

Utah Distributed Systems Meetup and Reading Group - Metrics/Tracing

JT Olds

Space Monkey
Vivint R&D

June 16 2015

Outline

- 1 Introduction
- 2 Metrics
- 3 Tracing

Outline

1 Introduction

2 Metrics

3 Tracing

Introduction

1 Introduction

The metrics system is Netflix's biggest cloud spend.
"We're a monitoring platform that happens to stream movies." -@royrapoport #monitorama

– @dberkholz

Why metrics or tracing?

- Very hard to get a distributed system right first try.
- Traditional debugging approaches don't scale.
- Not knowing what your system is doing is flying blind or stabbing in the dark.
- Mental models can be great but can also be wrong.

Why metrics or tracing?

- Very hard to get a distributed system right first try.
- Traditional debugging approaches don't scale.
- Not knowing what your system is doing is flying blind or stabbing in the dark.
- Mental models can be great but can also be wrong.

Why metrics or tracing?

- Very hard to get a distributed system right first try.
- Traditional debugging approaches don't scale.
- Not knowing what your system is doing is flying blind or stabbing in the dark.
- Mental models can be great but can also be wrong.

Why metrics or tracing?

- Very hard to get a distributed system right first try.
- Traditional debugging approaches don't scale.
- Not knowing what your system is doing is flying blind or stabbing in the dark.
- Mental models can be great but can also be wrong.

Outline

1 Introduction

2 Metrics

3 Tracing

Metrics

2 Metrics

- What?
- Instrumentation
- Graphite
- Graphite Replacements
- Graphite Extensions
- Graphite Helpers

Metrics

2 Metrics

■ What?

- Instrumentation
- Graphite
- Graphite Replacements
- Graphite Extensions
- Graphite Helpers

What do I mean by metrics?

- Any kind of data collection about your system
- ...but usually time-series key/value, where the value is some number.
- Sometimes called telemetry.

What do I mean by metrics?

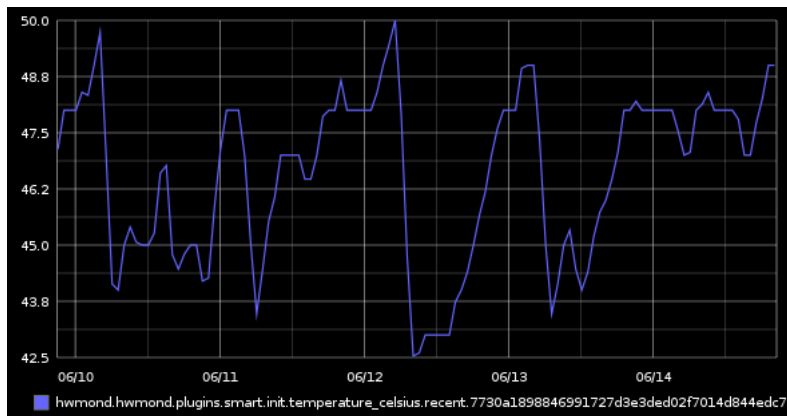
- Any kind of data collection about your system
- ...but usually time-series key/value, where the value is some number.
- Sometimes called telemetry.

What do I mean by metrics?

- Any kind of data collection about your system
- ...but usually time-series key/value, where the value is some number.
- Sometimes called telemetry.

Example: Graphite wire format

```
smart.temp_cels.7730a 48.0 1434314640  
smart.temp_cels.7730a 46.3 1434314848  
smart.temp_cels.7730a 46.8 1434315040
```

Metrics

2 Metrics

- What?
- **Instrumentation**
- Graphite
- Graphite Replacements
- Graphite Extensions
- Graphite Helpers

Libraries

- Lots of library choices
- Protocol so simple, you can write your own!
- (We did - `gopkg.in/spacemonkeygo/monitor.v1`)

Libraries

- Lots of library choices
- Protocol so simple, you can write your own!
- (We did - `gopkg.in/spacemonkeygo/monitor.v1`)

Libraries

- Lots of library choices
- Protocol so simple, you can write your own!
- (We did - `gopkg.in/spacemonkeygo/monitor.v1`)

Library commonalities

- Time measurement
- Event tracking
- Arbitrary value tracking
- Good ones let you measure quantiles

Library commonalities

- Time measurement
- Event tracking
- Arbitrary value tracking
- Good ones let you measure quantiles

Library commonalities

- Time measurement
- Event tracking
- Arbitrary value tracking
- Good ones let you measure quantiles

Library commonalities

- Time measurement
- Event tracking
- Arbitrary value tracking
- Good ones let you measure quantiles

- I'm mostly skipping how you get data cause there's so many ways.
- This means I'm also skipping monitoring systems, like Nagios or Zabbix or Ganglia or Munin or Monit or Zenoss or Cacti or Collectd or cAdvisor or Sensu or...

- I'm mostly skipping how you get data cause there's so many ways.
- This means I'm also skipping monitoring systems, like Nagios or Zabbix or Ganglia or Munin or Monit or Zenoss or Cacti or Collectd or cAdvisor or Sensu or...

Metrics

2 Metrics

- What?
- Instrumentation
- **Graphite**
- Graphite Replacements
- Graphite Extensions
- Graphite Helpers

The 800 pound Gorilla

Live demo

Metrics

2 Metrics

- What?
- Instrumentation
- Graphite
- **Graphite Replacements**
- Graphite Extensions
- Graphite Helpers


Full graphite replacements

- **InfluxDB** - <http://influxdb.com/>
- Prometheus - <http://prometheus.io/>
- OpenTSDB - <http://opentsdb.net/>
- Atlas - <https://github.com/Netflix/atlas>
- Blueflood <http://blueflood.io/>

Read Points

Query

data



time	sequence_number	bar	foo
1384185425284	29696	182	205
1384185424149	29695	234	293
1384185423055	29694	231	246
1384185421955	29693	236	130
1384185420678	29692	235	205

Data Interface

Read Points


Query

```
select * from "influxdb/memory/memory-used"
```

InfluxDB features a [SQL-like query language](#)

Execute Query

influxdb/memory/memory-used



time	sequence_number	type_instance	dsname	dstype	value	host	plugin	plugin_instance	type
1432501561000	173150001	used	value	gauge	154021888	influxdb	memory		memory
1432501551000	172870001	used	value	gauge	154034176	influxdb	memory		memory

Data Interface

Read Points


Query

```
select mean from ping.ams1.www_google_com.last_3h.group_1m where time > now()-5h
```

InfluxDB features a [SQL-like query language](#)

Execute Query

ping.ams1.www_google_com.last_3h.group_1m



Full graphite replacements

- **InfluxDB** - <http://influxdb.com/>
- **Prometheus** - <http://prometheus.io/>
- OpenTSDB - <http://opentsdb.net/>
- Atlas - <https://github.com/Netflix/atlas>
- Blueflood <http://blueflood.io/>

Prometheus

- Pull instead of push
- Alerting and precomputation rules built in

Prometheus


- Pull instead of push
- Alerting and precomputation rules built in

Prometheus Alerts Graph Status Help

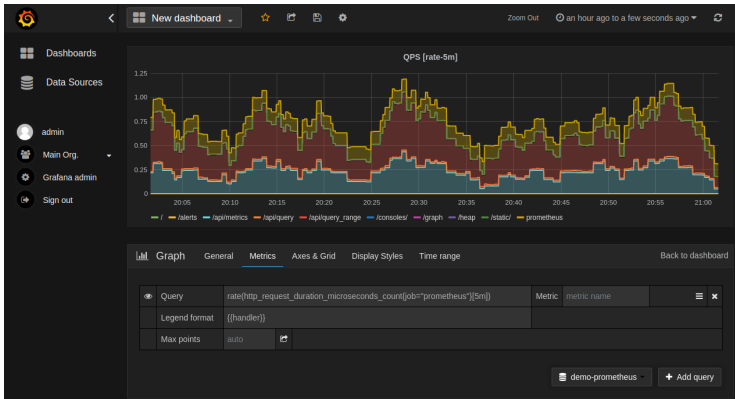
```
topk(3, sum(rate(bazooka_instance_cpu_time_ns[5m])) by (app, proc))
```

Execute

Graph Console

- 2h + ◀ Until ▶ Res. (s) stacked





Full graphite replacements

- InfluxDB - <http://influxdb.com/>
- Prometheus - <http://prometheus.io/>
- OpenTSDB - <http://opentsdb.net/>
- Atlas - <https://github.com/Netflix/atlas>
- Blueflood <http://blueflood.io/>

TSD

Time Series Database

Graph Stats Logs Version

From To [now](#) Autoreload WxH:

Metric: Rate Rate Ctr Right Axis

Tags Rate Ctr Max:

Rate Ctr Reset:

Aggregator: Downsample

Label

Format

Range

Log scale

Please specify a metric:

- tcollector.collector.lines_invalid
- tcollector.collector.lines_received
- tcollector.collector.lines_sent
- tcollector.reader.lines_collected
- tcollector.reader.lines_dropped
- tsd.compaction.count
- tsd.compaction.deletes
- tsd.compaction.errors
- tsd.compaction.queue.size
- tsd.compaction.writes
- tsd.connectionmgr.connections
- tsd.connectionmgr.exceptions
- tsd.hbase.connections.created
- tsd.hbase.flushes
- tsd.hbase.latency_50pct

TSD

 Time Series Database[Graph](#) [Stats](#) [Logs](#) [Version](#)

From To [\(now\)](#) Autoreload WxH:

Every: seconds

[client_conn](#) [1min](#) [cache_hit](#) +


Metric: Rate Right Axis

Tags Aggregator:

Downsample

Smooth

4853 points retrieved, 841 points plotted in 37ms.



TSD

Time Series Database

Graph Stats Logs Version

From To Autoreload WxH:

transfer_time +

Metric: Rate Rate Ctr Right Axis
 Rate Ctr Max:
 Rate Ctr Reset:


Tags X

Aggregator: Downsample

Key location:

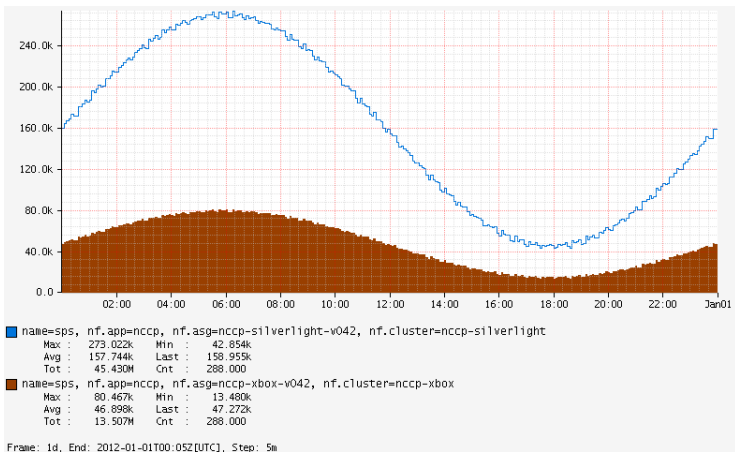
 Horizontal layout
 Box
 No key (overrides others)

5656 points retrieved, 655 points plotted in 46ms.

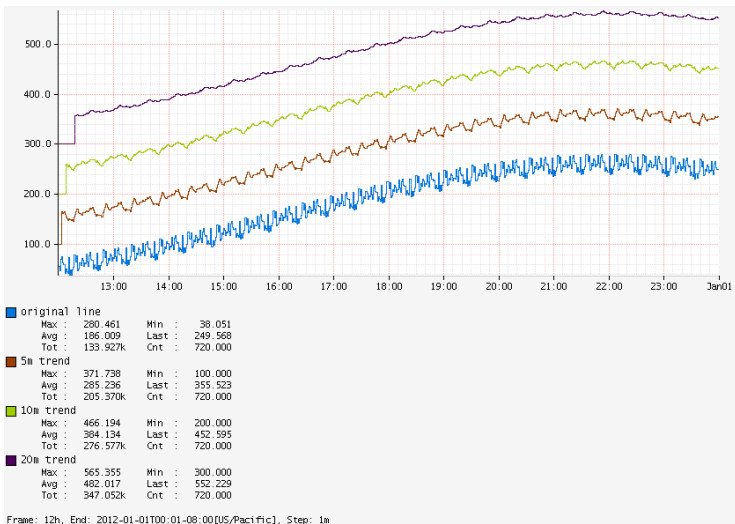


Full graphite replacements

- **InfluxDB** - <http://influxdb.com/>
- **Prometheus** - <http://prometheus.io/>
- **OpenTSDB** - <http://opentsdb.net/>
- **Atlas** - <https://github.com/Netflix/atlas>
- **Blueflood** <http://blueflood.io/>







Full graphite replacements

- **InfluxDB** - <http://influxdb.com/>
- **Prometheus** - <http://prometheus.io/>
- **OpenTSDB** - <http://opentsdb.net/>
- **Atlas** - <https://github.com/Netflix/atlas>
- **Blueflood** <http://blueflood.io/>

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- **Instrumental** - <https://instrumentalapp.com/>
- **Datadog** - <https://www.datadoghq.com/>
- etc.

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- **Instrumental** - <https://instrumentalapp.com/>
- **Datadog** - <https://www.datadoghq.com/>
- etc.

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- Instrumental - <https://instrumentalapp.com/>
- Datadog - <https://www.datadoghq.com/>
- etc.

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- **Instrumental** - <https://instrumentalapp.com/>
- **Datadog** - <https://www.datadoghq.com/>
- etc.

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- **Instrumental** - <https://instrumentalapp.com/>
- **Datadog** - <https://www.datadoghq.com/>
- etc.

Hosted services

- **Stathat** - <http://www.stathat.com/>
- **Librato** - <https://www.librato.com/>
- **Hosted Graphite** -
<https://www.hostedgraphite.com/>
- **Instrumental** - <https://instrumentalapp.com/>
- **Datadog** - <https://www.datadoghq.com/>
- etc.

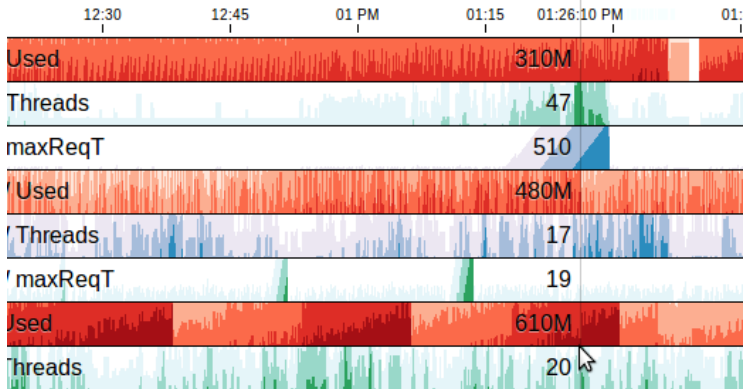
Metrics

2 Metrics

- What?
- Instrumentation
- Graphite
- Graphite Replacements
- **Graphite Extensions**
- Graphite Helpers

Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>

Descartes Home **Graphs** Dashboards Metrics

Graphs + search for "cpu"

cpu

Import Graphs

Add to Dashboard

Interval

1h

3h

6h

24h

Columns

1

2

3

4


🔄

🗑️

switch.cpu



redis2.cpu



collectd.cpu_by_processes



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>

My Dashboard

Dashboards ▾

Data Browser

Email Metrics

Hourly metrics for the email system



UPDATED FRI OCT 07 21:41:43 +0100 2011

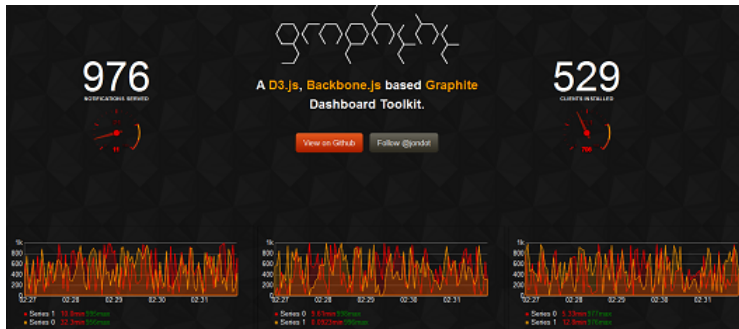
Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



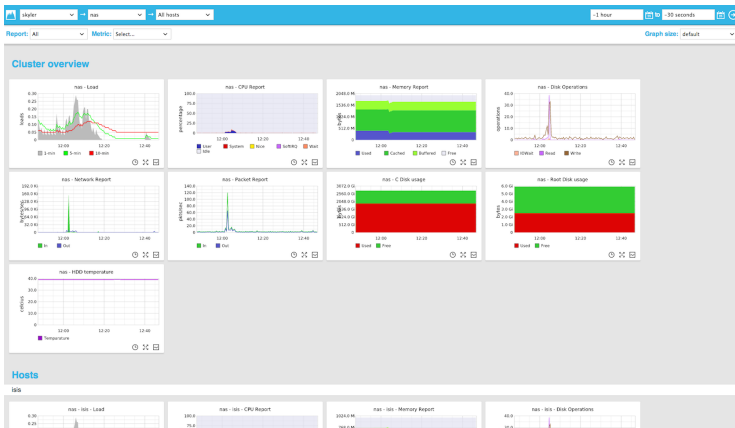
Different frontends

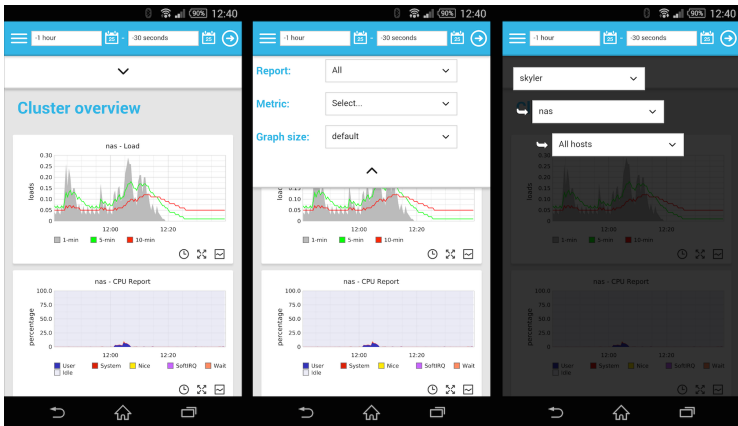
- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
[-https://vimeo.github.io/graph-explorer/](https://vimeo.github.io/graph-explorer/)
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>





Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



Different frontends


- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>

Graphiti

Dashboards

Graphs

+ New Graph



Graph

Update

Save New

Dashboards

Add to Mikes ↓

Save

Create New

Options

Metrics

Editor

```
1 {
2   "options": {
3     "title": "Resque Job Complete",
4     "from": "-5d",
5     "until": "-1d",
6     "bgcolor": "white",
7     "fgcolor": "black",
8     "thickness": 2
9   }
}
```

Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>

- **Tasseo** -

Leonardo DB Cluster - Network - System - Dashboards - Release - Site Performances -







Dashboard...



Server-1 System Metrics


30 minutes 2 hours 8 hours 24 hours 1 week 1 month 1 year Select Date

+ - reset



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>



Different frontends

- **Cubism.js** - <https://square.github.io/cubism/>
- **Dashing** - <https://shopify.github.io/dashing/>
- **Descartes** -
<https://github.com/obfuscurity/descartes>
- **Gdash** - <https://github.com/ripienaar/gdash>
- **Grafana** - <http://grafana.org/>
- **Graphene** - <https://jondot.github.io/graphene/>
- **Graphsky** - <https://github.com/TMG-nl/graphsky>
- **Graph-explorer**
-<https://vimeo.github.io/graph-explorer/>
- **Graphiti** -
<https://github.com/paperlesspost/graphiti>
- **Leonardo** -
<https://github.com/PrFalken/leonardo>
- **Tasseo** -
<https://github.com/obfuscurity/tasseo>


Infrastructure Carbon

02:35 PM
SUNDAY 22 JUNE 2014


Past Day



Metrics Received



Metrics Created

TOTAL METRICS
RECEIVED

1.85G

PEAK METRICS
RECEIVED

5.6k /sec

NEW METRICS
CREATED

980

POINTS
COMMITTED

1.85G

TOTAL QUERIES
EXECUTED


1.35M

MAXIMUM QUERY
RATE


50 /sec

Agents


Memory Usage



CPU Usage



Points per Update



Tosca v0.1.0

Different backends

- **Whisper → Ceres**
- Cyanite (Cassandra-backed)
- KairosDB (Cassandra-backed, OpenTSDB rewrite)
- level-tds (LevelDB-backed)
- Blueflood (Cassandra-backed)

Different backends

- Whisper → Ceres
- Cyanite (Cassandra-backed)
- KairosDB (Cassandra-backed, OpenTSDB rewrite)
- level-tds (LevelDB-backed)
- Blueflood (Cassandra-backed)

Different backends

- Whisper → Ceres
- Cyanite (Cassandra-backed)
- KairosDB (Cassandra-backed, OpenTSDB rewrite)
- level-tds (LevelDB-backed)
- Blueflood (Cassandra-backed)

Different backends

- Whisper → Ceres
- Cyanite (Cassandra-backed)
- KairosDB (Cassandra-backed, OpenTSDB rewrite)
- level-tds (LevelDB-backed)
- Blueflood (Cassandra-backed)

Different backends

- Whisper → Ceres
- Cyanite (Cassandra-backed)
- KairosDB (Cassandra-backed, OpenTSDB rewrite)
- level-tds (LevelDB-backed)
- Blueflood (Cassandra-backed)

Metrics

2 Metrics

- What?
- Instrumentation
- Graphite
- Graphite Replacements
- Graphite Extensions
- **Graphite Helpers**

Collection services

- **StatsD** - <https://github.com/etsy/statsd/>
- **Brubeck** - <https://github.com/github/brubeck/>
- **Heka** - <https://hekad.readthedocs.org/>

Collection services

- **StatsD** - <https://github.com/etsy/statsd/>
- **Brubeck** - <https://github.com/github/brubeck/>
- **Heka** - <https://hekad.readthedocs.org/>

Collection services

- **StatsD** - <https://github.com/etsy/statsd/>
- **Brubeck** - <https://github.com/github/brubeck/>
- **Heka** - <https://hekad.readthedocs.org/>

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>


skyline

stats.browse_pages.filters.miss.ship_to

red line shows detected anomalous datapoint

holy anomaly!

1 hour



24 hours:



metric name

anomalous datapoint

stats.browse_categories.feed ↗	timeseries datapoint that triggered detection	30
stats.browse_pages.cache.miss ↗		36
stats.browse_pages.filters.miss.ship_to ↗		30
stats.region.suggested.code.KG ↗		0
stats.shipping_labels.api.multi_labels.subsequent_package_purchase_duration.3 ↗		0
stats.shipping_labels.api.multi_labels.subsequent_package_purchase_duration.5 ↗		0
stats.shopstats.domain.ShopStatsRollup.etsy ↗		42
stats.shopstats.domain.ShopStatsRollup.etsy.activity ↗		170
stats.timers.api.findAllReceiptTransactions.200.lower ↗		166
stats.timers.api.findAllReceiptTransactions.200.mean ↗		166
stats.timers.api.findAllReceiptTransactions.200.mean_90 ↗		166
stats.timers.api.findAllReceiptTransactions.200.median ↗		166
stats.timers.api.findAllReceiptTransactions.200.lower ↗		166

list of anomalous metrics

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>

Tools to help analyze

- **Druid** - <http://druid.io/>
- **Kale (Skyline/Occulus)** - <https://codeascraft.com/2013/06/11/introducing-kale/>
- **Bosun** - <https://bosun.org/>
- **Heka** - <https://hekad.readthedocs.org/>
- **Riemann** - <http://riemann.io/>
- **Rocksteady** - <https://code.google.com/p/rocksteady/>

Etc

`https://graphite.readthedocs.org/en/latest/
tools.html`

Outline

1 Introduction

2 Metrics

3 Tracing

Tracing

- 3 Tracing
 - What?
 - Services

Tracing

- 3 Tracing
 - What?
 - Services

What is distributed tracing?

- Basically, stack traces across services from production.
- Instrument your RPC code to automatically pass current context information through to other calls.
- `trace_id`, `span_id`, `parent_span_id`, `sampled`, `debug`
- Annotations - `start`, `end`, `errored`, `URL`, `Status Code`, `others`

What is distributed tracing?

- Basically, stack traces across services from production.
- Instrument your RPC code to automatically pass current context information through to other calls.
- `trace_id`, `span_id`, `parent_span_id`, `sampled`, `debug`
- Annotations - `start`, `end`, `errored`, `URL`, `Status Code`, `others`

What is distributed tracing?

- Basically, stack traces across services from production.
- Instrument your RPC code to automatically pass current context information through to other calls.
- `trace_id`, `span_id`, `parent_span_id`, `sampled`, `debug`
- Annotations - `start`, `end`, `errored`, `URL`, `Status Code`, `others`

What is distributed tracing?

- Basically, stack traces across services from production.
- Instrument your RPC code to automatically pass current context information through to other calls.
- `trace_id`, `span_id`, `parent_span_id`, `sampled`, `debug`
- **Annotations** - `start`, `end`, `errored`, **URL**, **Status Code**, others

Tracing

- 3** Tracing
 - What?
 - Services

Tracing services

- Dapper - research.google.com/pubs/pub36356.html
- Zipkin - <https://twitter.github.io/zipkin/>
- Phosphor - <https://github.com/mattheath/phosphor>

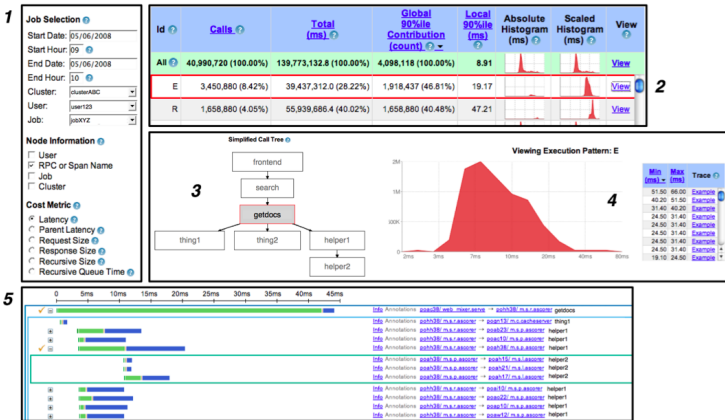
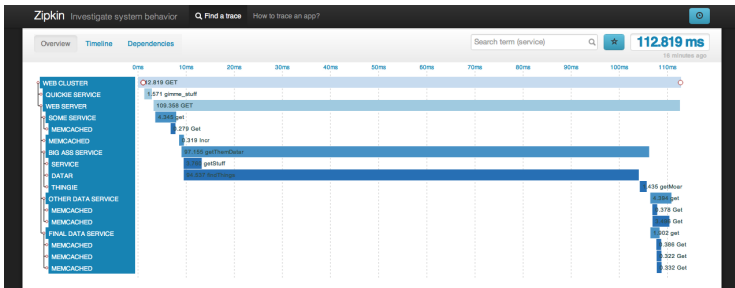
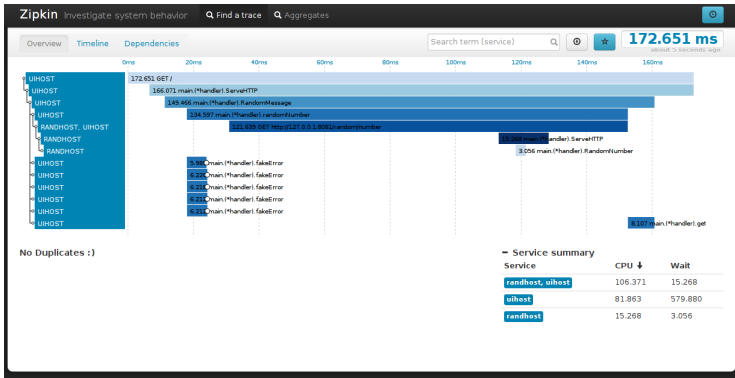


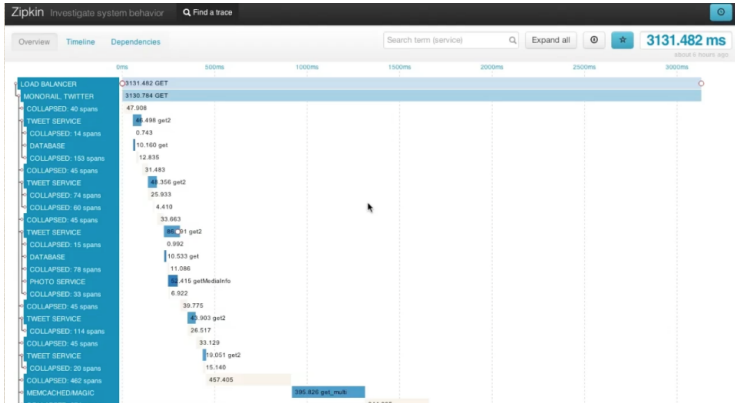
Figure 6: A typical user workflow in the general-purpose Dapper user interface.

Tracing services

- Dapper -
`research.google.com/pubs/pub36356.html`
- Zipkin - `https://twitter.github.io/zipkin/`
- Phosphor -
`https://github.com/mattheath/phosphor`







Tracing services

- Dapper -
`research.google.com/pubs/pub36356.html`
- Zipkin - `https://twitter.github.io/zipkin/`
- Phosphor -
`https://github.com/mattheath/phosphor`

Space MONKEY

Space Monkey!

- Distributed Hash Tables
- Consensus algorithms
- Reed Solomon
- Monitoring and sooo much data
- Security and cryptography engineering

Space Monkey!

Come work with us!