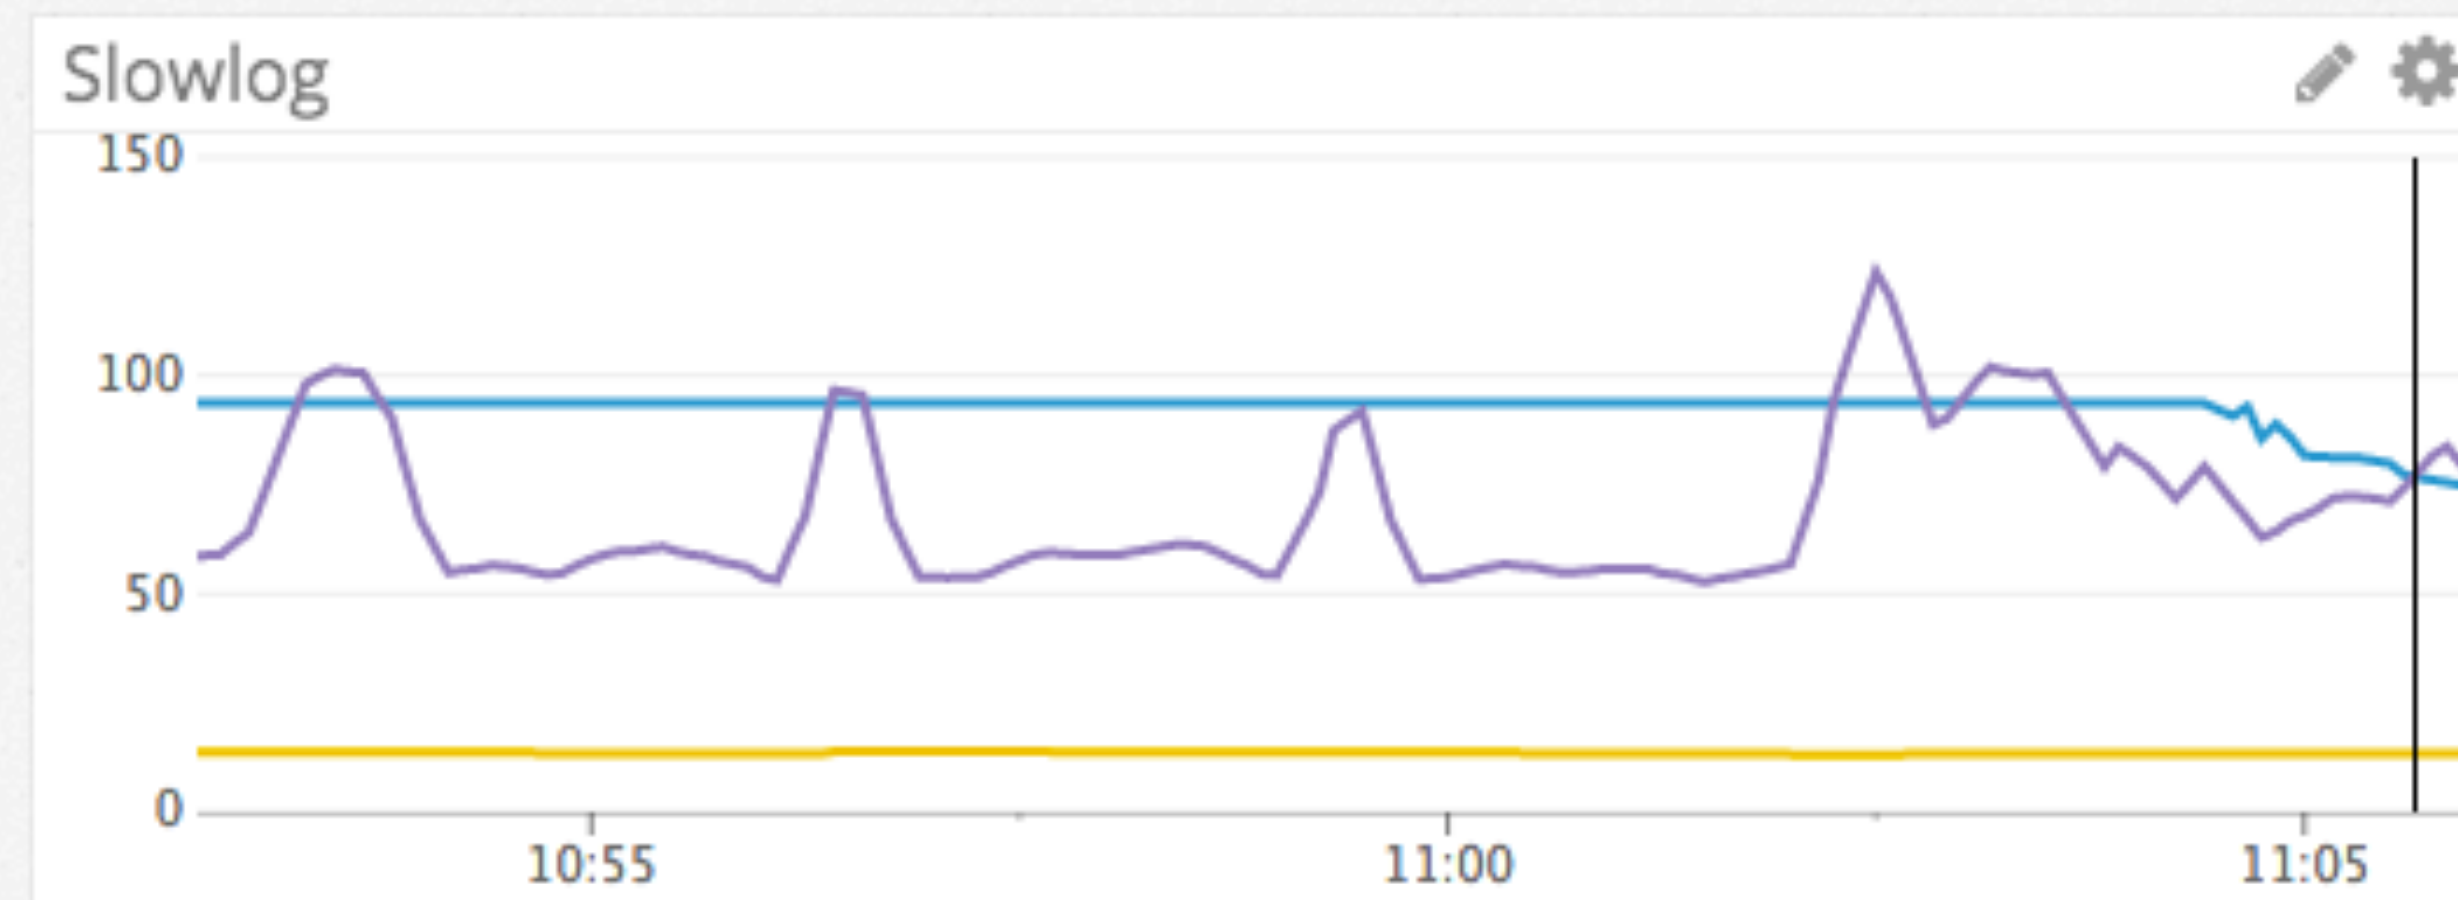

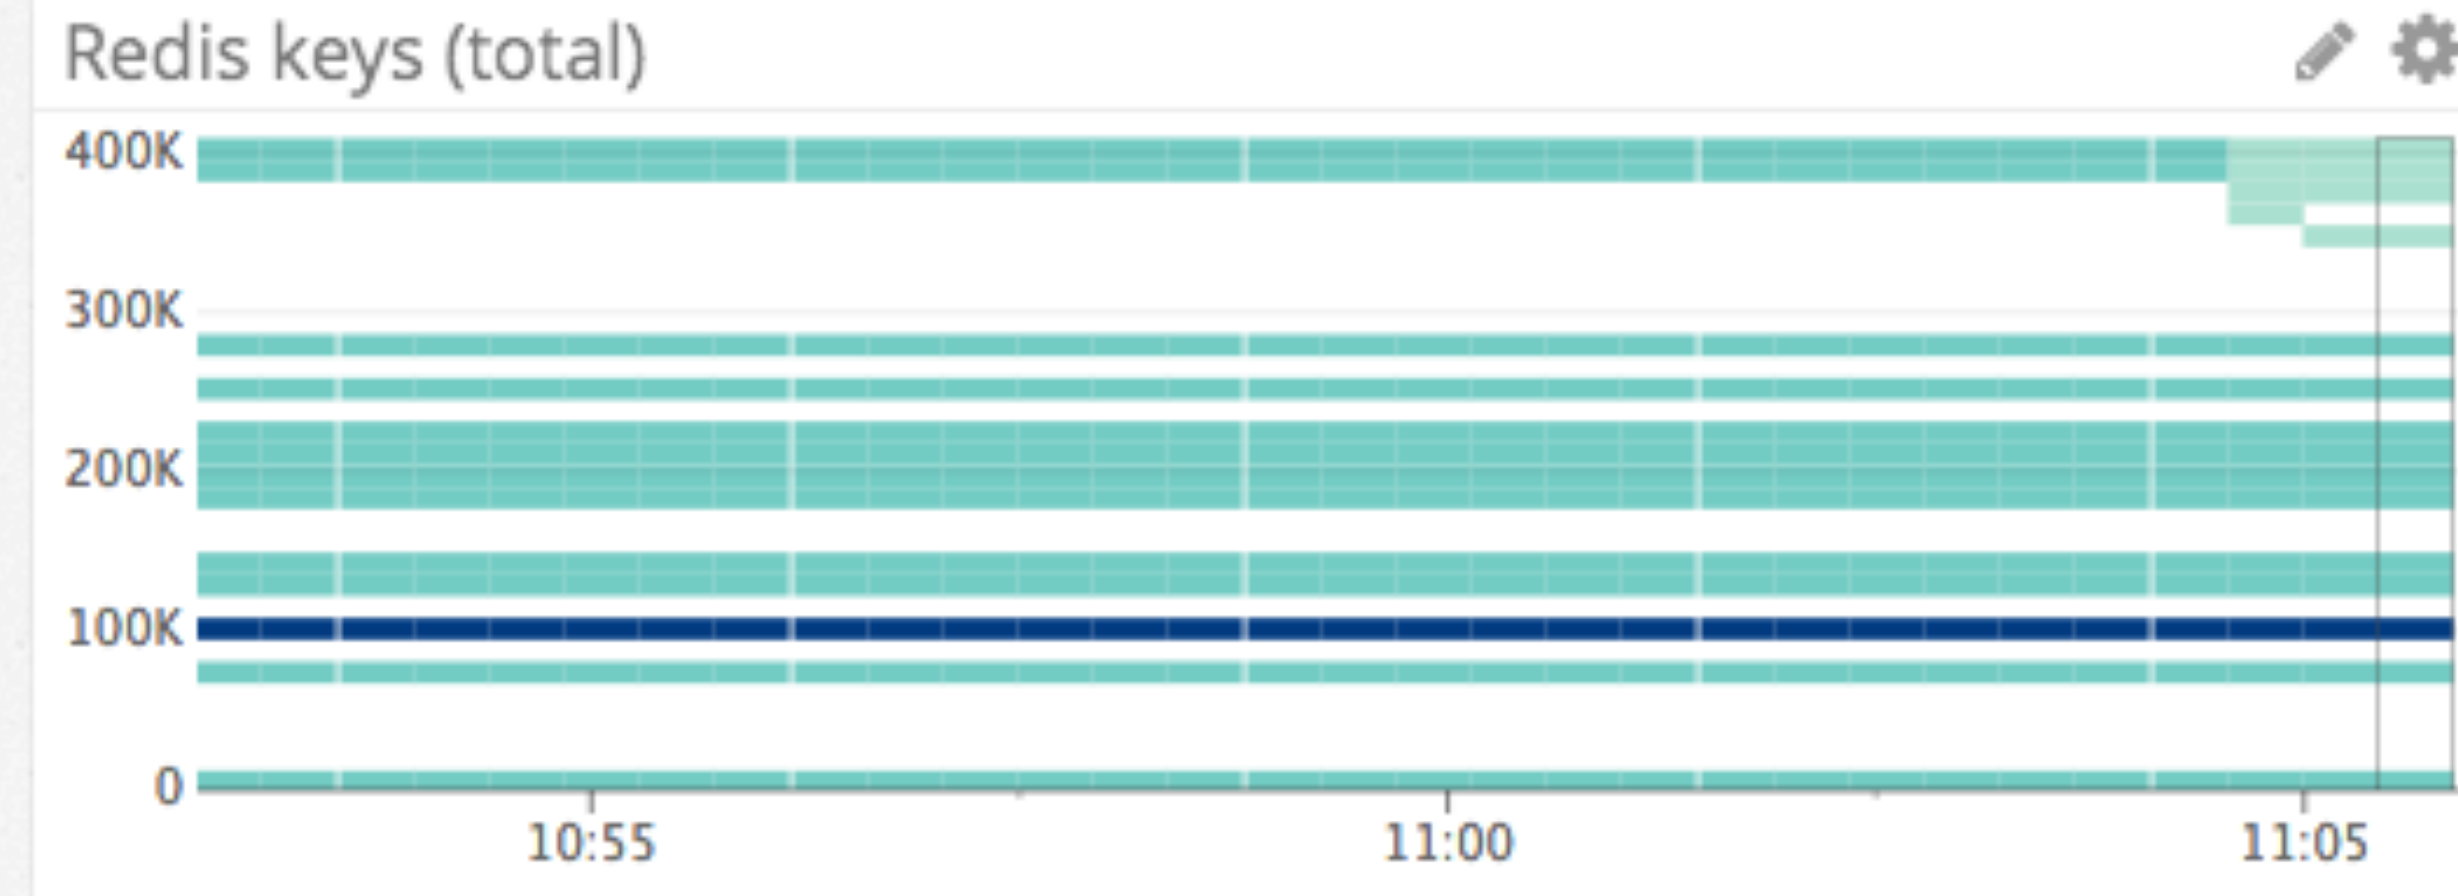
Tomcat



MongoDB



New Relic



Triggered Monitors

Manage Monitors

Manage

All

Filter monitors by name, groups,



STATUS

NAME ↑



ALERT

Load is high on a t...



ALERT



Load is high on a t...



ALERT

CPU usage >25%



ALERT

Load is high on a t...



ALERT



test celene system...



ALERT

Simple monitor



ALERT

Load is high on a t...



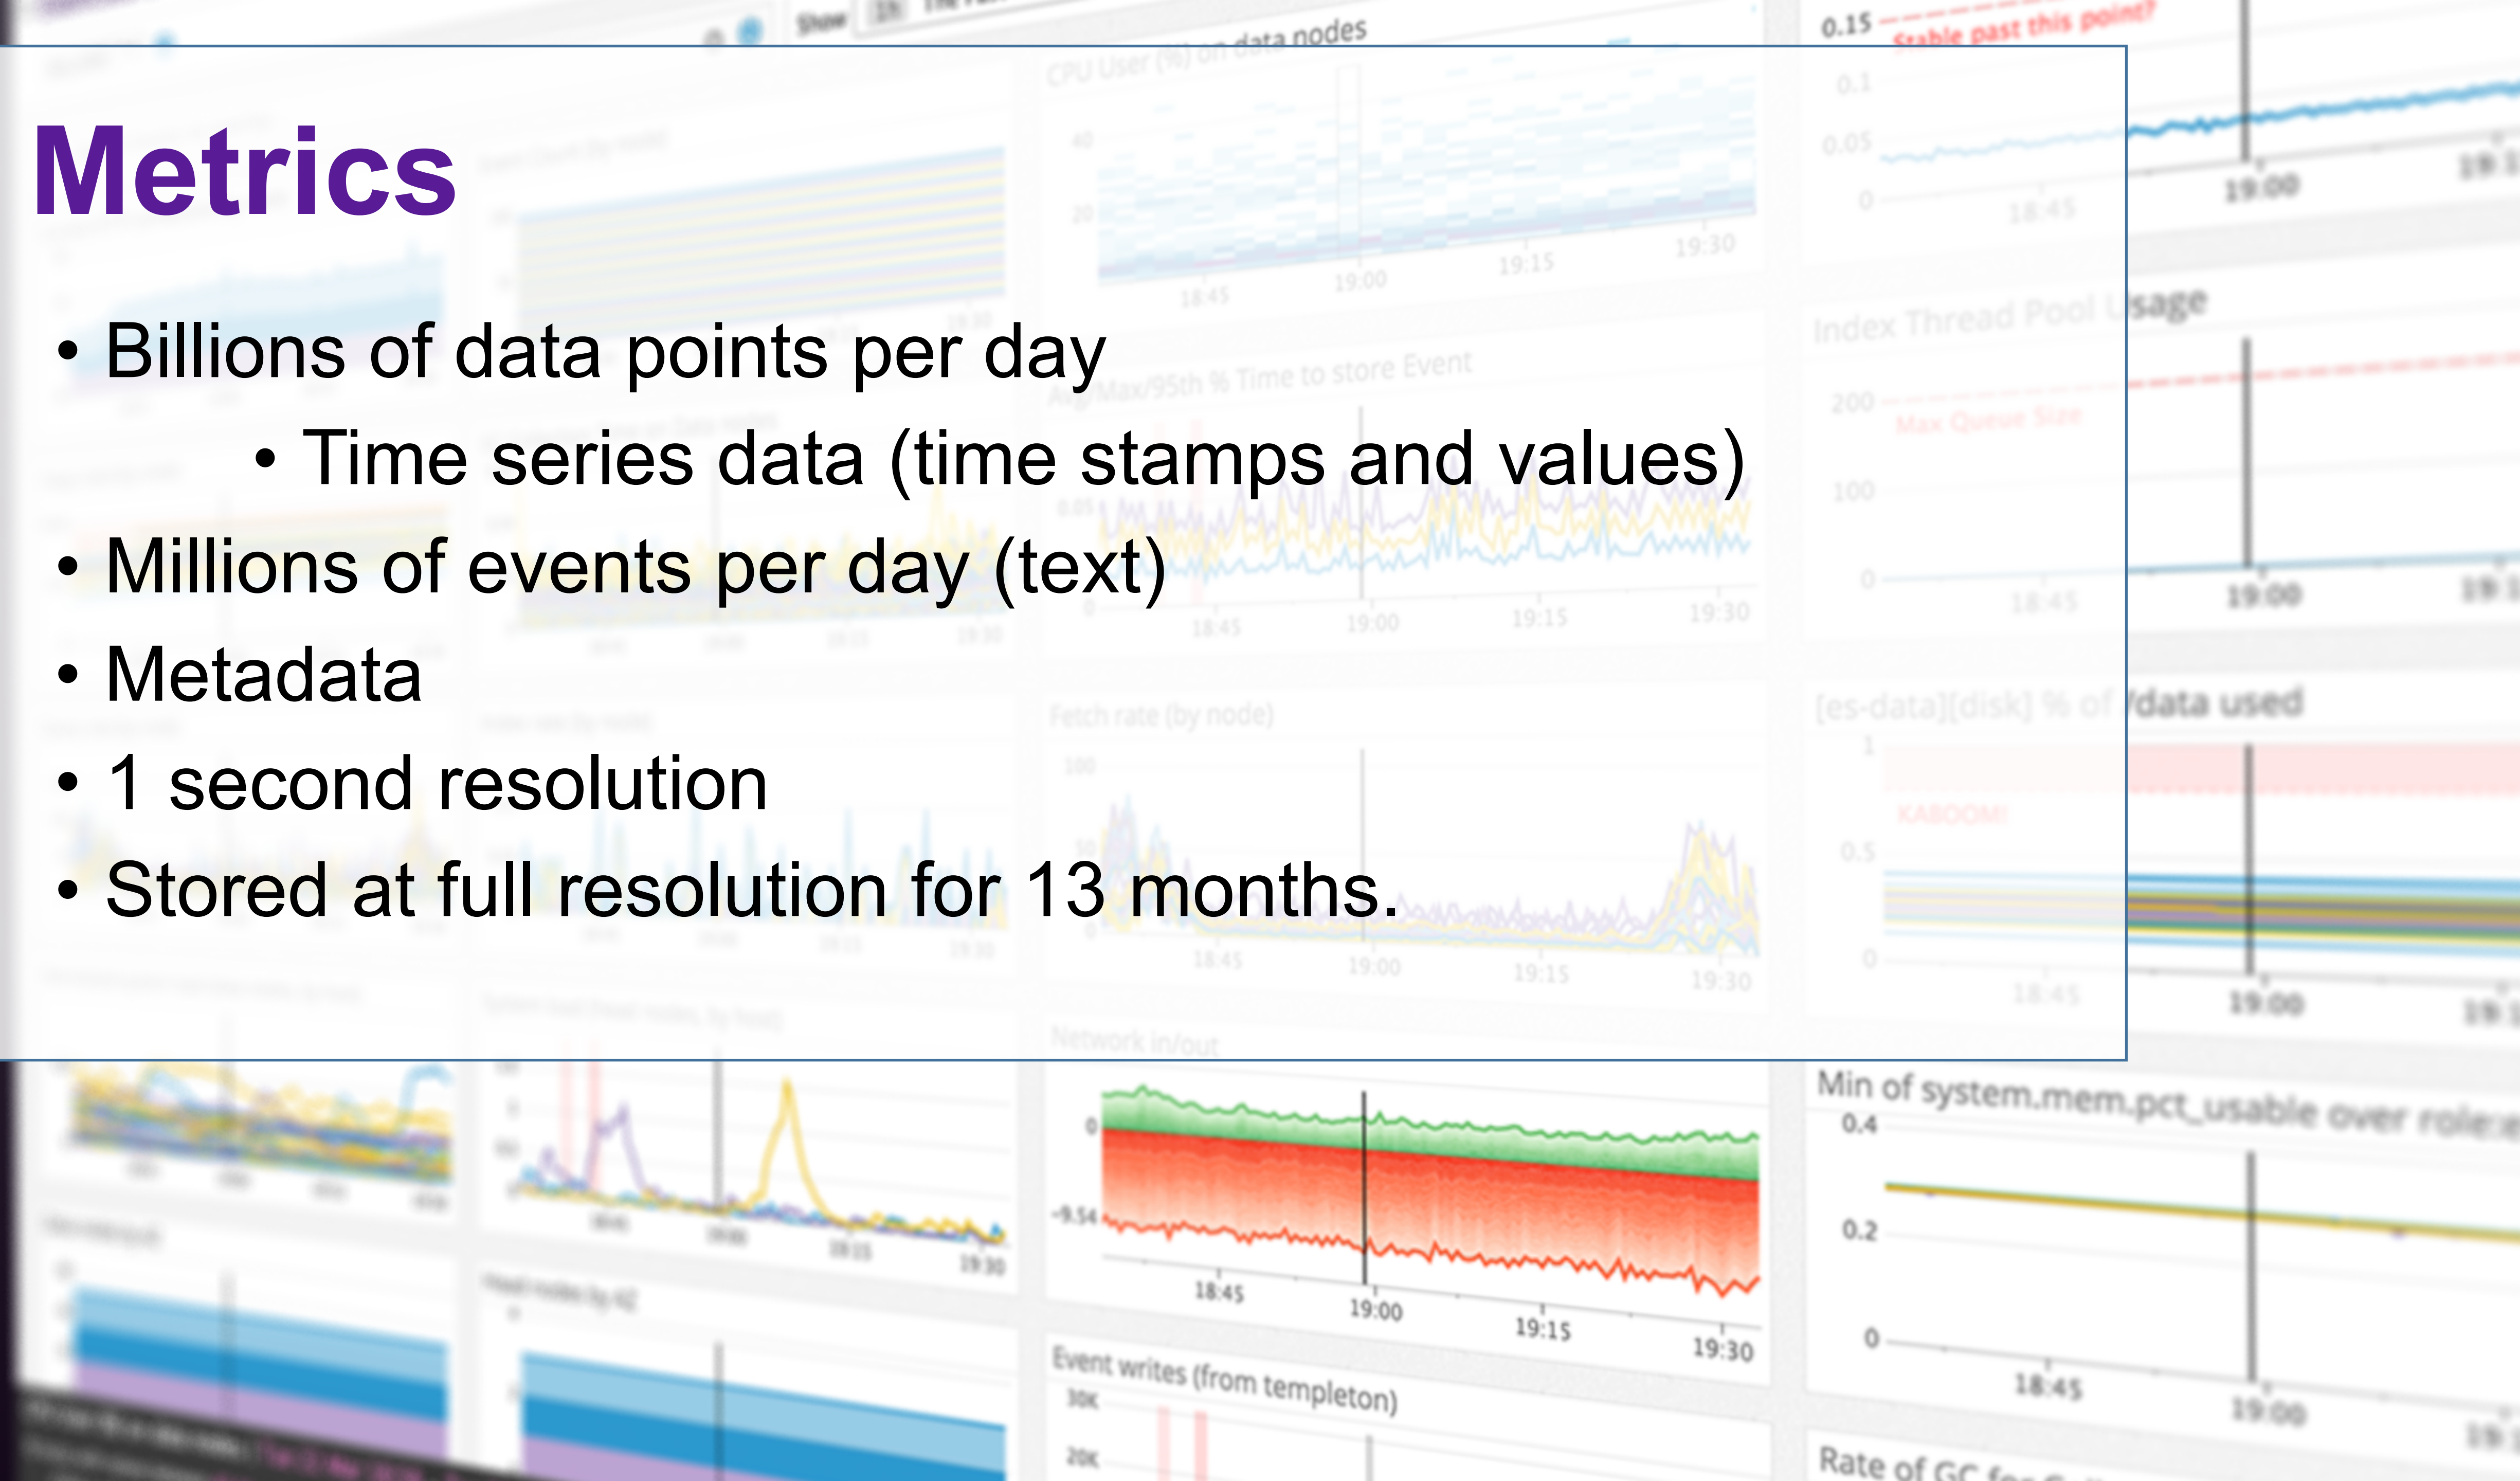
ALERT



test celene system...

Metrics

- Billions of data points per day
 - Time series data (time stamps and values)
- Millions of events per day (text)
- Metadata
- 1 second resolution
- Stored at full resolution for 13 months.





- Events
- Dashboards
- Infrastructure
- Monitors
- Metrics
- Integrations

- Help
- Team
- matt.williams (Datadog HQ)

Search Events...

Show 2d

Search full event text, or target specific event properties using these prefixes:

users: `pup@datadoghq.com`

sources: `github, chef`

tags: `env-prod, db`


hosts: `db1.myapp.com, db2.myapp.com`

status: `[error, warning, success]`

priority: `[low, normal, all]`

incident: `[all, open, claimed, resolved]`

[More search help →](#)



build #13136 started on roy.datadog.com #job:staging-dogorphans-kill

ago · Add comment · Lower priority

build #1375 started on roy.datadog.com #job:consul-config-pre-commit

ago · Add comment · Lower priority

- FlowDock
 - Github
 - HAProxy
 - Jenkins
 - Monitor Alert
 - My Apps
 - MySQL
 - New Relic
 - PagerDuty
 - Pingdom
 - Puppet
 - RabbitMQ
 - Sentry
 - Slack
 - StatusPage
 - Users
 - VictorOps
 - Webhooks
- PRIORITY
- All
 - Normal
 - Low
- STATUS

Arkelenia pushed to master at DataDog/consul-config #branch:master #event

There was 1 commit in this push by #Arkelenia

[Compare URL](#)

Commits:

Updated 4 days ago · Created 4 days ago · Add comment

Memory usage is

pid	command	memory usage
26992	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-0 --skip-maintenance-tasks	121 MB
27599	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-6 --skip-maintenance-tasks	121 MB
27328	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-3 --skip-maintenance-tasks	121 MB
23361	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-0 --skip-maintenance-tasks	121 MB
23353	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-7 --skip-maintenance-tasks	121 MB
14804	/usr/bin/python3 /usr/local/bin/delancie --config_file=/etc/datadog/dogweb.ini worker --flavor query-cache-warmer --log-root /mnt/log/delancie --name delancie-4 --skip-maintenance-tasks	121 MB

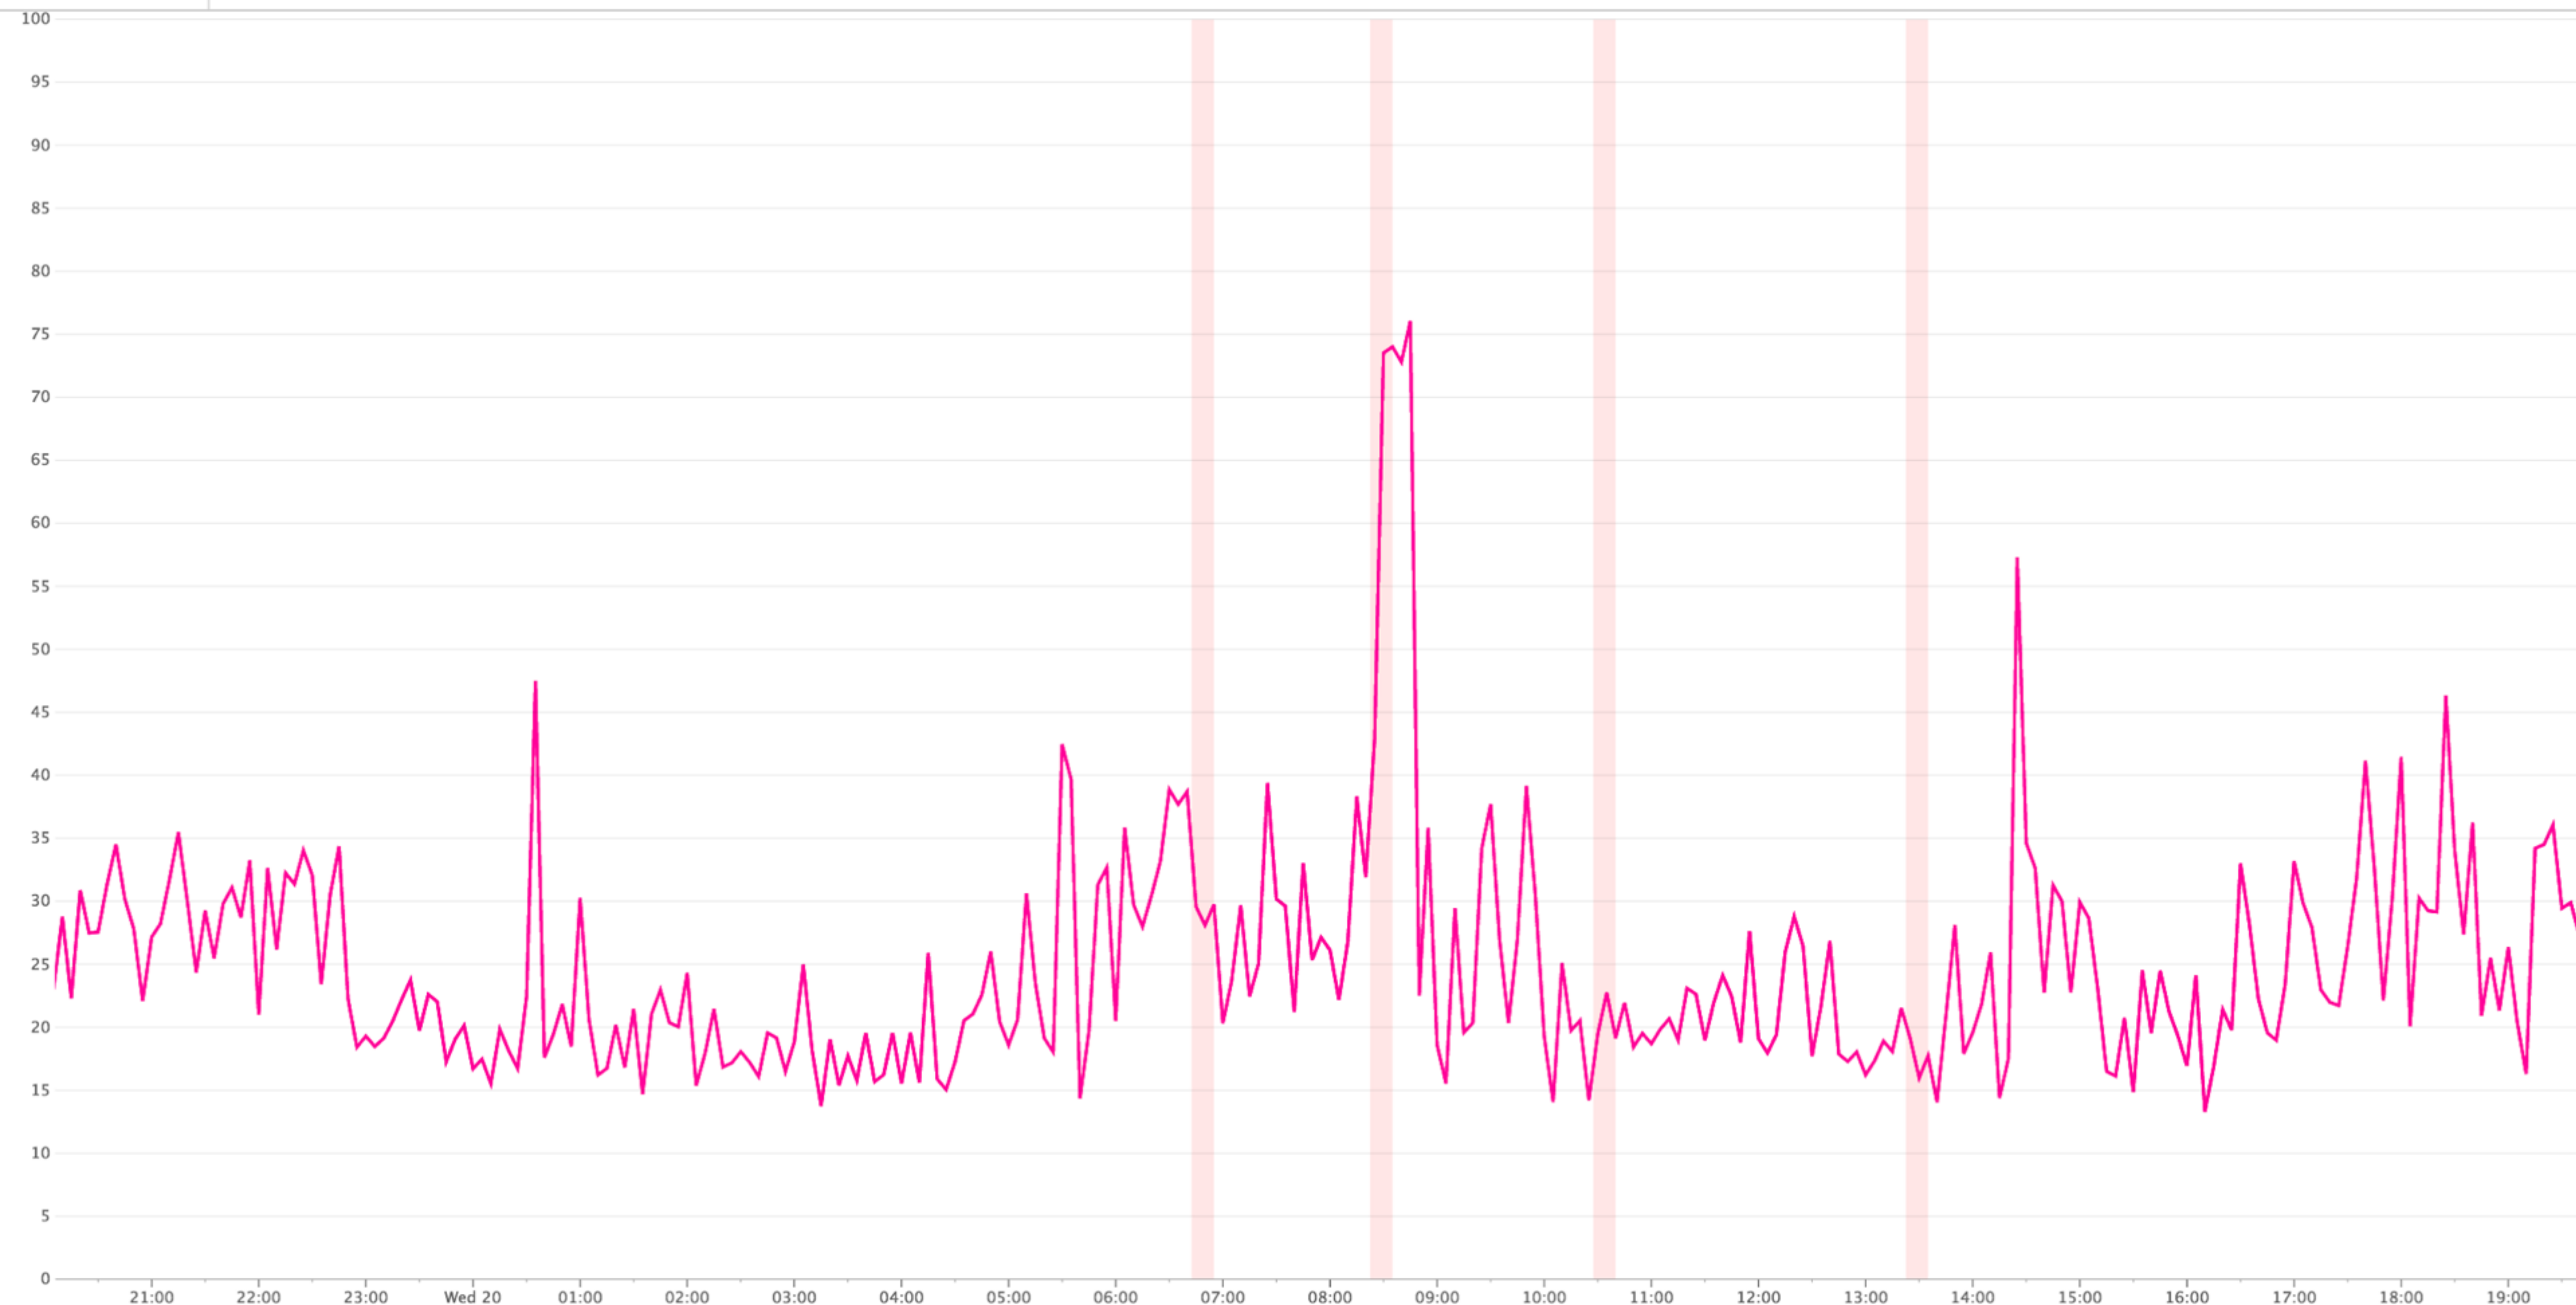
Show more

Updated 4 days ago · Created 4 days ago · Add comment · Lower priority

Primarily ElasticSearch

I/O wait (%)

Show 1d The Past Day




A metric is born...

- Open Source Python agent
- SDKs, Libraries, RESTful APIs



A metric is born...

- Open Source Python agent
- SDKs, Libraries, RESTful APIs
- SaaS Integrations



Agents / SDKs




Backend Systems
& Processing



```
{
  "series": [
    {
      "metric": "<metric-name>",
      "points": [
        [ <timestamp>, <value> ]
      ],
      "type": "gauge",
      "host": "<host name>",
      "tags": [ "<tags>" ]
    }
  ]
}
```

DATAPOINT

SYSTEM.NET.BYTES_RCVD 4 2016-03-02 15:00:00 ['FILE-SERVER']




Tags All the Way Down

Launch Instance **Connect** **Actions** ▾

🔍 Filter by tags and attributes or search by keyword

<input type="checkbox"/>	Name ▾	Instance ID ▾	Instance Type ▾	Availability Zone
<input type="checkbox"/>	m3.medium	us-east-1a
<input type="checkbox"/>	r3.large	us-east-1c
<input type="checkbox"/>	m3.xlarge	us-east-1a
<input type="checkbox"/>	t2.micro	us-east-1a



Asking Better Questions

“Monitor all containers running image web in region us-west-2 across all availability zones that use more than 1.5x the average memory on c3.xlarge”


Asking Better Questions

“90% of all web requests are taking more than 0.5s to process and respond.”

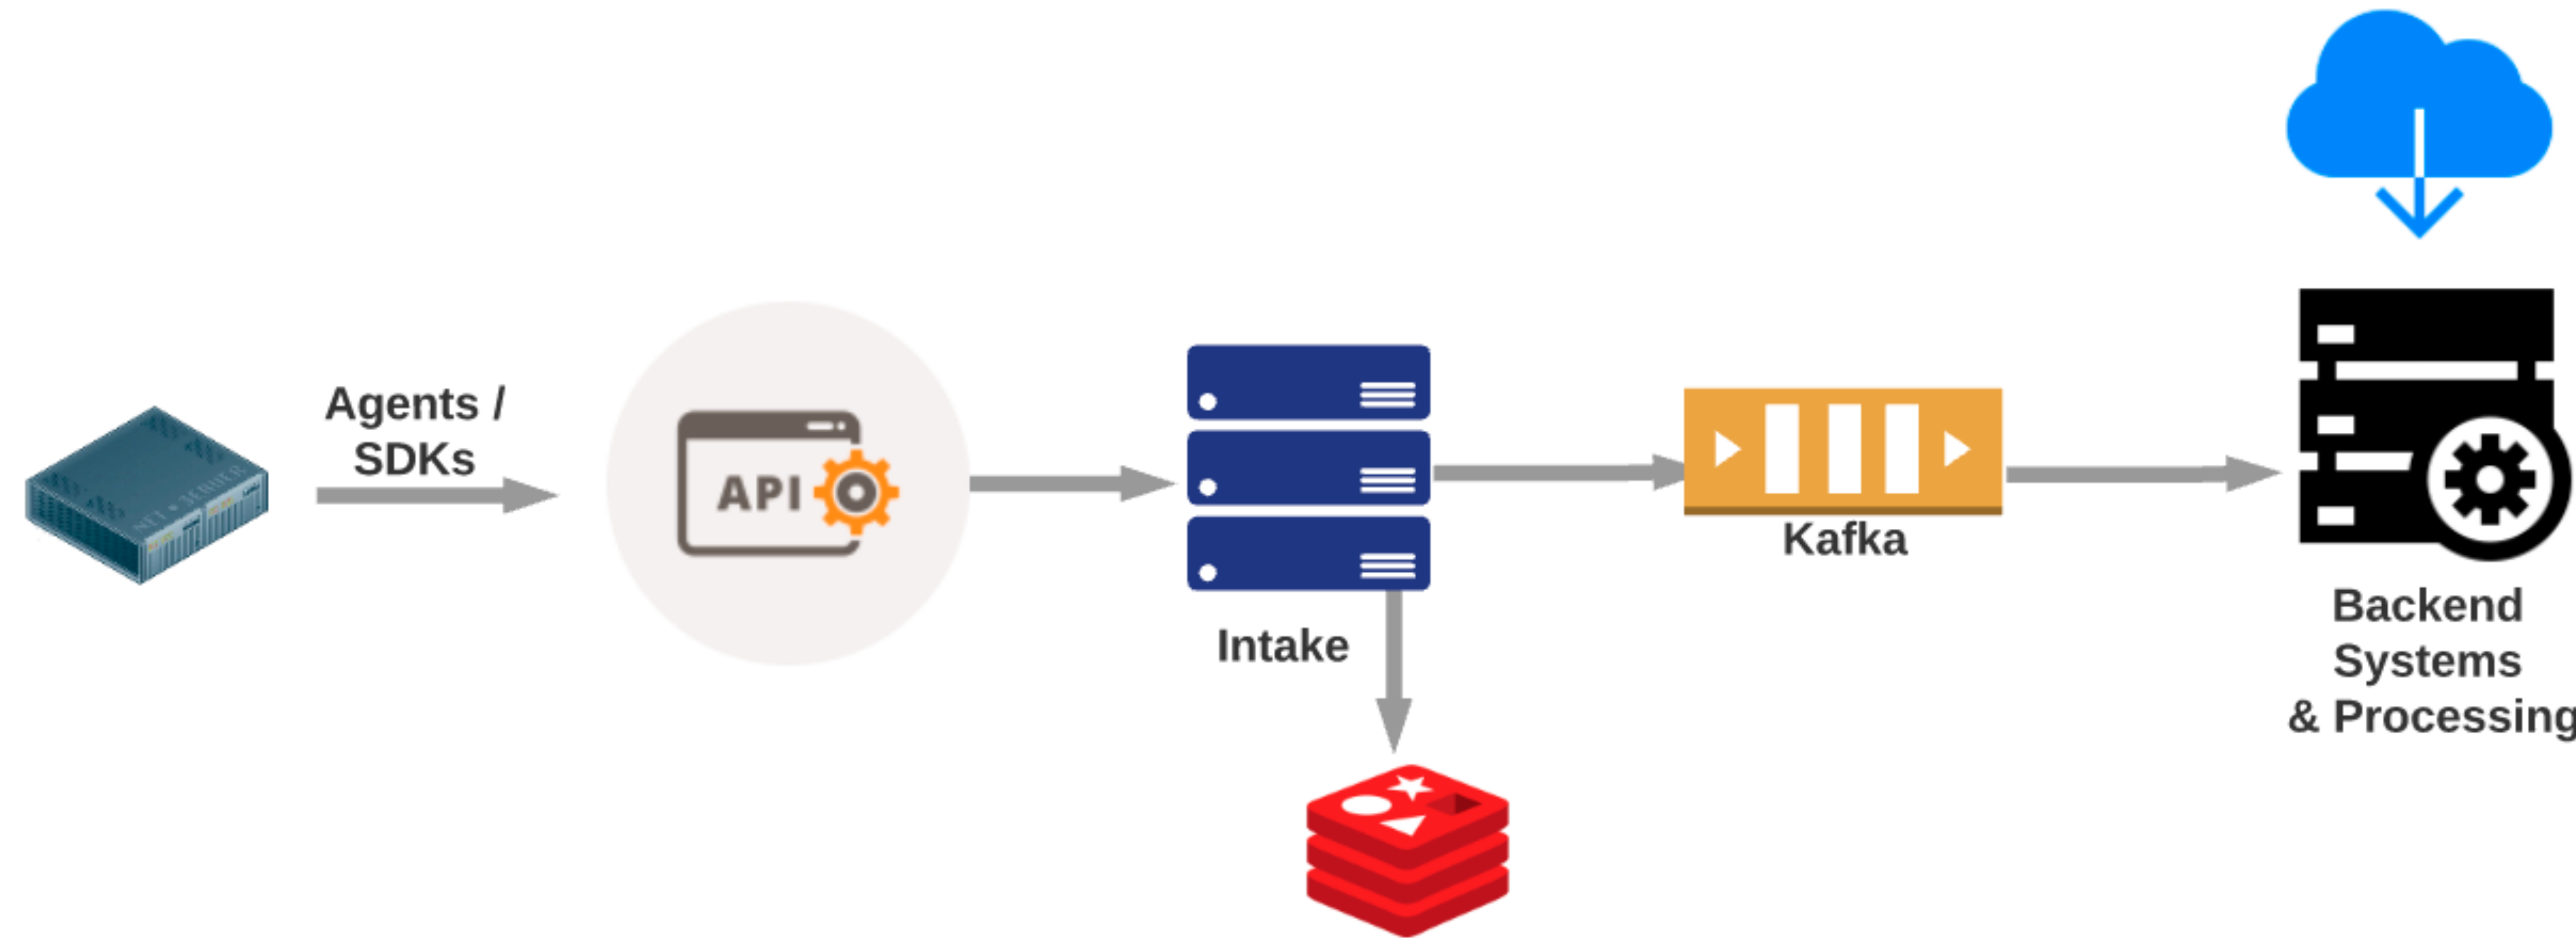
```
{
  "series": [
    {
      "metric": "system.cpu.system",
      "points": [
        [ 419707187, 5 ]
      ],
      "type": "gauge",
      "host": "test.example.com",
      "tags": [ "environment:prod" ]
    }
  ]
}
```

Intake APIs

- Constant stream of data
- Low latency requirements (30-60s)
- Data needed by multiple consumers/systems.



Redis Use Case #1: Queuing



Why Redis?

- Easy to scale vertically
- Simple push/pop interaction for queues
- Simple clustering via twemproxy


Work Metrics - Queuing

- **Queue Depth**
- **Message Latency**
- **Cmd Latency**
- **Read vs Write Calls per sec**

Work Metrics - Queuing


Intake-queue calls by command

Show 1h The Past Hour



Work Metrics - Queuing


Redis Commands Per Second

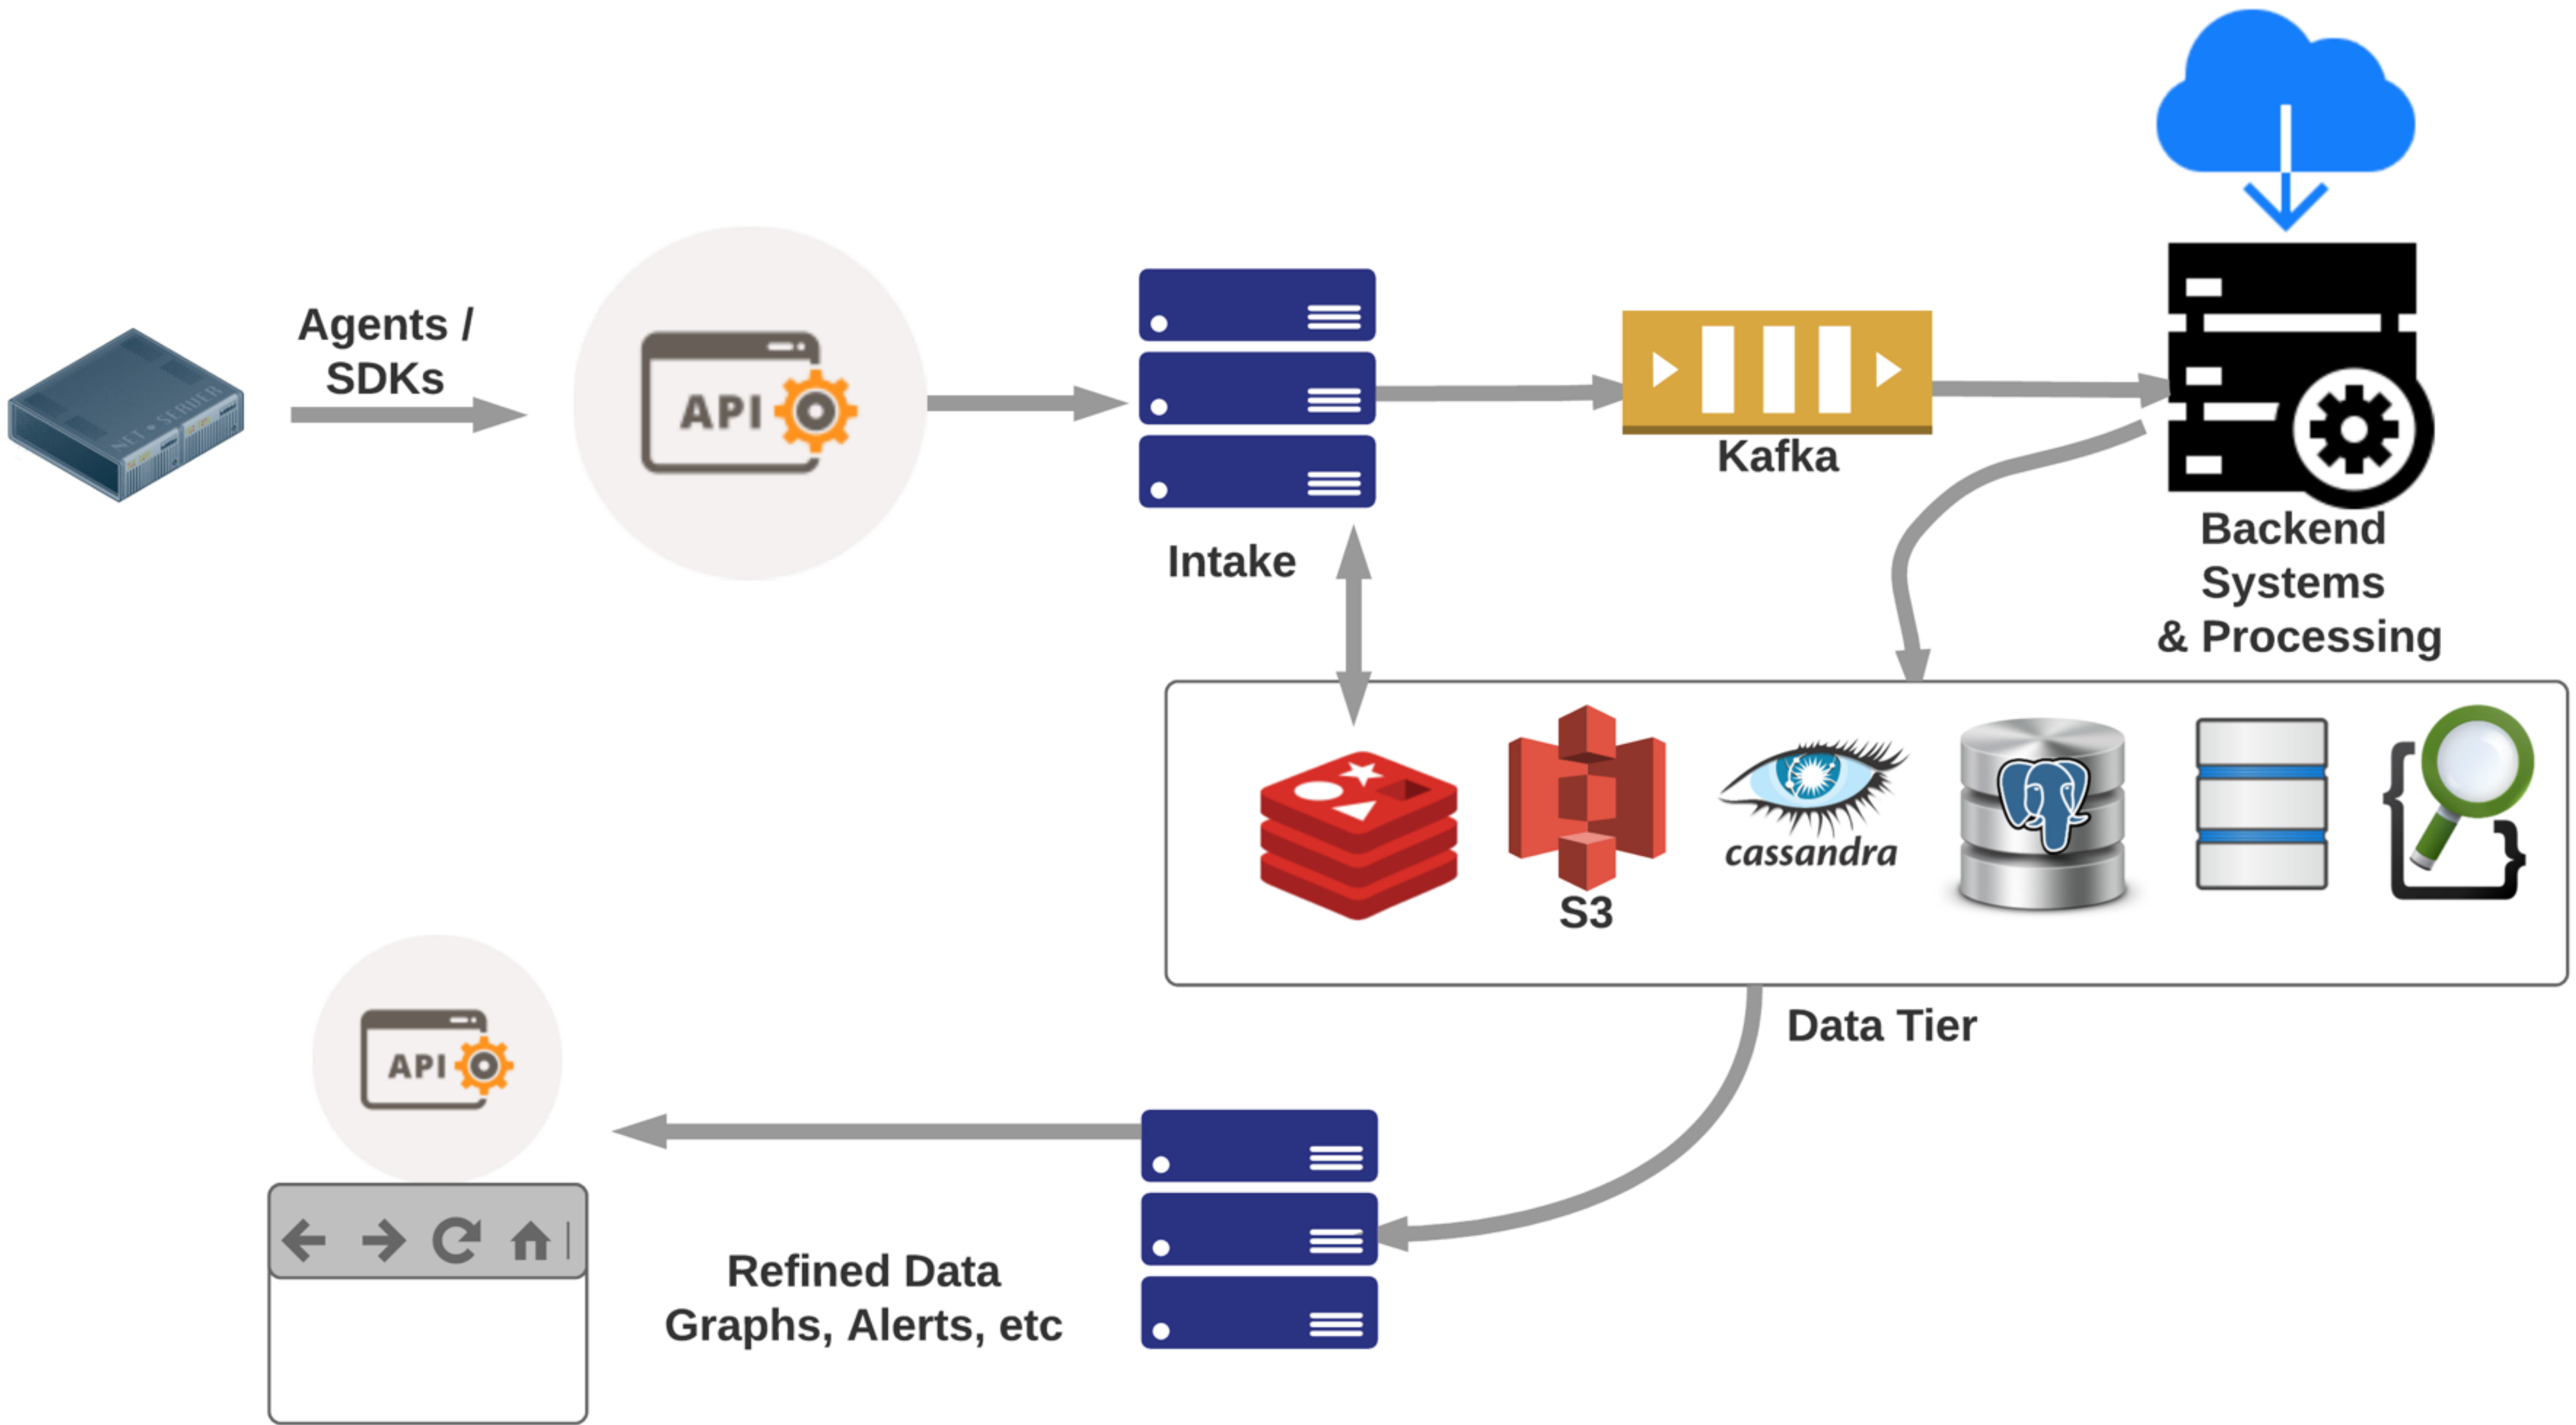


Resource Metrics - Queuing

- **Network Utilization**
- **used_memory**
- **Disk IO**
- **Disk Space**
- **connected_clients**
- **keyspace**
- **rejected_connection**

Redis Keys by Host





Use Case #2 Caching

Initial Architecture

- **Cache Aside**
- **Single Redis per Worker**
- **Local Cache**
- **LRU Caching (allkeys-lru)**




Dogpound

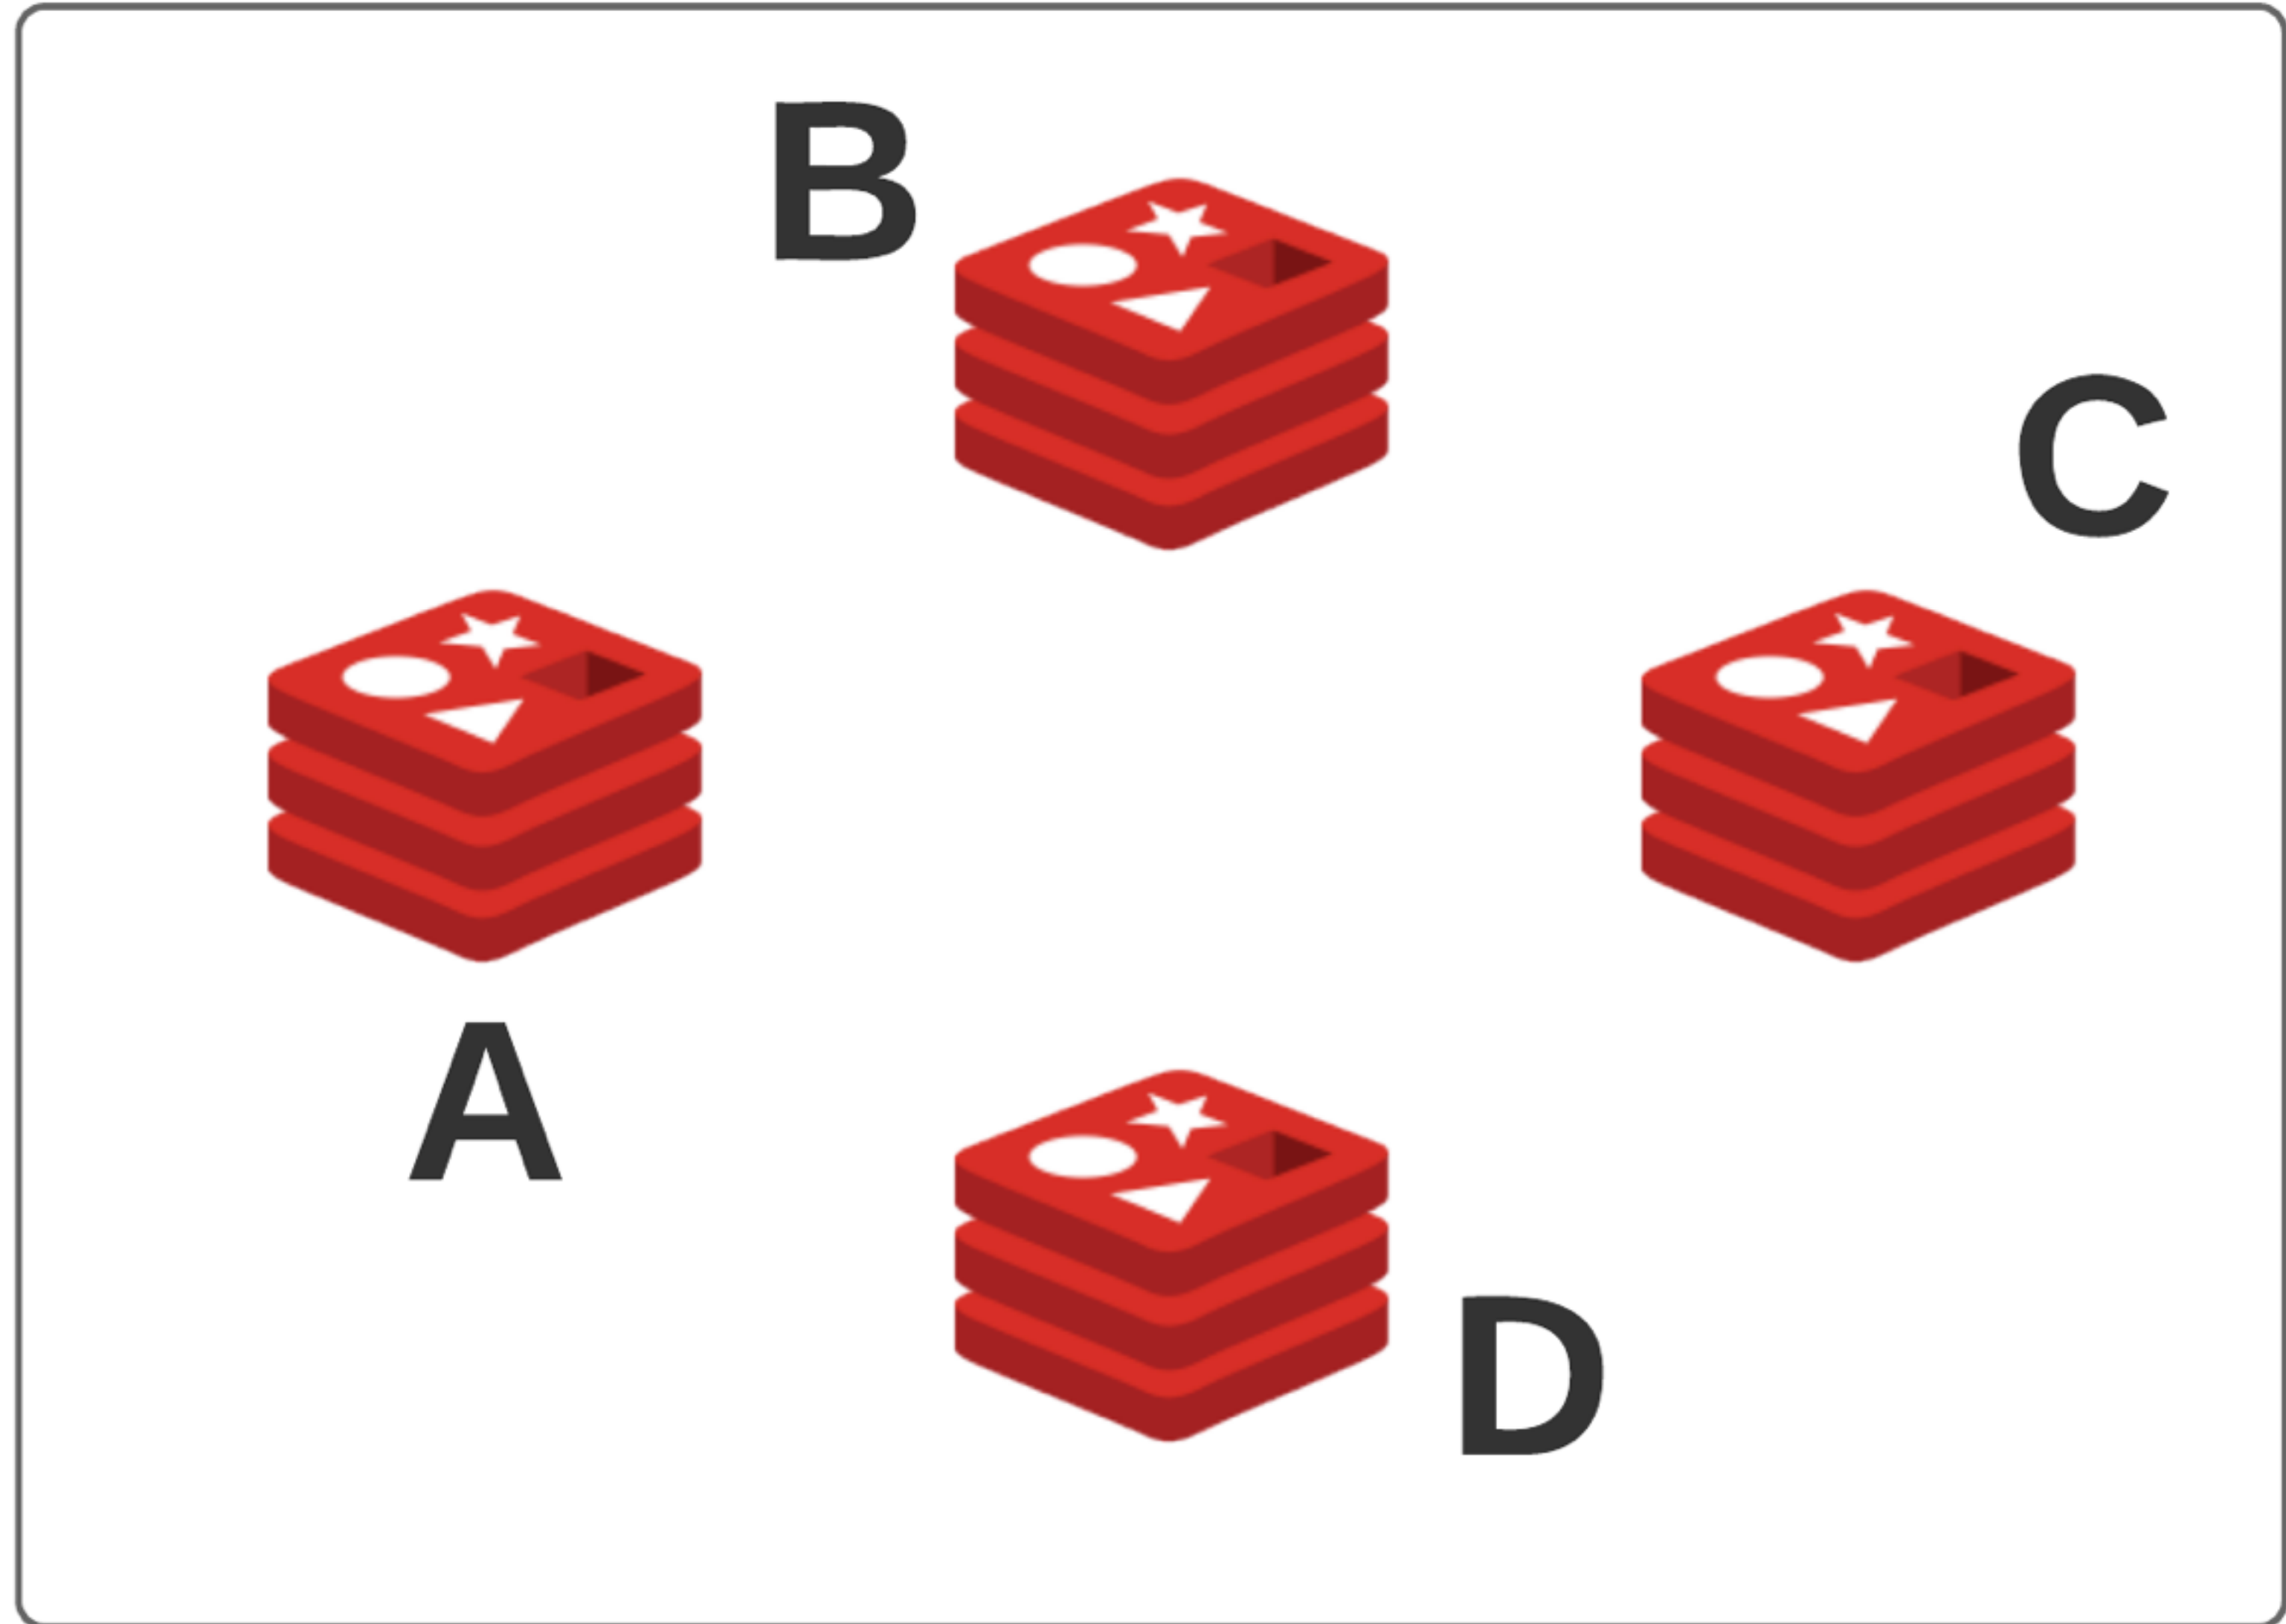

- **Uniform API for caching backends**
- **Pluggable Architecture**
 - **memory**
 - **redis**
 - **sharded redis**
 - **tiered caching**
- **Service Discovery via Consul**



Sharded Redis

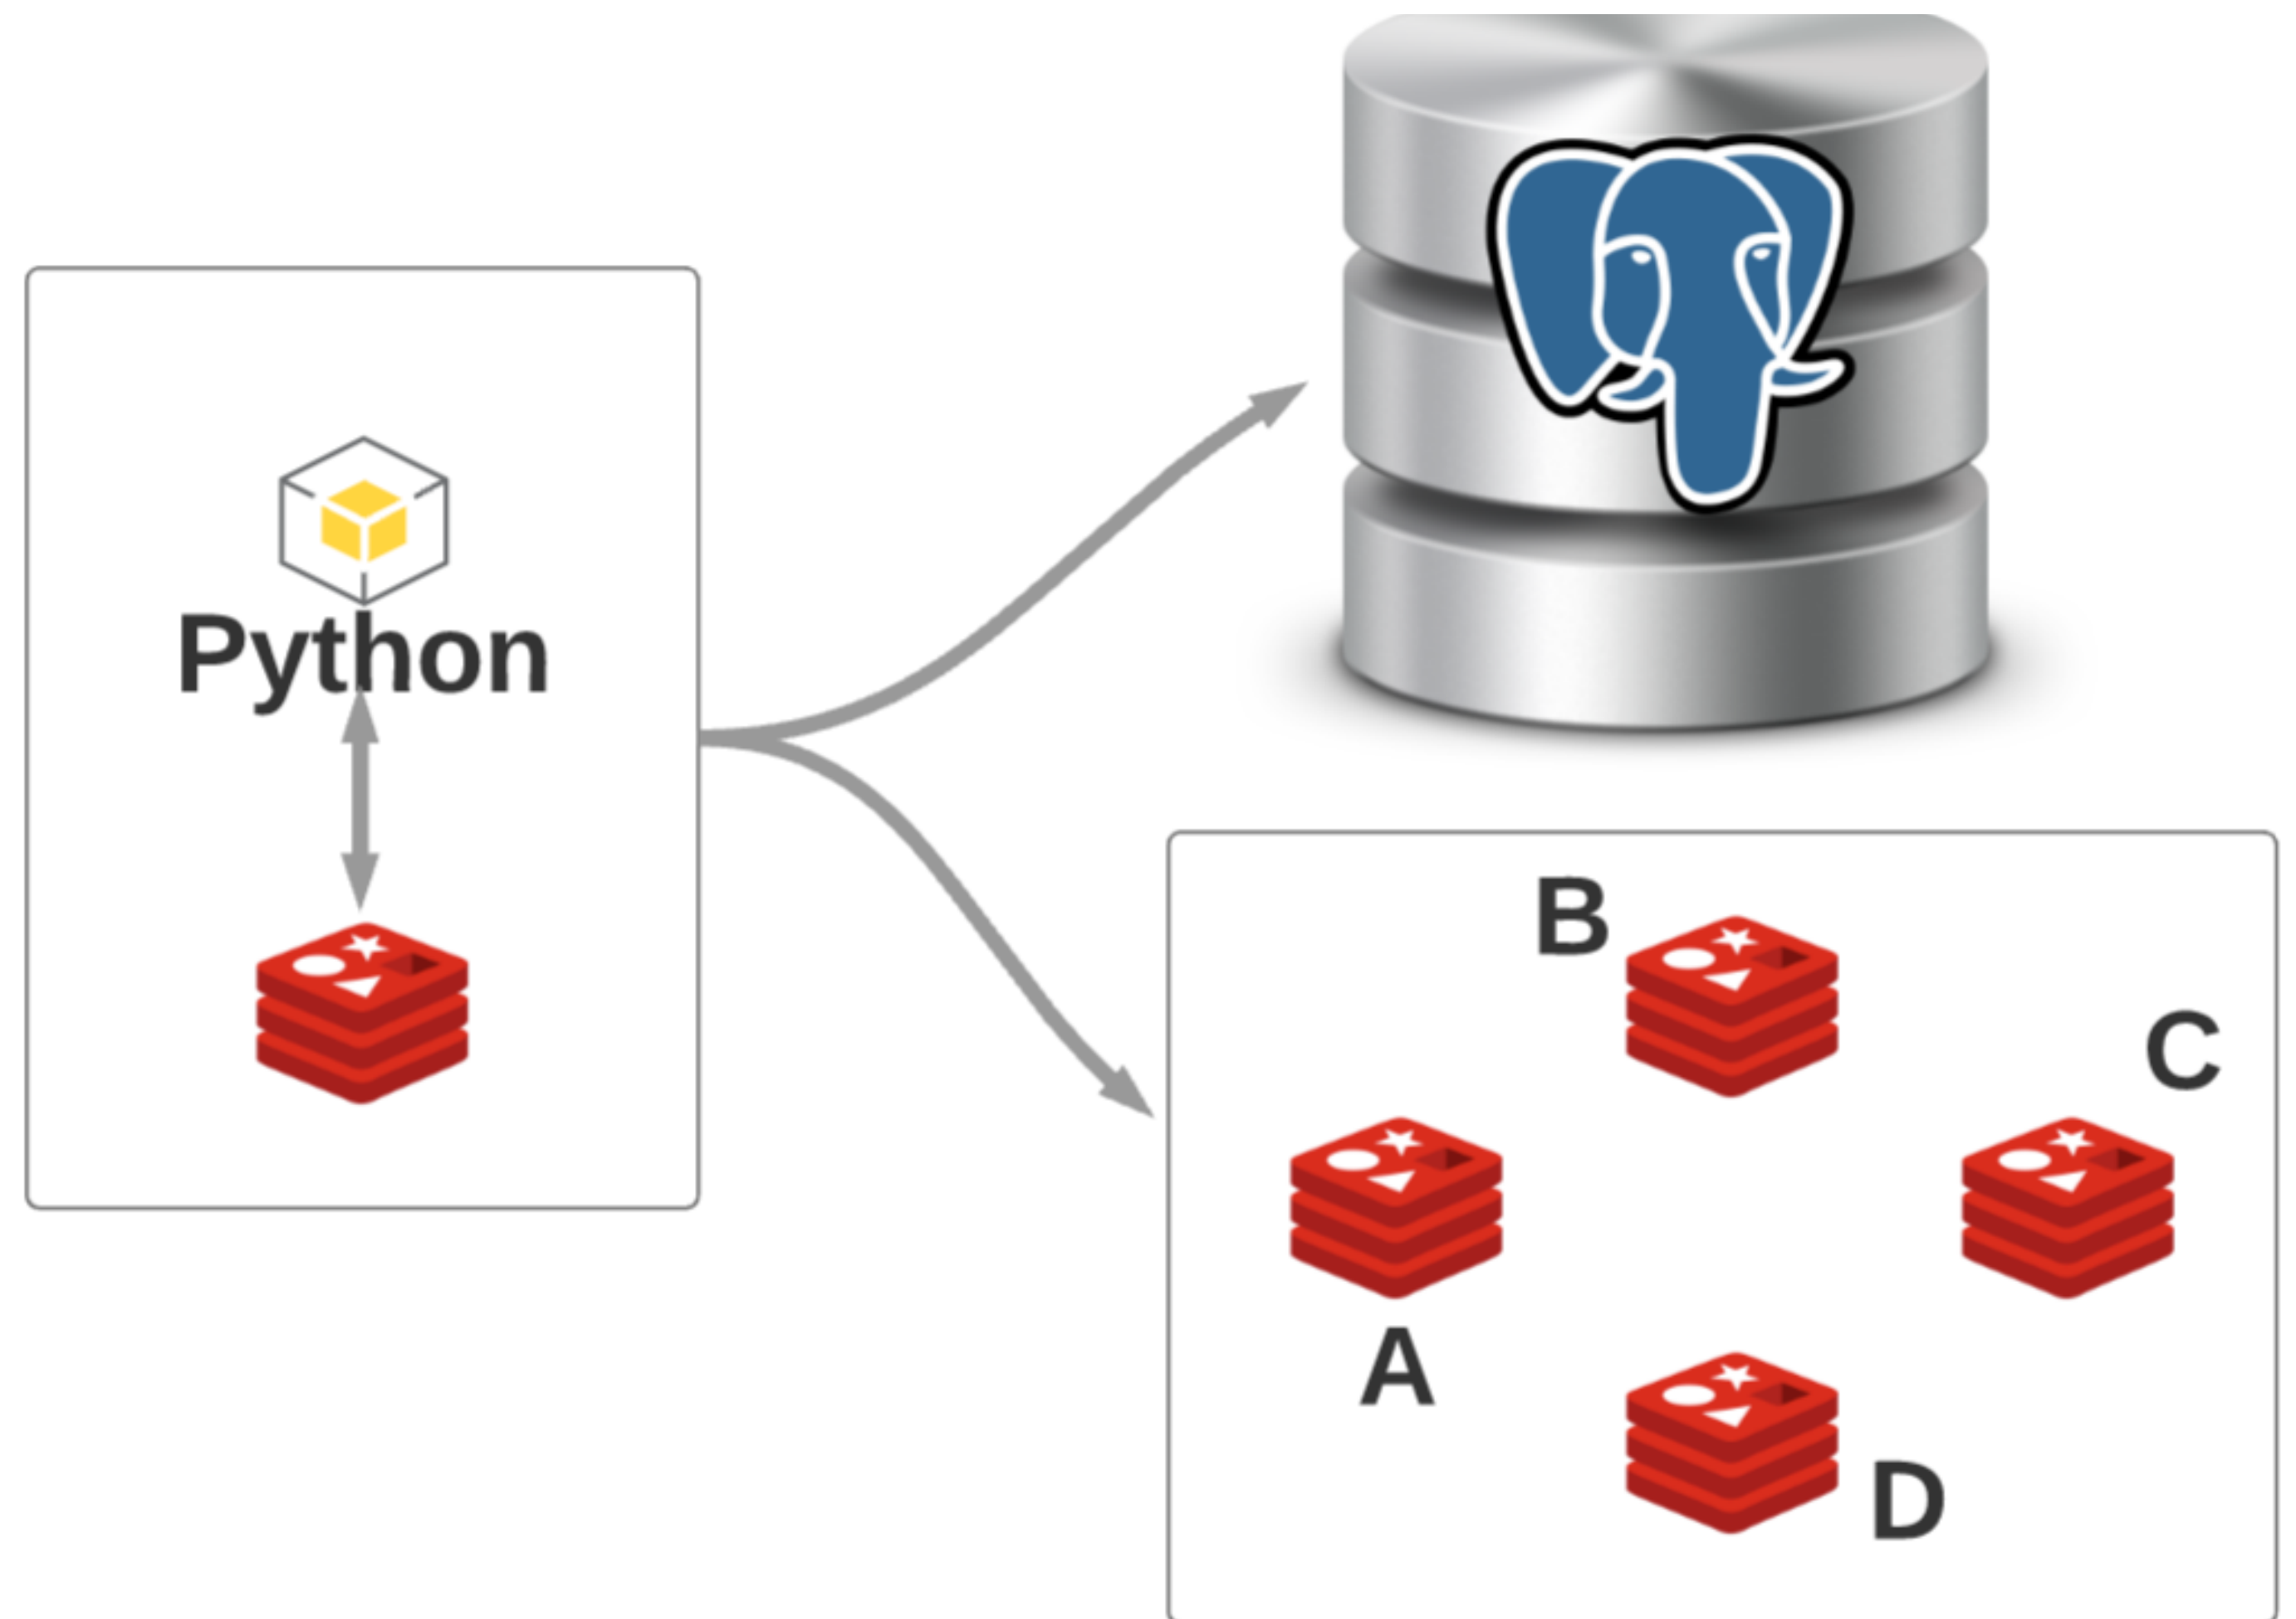


HashRing



Tiered Cache Redis

- Local Redis
- Shared Cache
 - Tiered
 - HA Pair
- Fallback to primary data store




Work Metrics - Caching


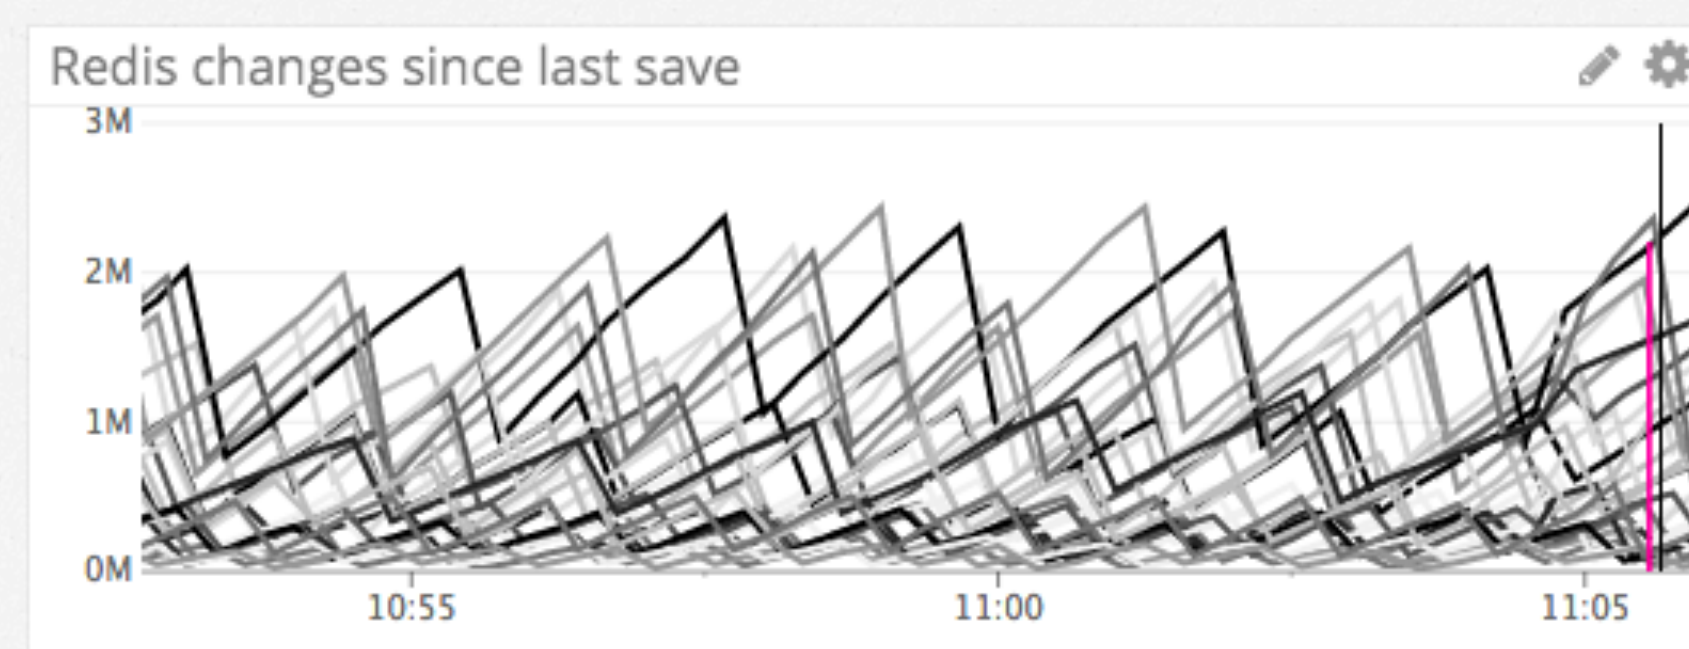


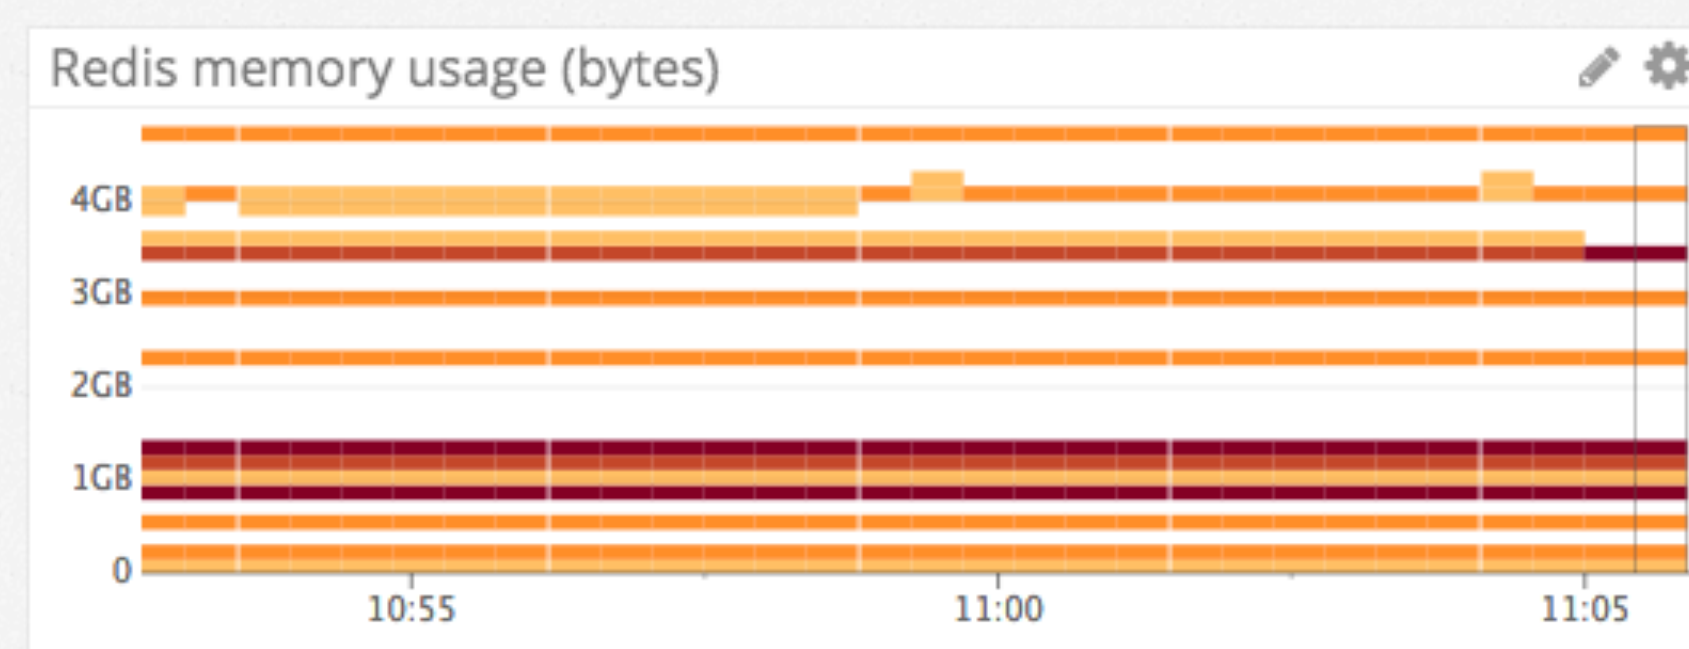


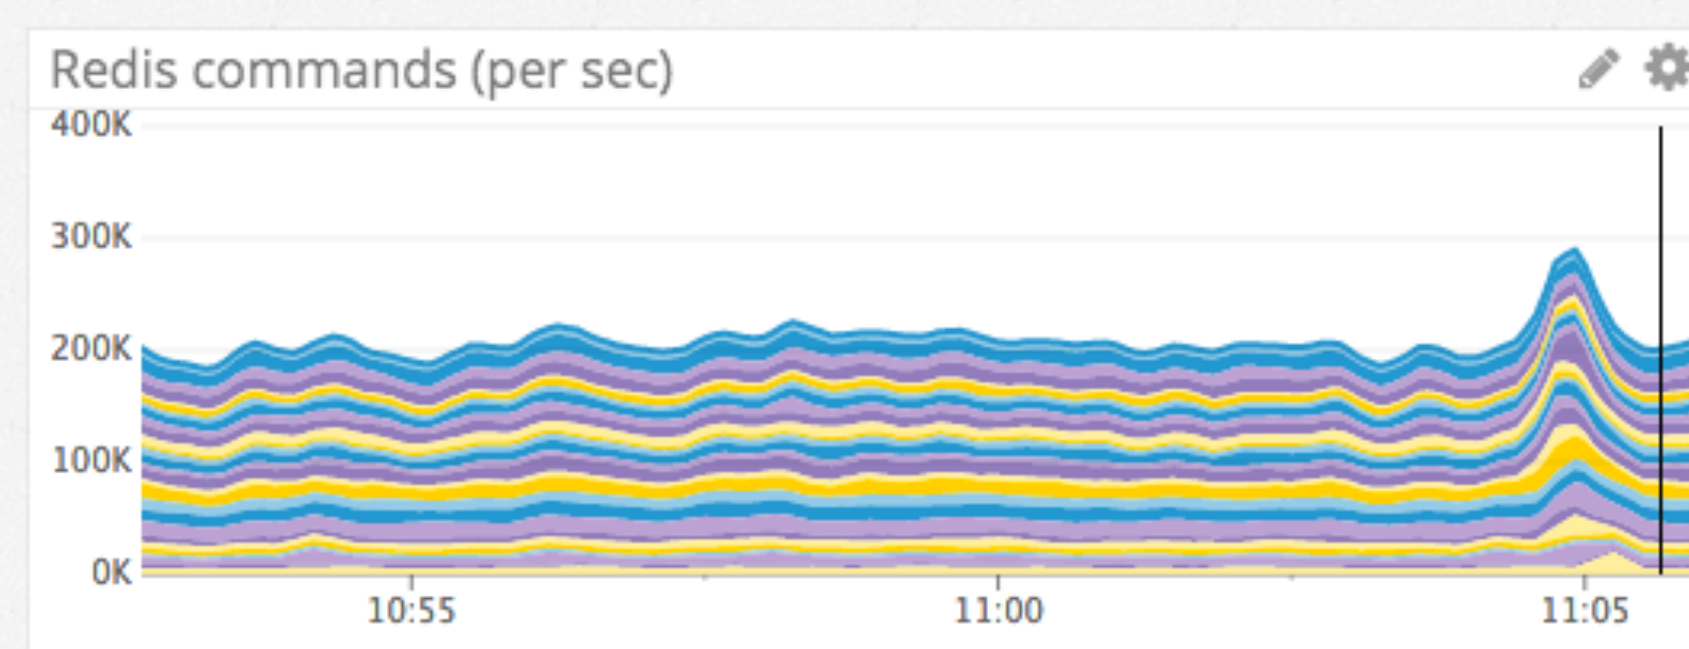

- **Cache Hit to Miss Ratio**
- **cmds / second**
- **Keys stored by host**
- **request latency**

Resource Metrics - Caching

- **evicted_keys**
- **used_memory**
- **network utilization**
- **connected_clients**
- **keyspace**
- **utilization on backend data store**

Redis Memory Used by Host





Events

- **Adding / Removing Nodes**
- **Incidents**
- **Code Deploys**
- **Config Changes**



Source: <http://bit.ly/1NoW6aj>

Resources

[Monitoring 101: Alerting](https://www.datadoghq.com/blog/monitoring-101-alerting/)

<https://www.datadoghq.com/blog/monitoring-101-alerting/>

[Monitoring 101: Collecting the Right Data](https://www.datadoghq.com/blog/monitoring-101-collecting-data/)

<https://www.datadoghq.com/blog/monitoring-101-collecting-data/>

[Monitoring 101: Investigating performance issues](https://www.datadoghq.com/blog/monitoring-101-investigation/)

<https://www.datadoghq.com/blog/monitoring-101-investigation/>

[The Power of Tagged Metrics](https://www.datadoghq.com/blog/the-power-of-tagged-metrics/)

<https://www.datadoghq.com/blog/the-power-of-tagged-metrics/>

[Monitoring Redis: Collecting Performance Metrics](https://www.datadoghq.com/blog/how-to-monitor-redis-performance-metrics/)

<https://www.datadoghq.com/blog/how-to-monitor-redis-performance-metrics/>

[HashRing](https://pypi.python.org/pypi/hash_ring)

https://pypi.python.org/pypi/hash_ring