

Distributed Metrics and Log Aggregation

Alexander Heusingfeld & Tammo van Lessen



Monolith

Clarendon Boulevard
Rosslyn (VA) 2011



MICROSERVICES...


MICROSERVICES EVERYWHERE

memegenerator.net

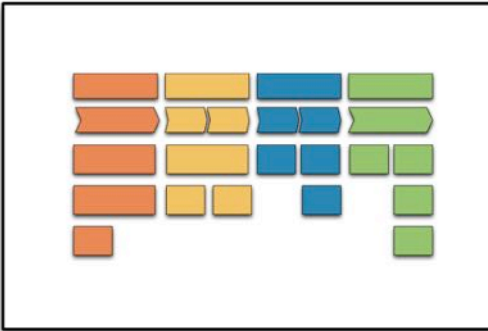
...so you start to disassemble your monoliths...



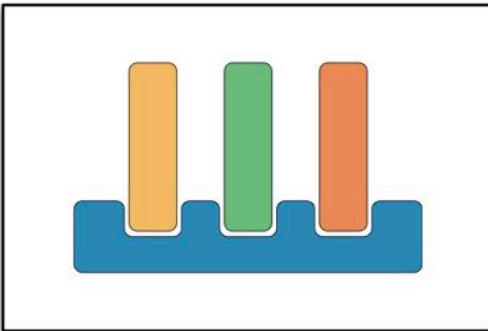
...and then you think about the infrastructure



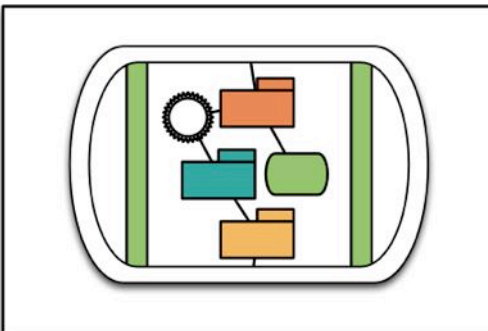
Architectural Decisions



> Domain architecture



> Macro architecture



> Micro architecture

Scenario: Big Shop



That wouldn't have happened with proper logging!



What makes good logging?

- What identifies a good log message?
- Which log level should I use when?
- Should I log into files? What format?

Some recommendations

- › Log messages should have a uniform style.
- › Log violations of assumptions.
- › Use markers to make log streams filterable.
- › Prefer machine-readable log formats over human-readable.
- › Identify correlation tokens and attach them to the log event.
- › Collect and store logs in a central repository.

Default Levels

Files? Warn only.

Logstash & Co? Info.

Magic bugs + advanced setup? Debug, or even trace.



- > Async Appenders (LMAX, MemoryMappedFileAppender)
- > Routing
- > Properties
- > Reconfiguration (Auto load, JMX,...)
- > Audit logs
- > Markers / Log levels
- > ...

Thread Context

```
ThreadContext.put("loginId", login);  
logger.error("Something bad happened!");  
ThreadContext.clear();
```

+ Layout:

```
%-5p: [%X{loginId}] %m%n
```

Log:

```
ERROR: [John Doe] Something bad happened!
```

Thread Context (2)

```
ThreadContext.put("loginId", login);  
logger.error("Something bad happened!");  
ThreadContext.clear();
```

+ JSON Layout:

Log:

```
{  
  "@version" => "1",  
  "@timestamp" => "2014-04-29T14:21:14.988-07:00",  
  "logger" => "com.example.LogStashExampleTest",  
  "level" => "ERROR",  
  "thread" => "Test worker",  
  "message" => "Something bad happened!",  
  "Properties" => {  
    "loginId" => "John Doe"  
  }  
}
```


Requirements in a distributed environment

- Aggregate logs in different formats from different systems.
- Search & Correlate
- Visualize
- Alert on complex correlations.



Logstash Architecture

inputs	codecs	filters	outputs
<ul style="list-style-type: none">• collectd• drupal_dblog• elasticsearch• eventlog• exec• file• ganglia• gelf• gemfire• generator• graphite• heroku• imap• invalid_input• irc• jmx• log4j• lumberjack• pipe• puppet_facter• rabbitmq• rackspace	<ul style="list-style-type: none">• cloudtrail• collectd• compress_spooler• dots• edn• edn_lines• fluent• graphite• json• json_lines• json_spooler• line• msgpack• multiline• netflow• noop• oldlogstashjson• plain• rubydebug• spool	<ul style="list-style-type: none">• advisor• alter• anonymize• checksum• cidr• cipher• clone• collate• csv• date• dns• drop• elapsed• elasticsearch• environment• extractnumbers• fingerprint• gelfify• geoip• grep• grok• grokdiscovery	<ul style="list-style-type: none">• boundary• circonus• cloudwatch• csv• datadog• datadog_metrics• elasticsearch• elasticsearch_http• <u>elasticsearch_river</u>• email• exec• file• ganglia• gelf• gemfire• google_bigquery• google_cloud_storage• graphite• graphtastic• hipchat• http• irc

Logstash – Hands on!

A Logstash Cluster

Legend:



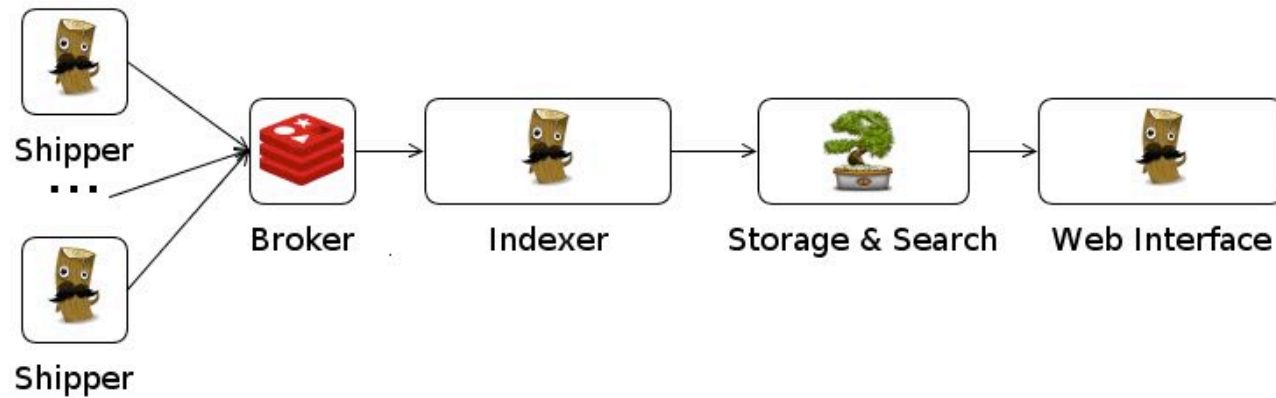
Logstash



Redis



ElasticSearch



... and there are others, too!

Apache Flume (ASL 2.0)

FluentD (ASL 2.0)

Graylog 2 (GPL)

Loggly (commerical)

Splunk (commerical)

A large iceberg floats in dark water, with its jagged, textured surface visible above the surface. The image is overlaid with a semi-transparent blue filter. The text is centered on the iceberg's surface.

Logging is cool.

And I can use it to collect metrics as well, right?

Yes, you can!

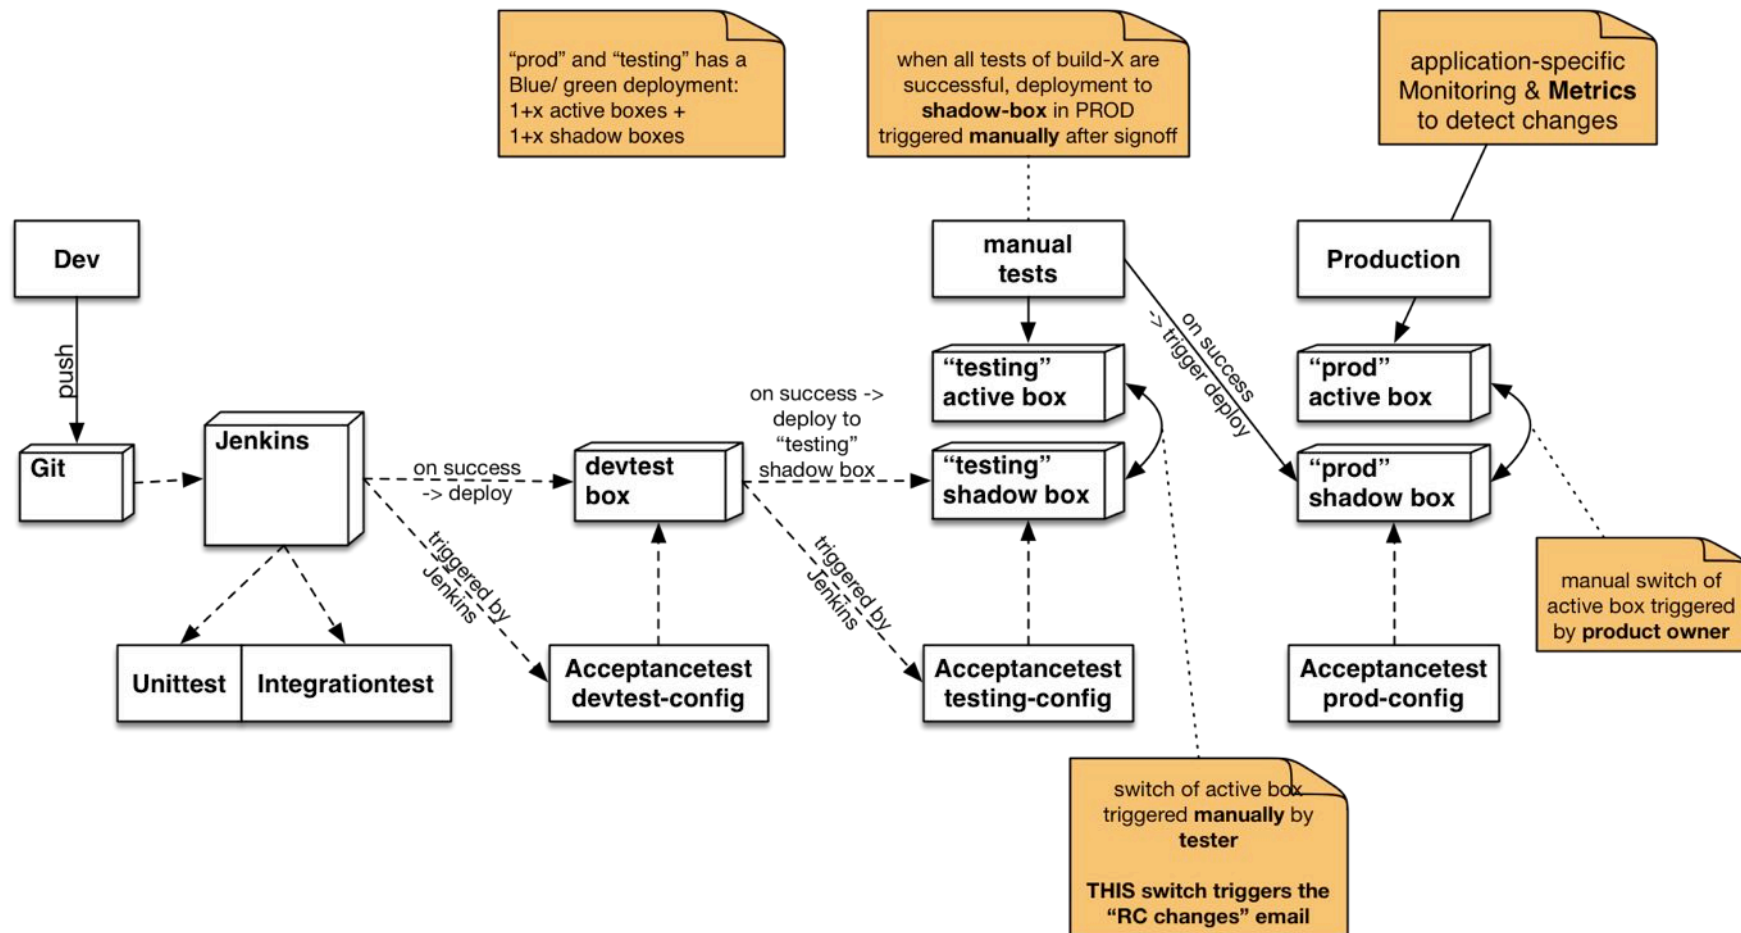
But you shouldn't!

Metrics

- Business Metrics
- Application Metrics
- System Metrics

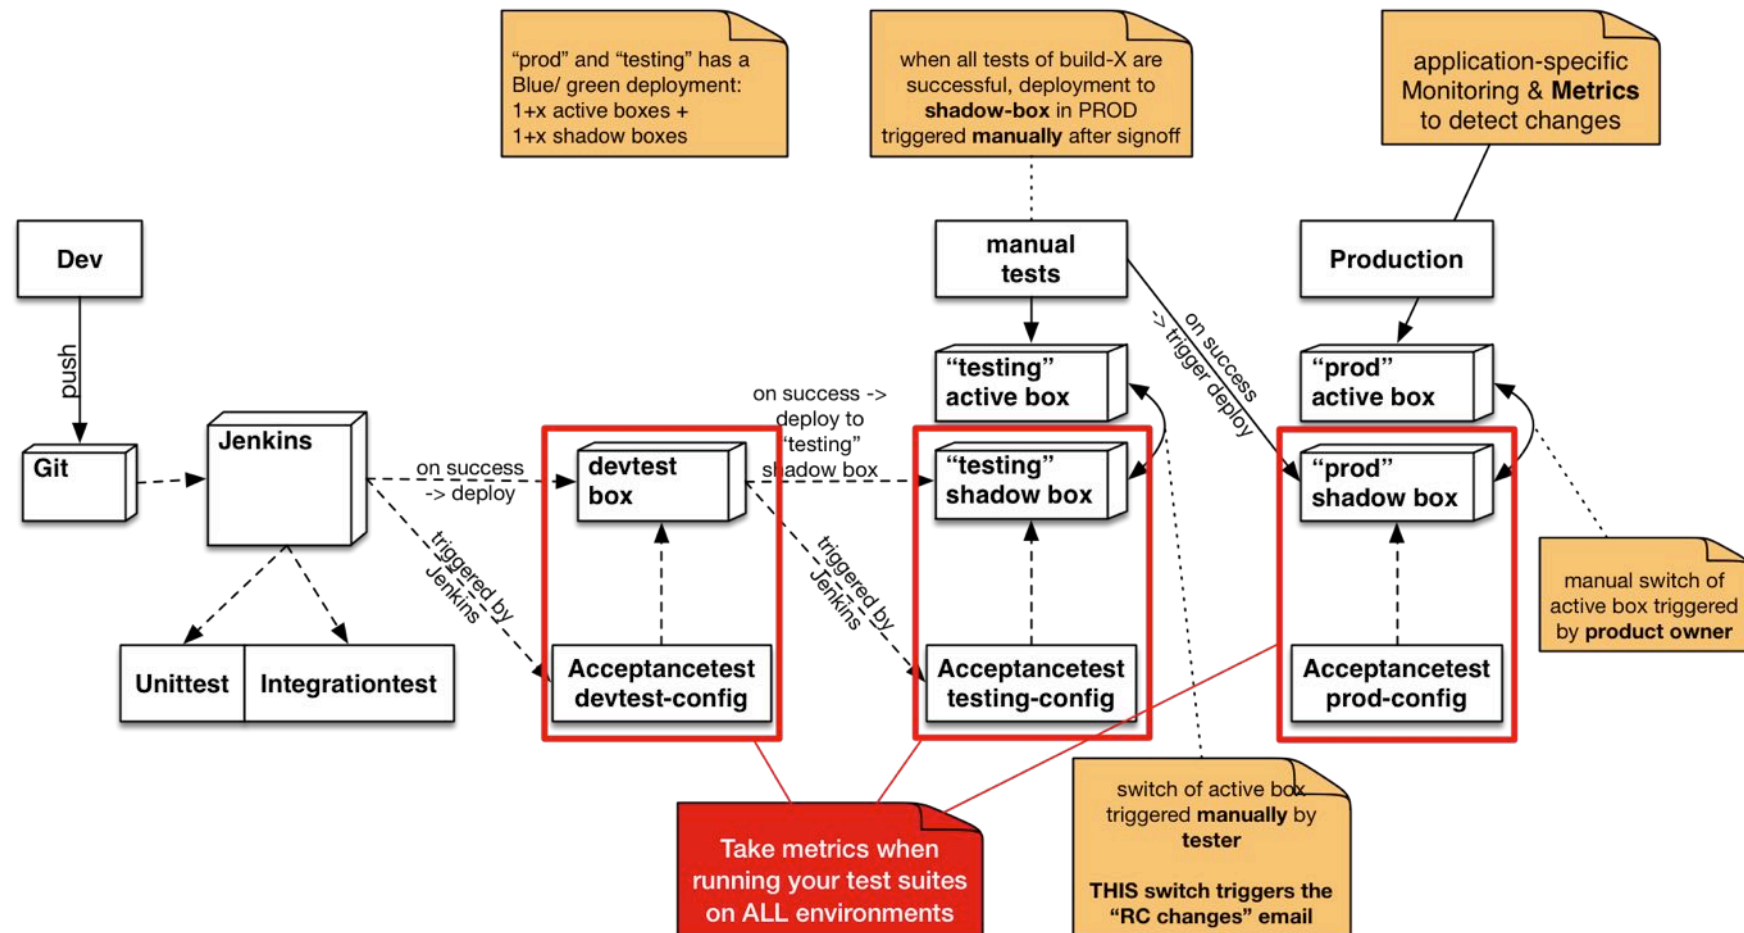
Continuous Delivery & Metrics?

Sample of a deployment-pipeline




Continuous Delivery & Metrics?

Sample of a deployment-pipeline



Gauges

An instrument that measures a value.



PANEL
LIGHT

Counters

A counter is a simple incrementing and decrementing integer.




Meters

A meter measures the rate at which a set of events occur.

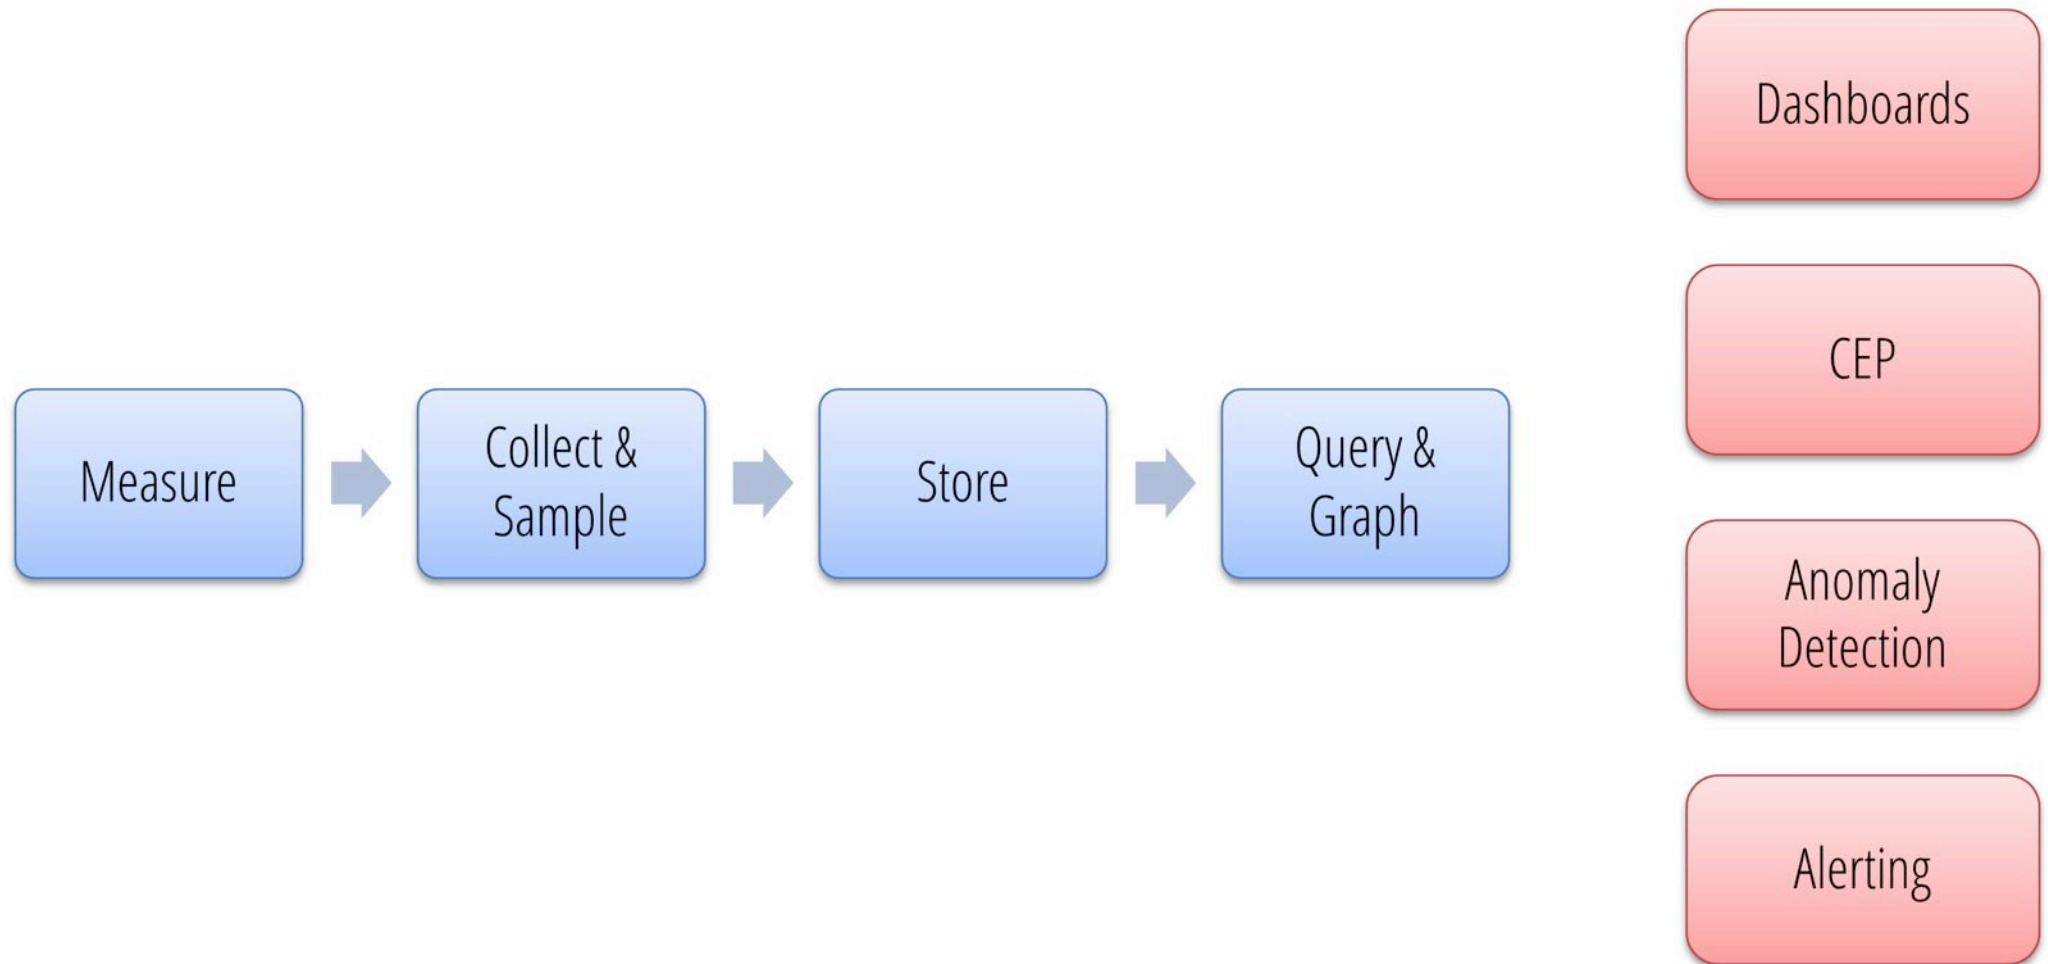
Histograms



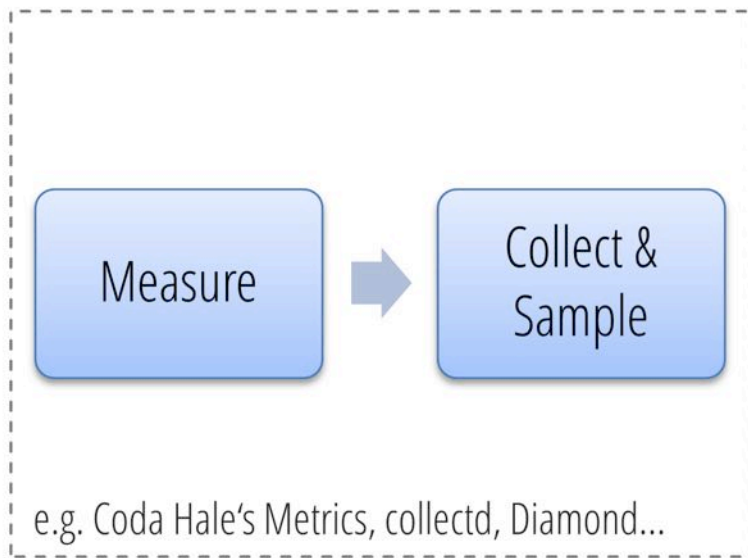
A Histogram measures the distribution of values.

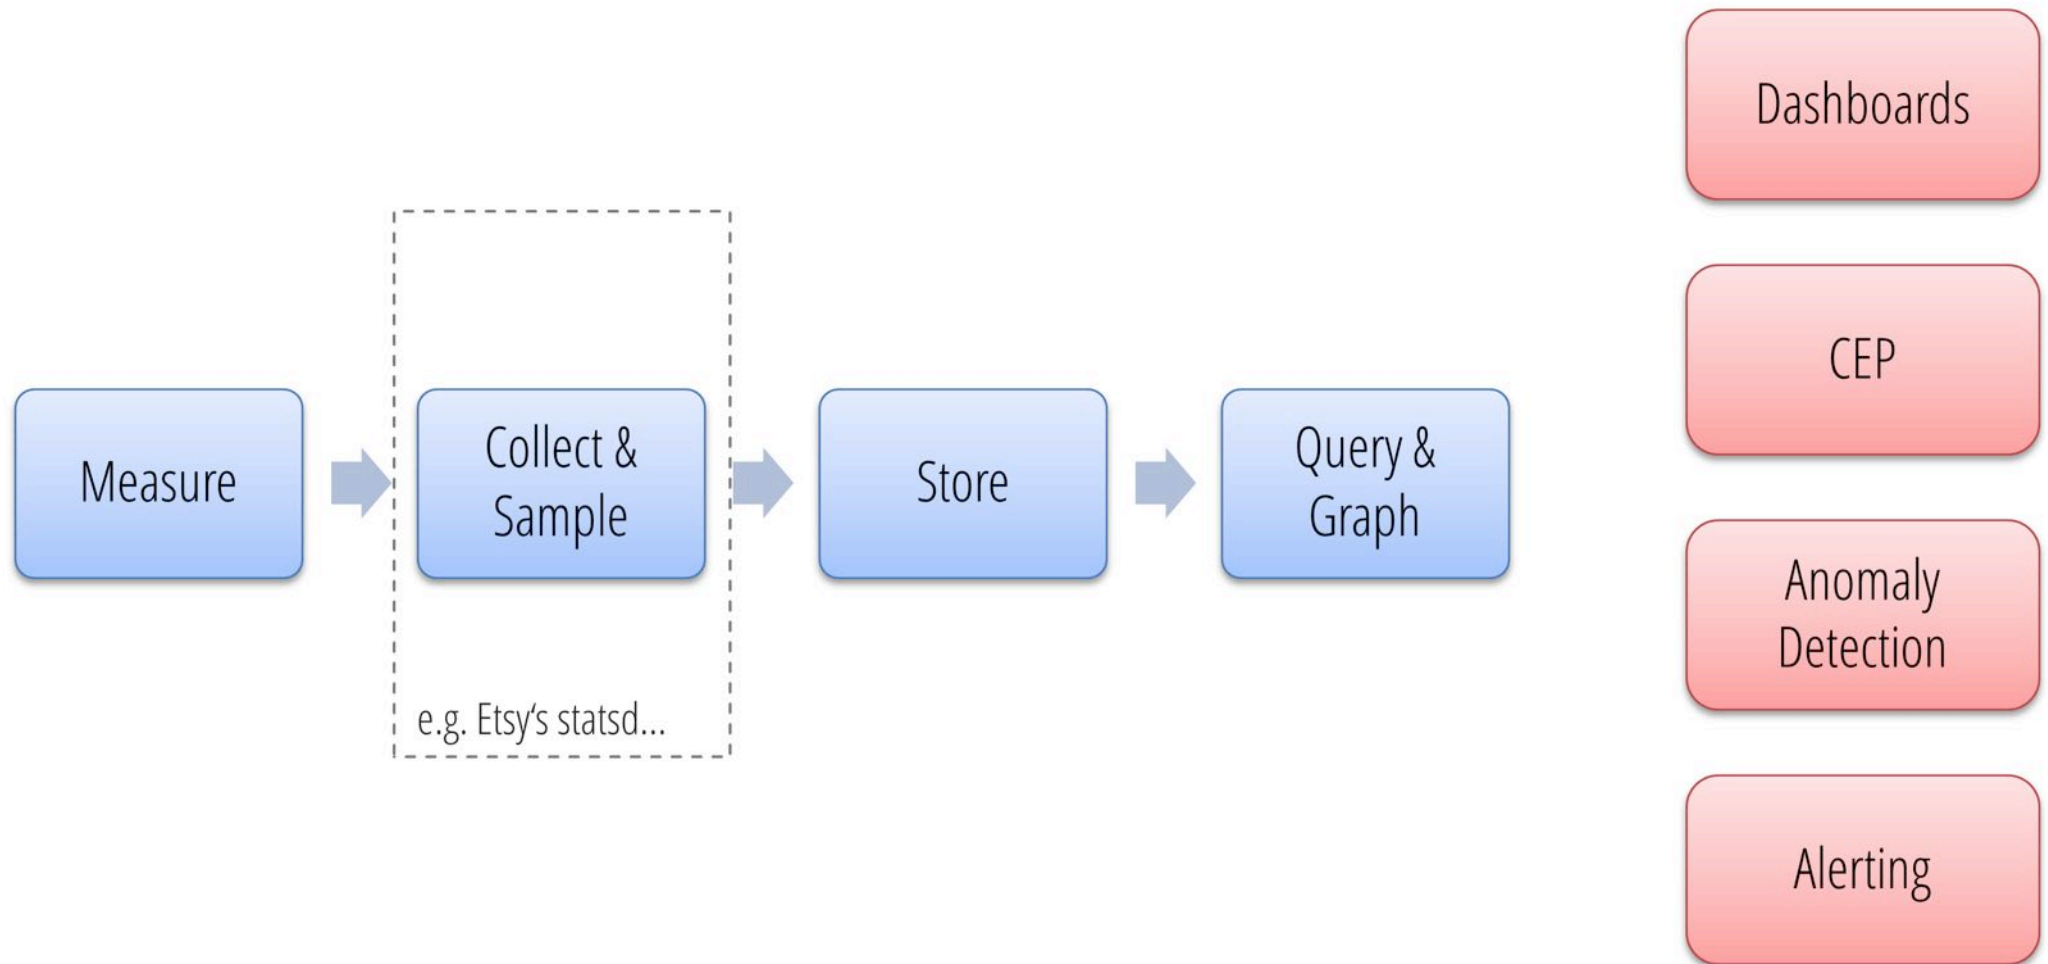


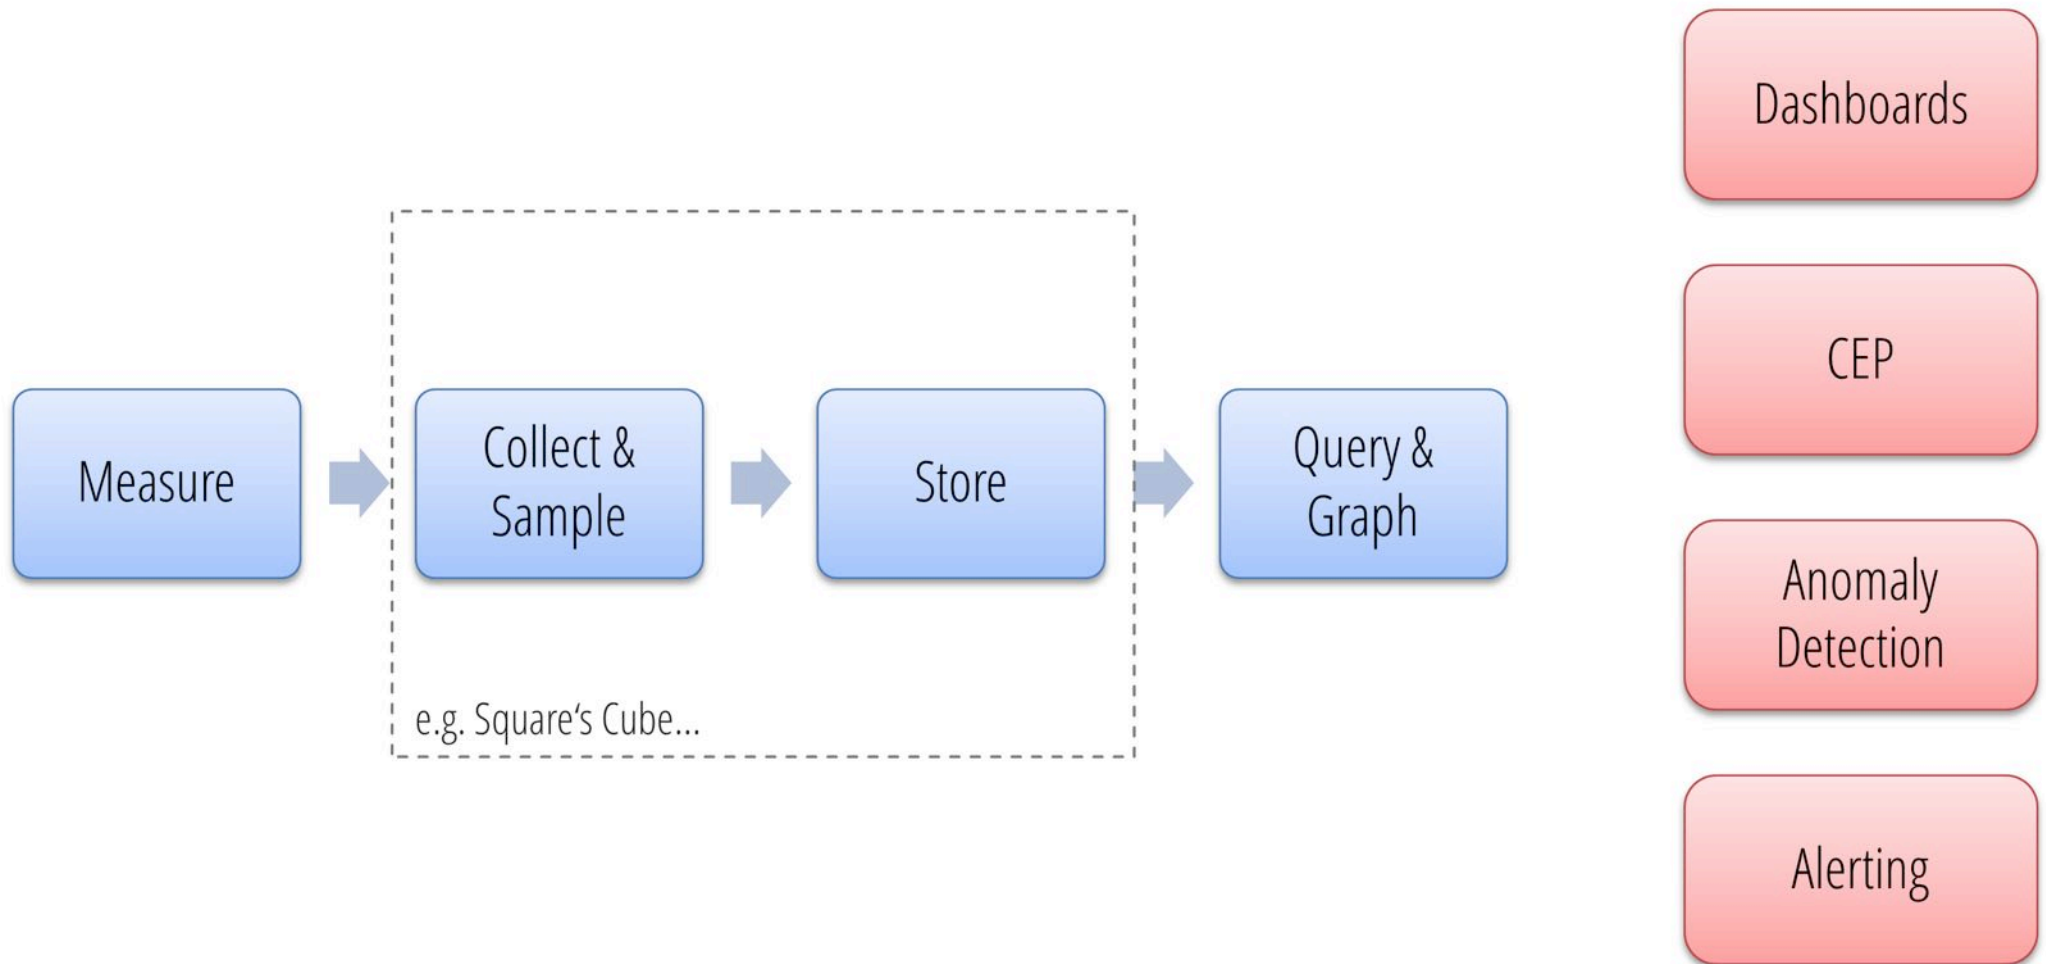
Timers

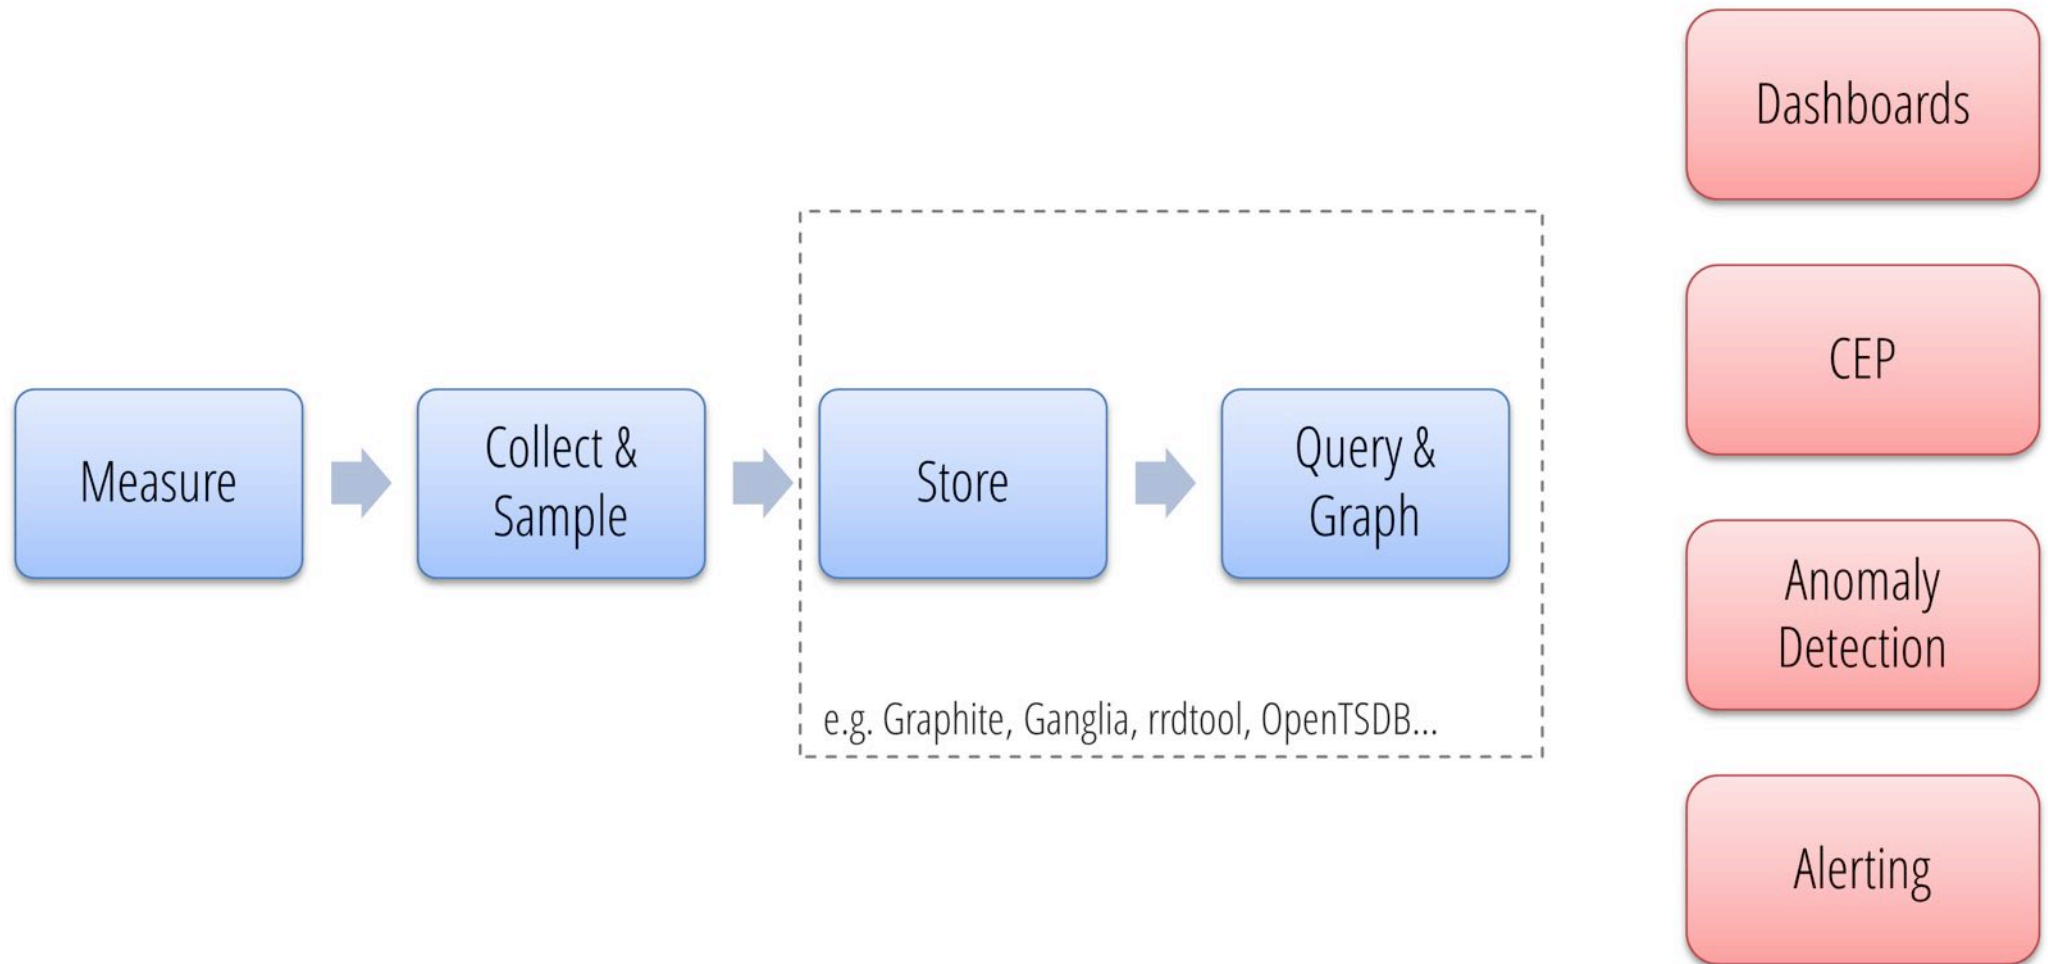


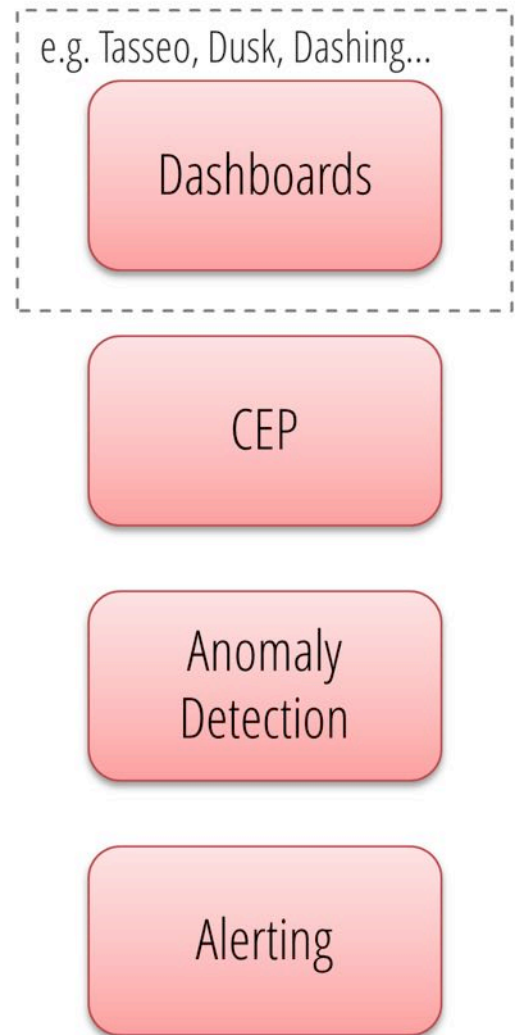



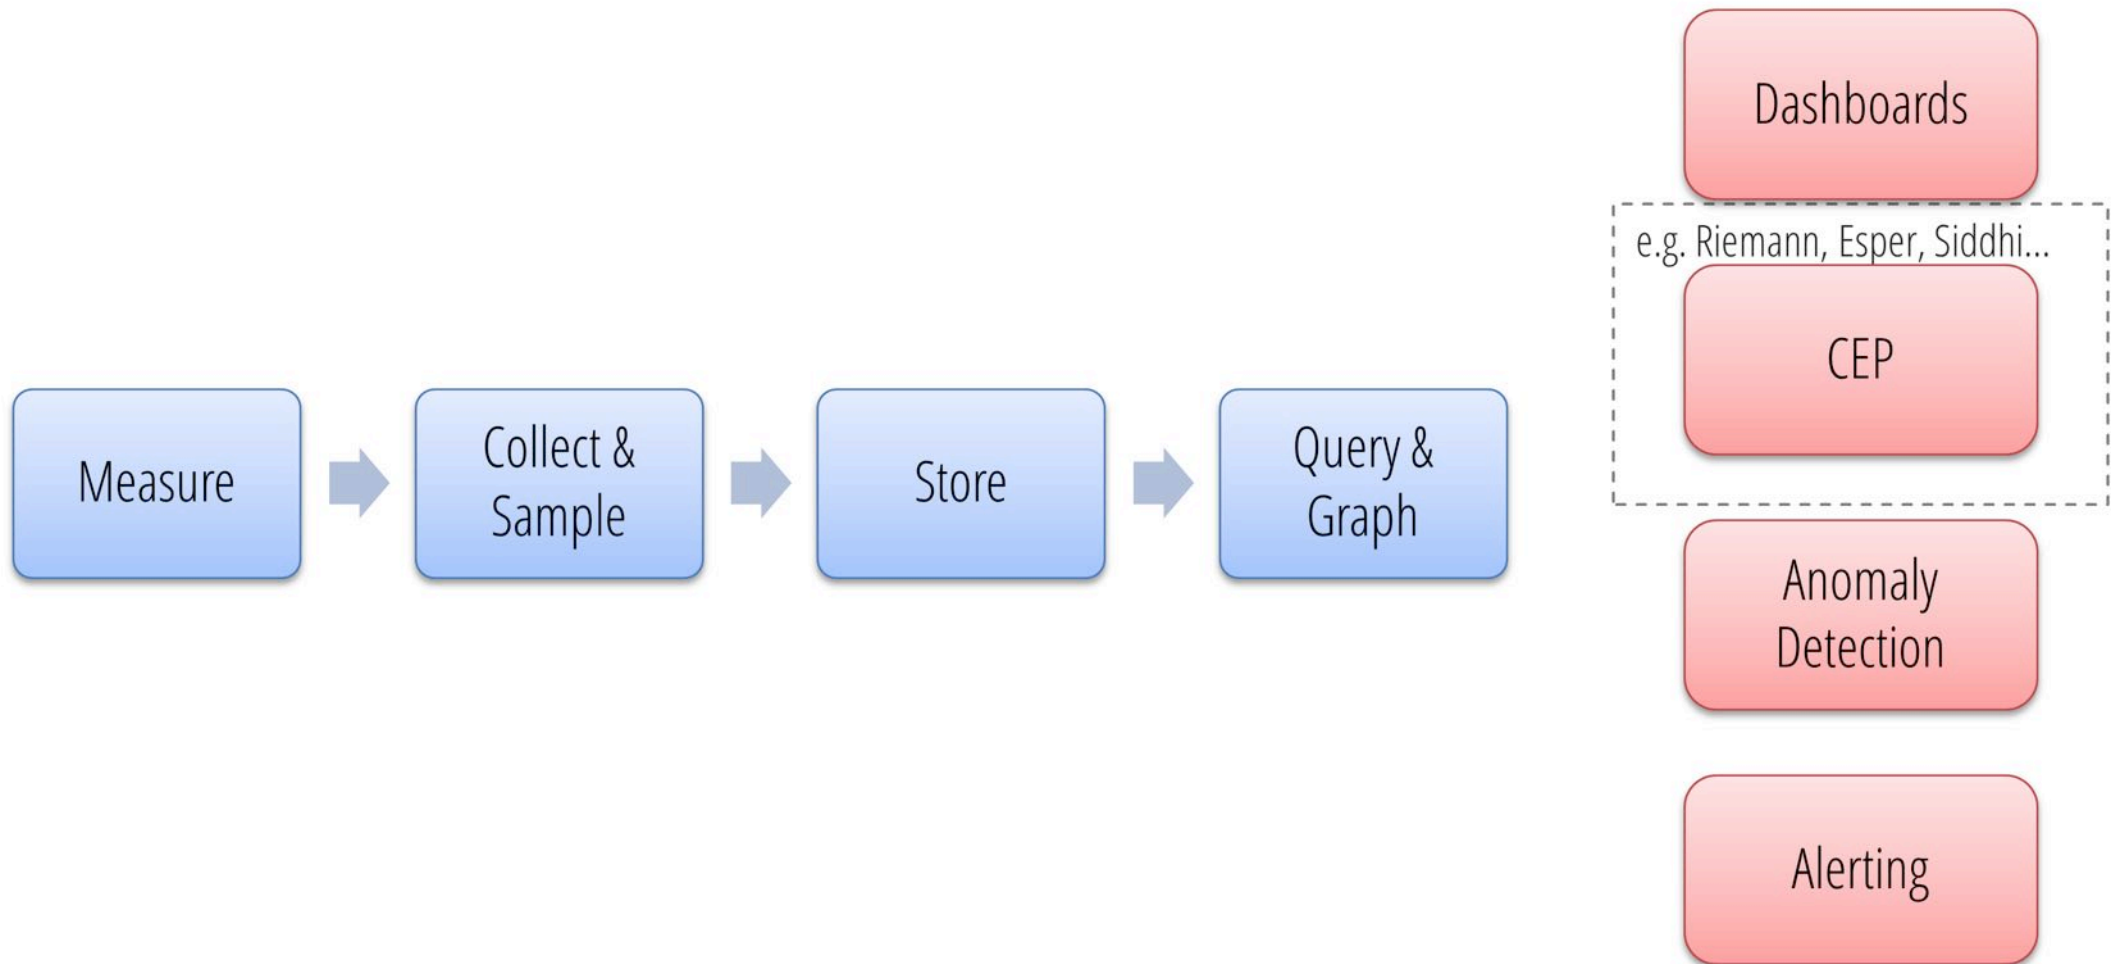


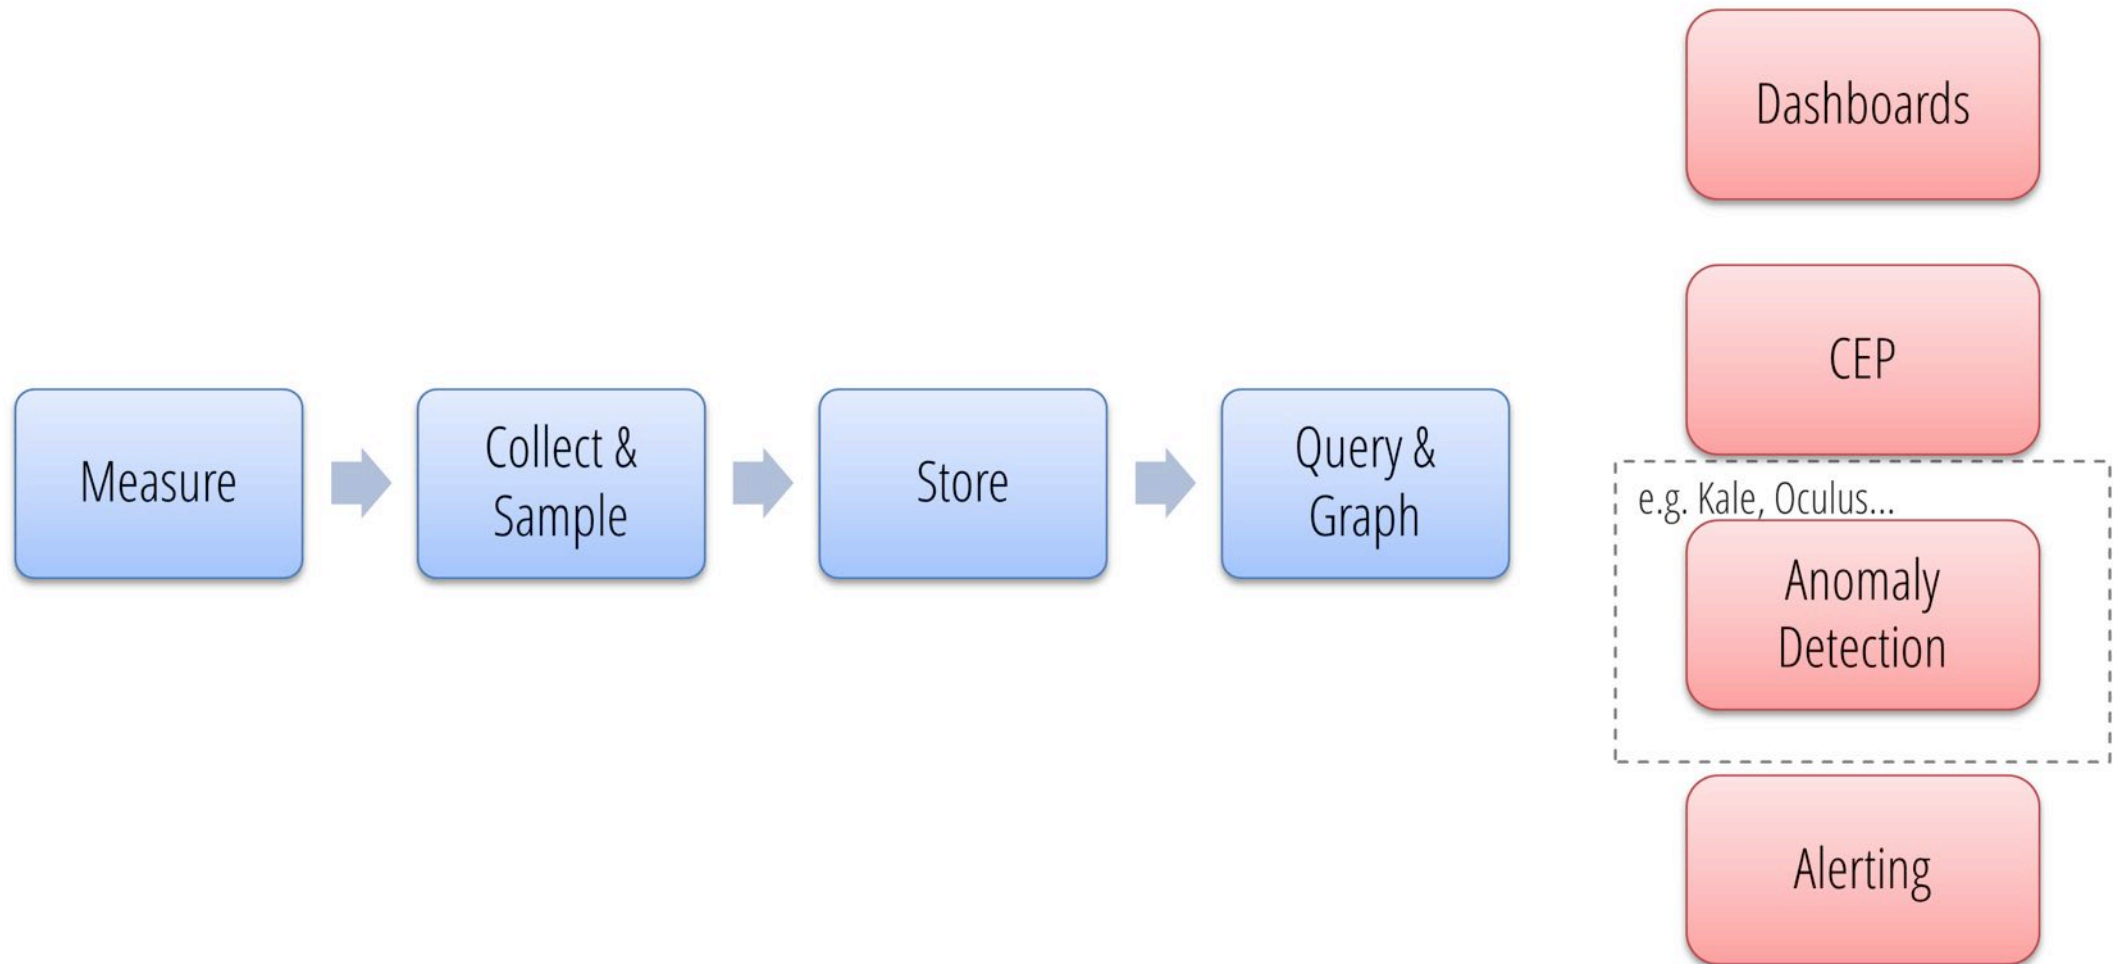


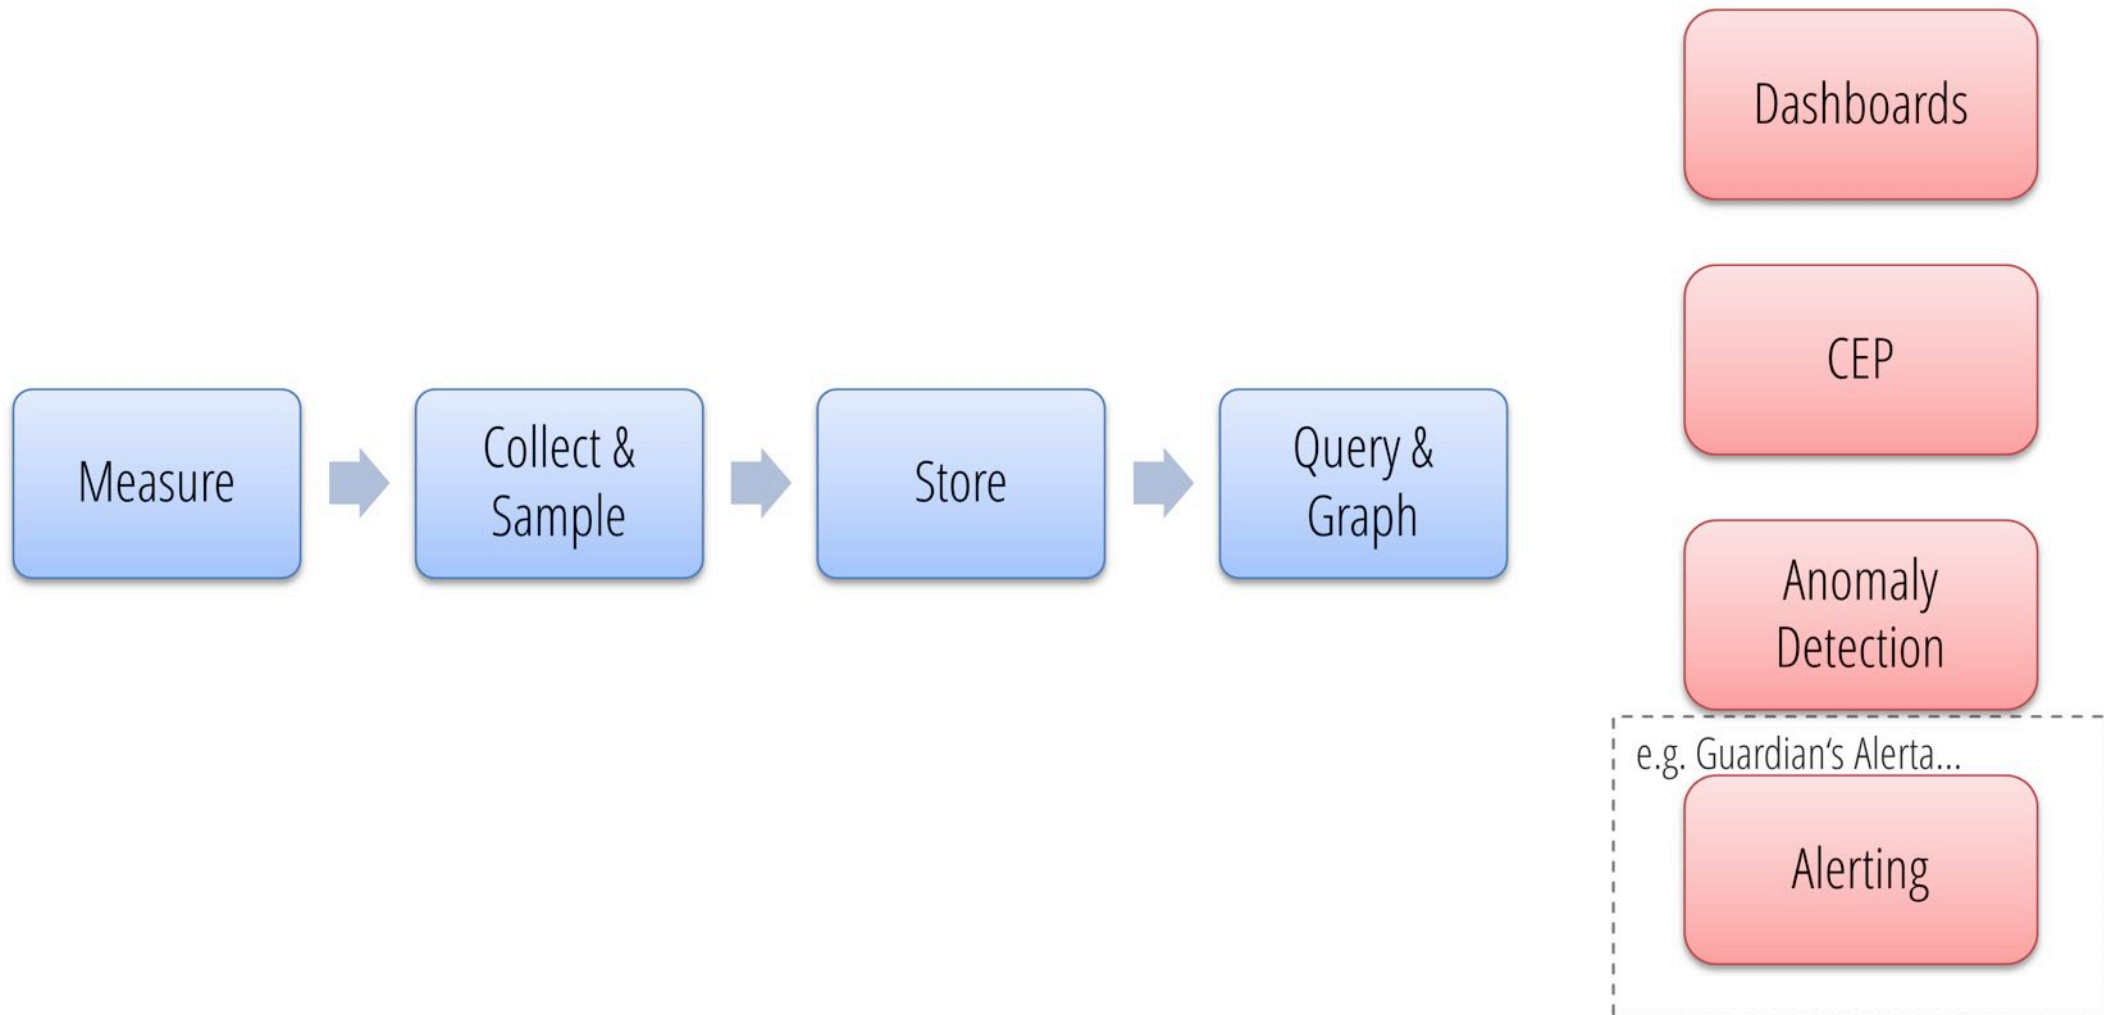












Dashboards




Auto Completer

Graphite Processing Throughput

FROM: 2012-01-01 00:00:00 UNTIL: 2012-01-01 00:00:00

Graphite Processing Throughput



Graph Title

Y Axis Label

Canvas Colors

foreground color

graph options

background color

Text Options

graph cache size

canvas

targets

update every

1

min

carbon.agents.graph02vip.committedPoints

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations


carbon.agents.graph02vip.updateOperations

carbon.agents.graph02vip.updateOperations

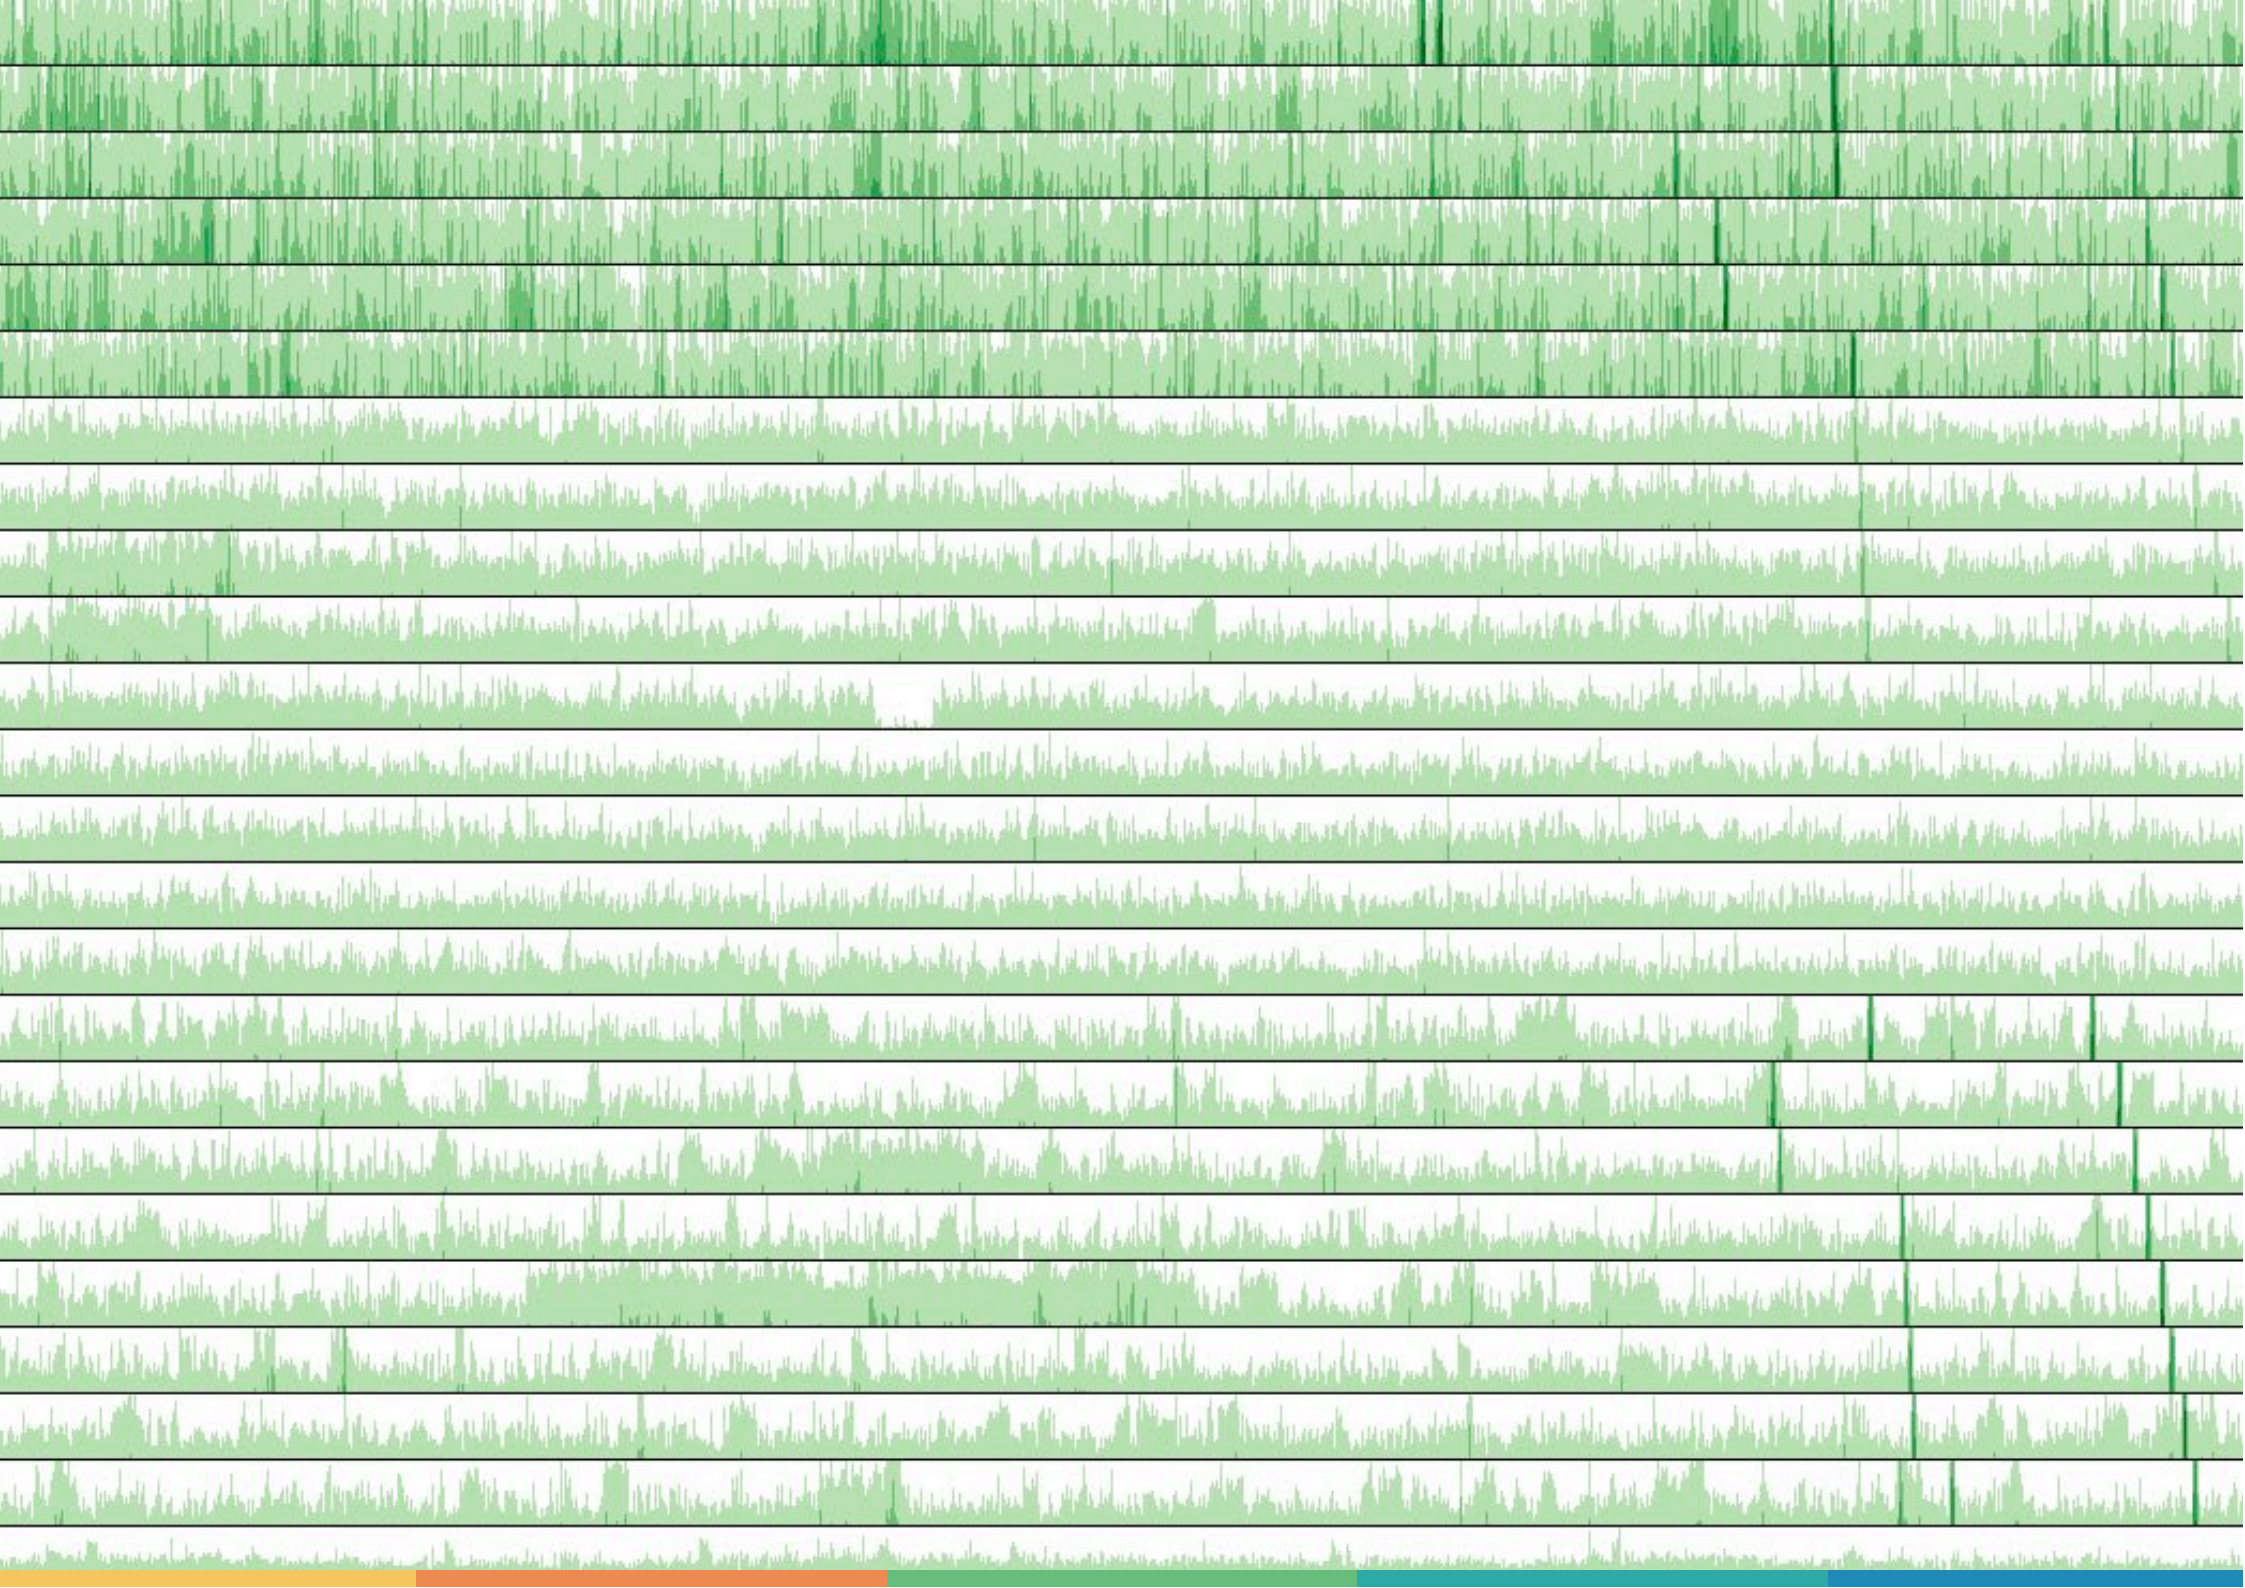
Cubism.js

☐ Mirror ☒ Offset

1 - +




Credits: Michael Bostock



Comparisons

```
var cube = context.cube("http://..."),  
primary = cube.metric("sum(request)"),  
secondary =  
    primary.shift(-7 * 24 * 60 * 60 * 1000);
```




Dashing

Hello

...
This is your shiny new dashboard.

Protip: You can drag the widgets around!

Synergy



Last updated at 17:34

Buzzwords

*Pivoting
Streamlininess
Turn-key
Paradigm shift
Web 2.0
Enterprise
Synergy
Exit strategy
Leverage*

*# of times said around the off
Last updated at 17:34*

Current Valuation

\$58

↑ 142%

In billions

Last updated at 17:34

40

30

20

10

Convergence

43

46s

47s

48s

49s

50s

51s

52s

53s

54s

Best practices

- Measure everything!
- Counters ./ Meters
- Metrics are cheap, but not for free.
- Retention Policies
- Get rid of silos
- Correlate your data
- ...to make better decisions

Prevent the apocalypse!

Logging shows events.

Metrics shows state.

Don't fly blind!



Thanks for your attention!

Alexander Heusingfeld |  @goldstift

Tammo van Lessen |  @taval



<https://www.innoq.com/>

Credits

- › Buuz and Woody
- › Monolith by Ron Cogswell
- › Dave - Wrapping up monolith tins
- › Pleuntje - connected
- › CPU by mbostock
- › Mess by Rev Stan
- › Pay Here by Marc Falardeau
- › Cockpit by Ronnie Rams
- › Stream by Phil Whitehouse
- › Magnifier by John Lodder (Flickr)
- › Flying Saucer, Cup, and Teapot! by Mr Thinktank
- › Ice berg by Derek Keats
- › Gas Meters by mxmstryo (Flickr)
- › Gauge Stock by Andrew Taylor (Flickr)
- › Counter by Marcin Wichary (Flickr)
- › Histogram of legos by color frequency by Jeff Boulter (Flickr)
- › pomodoro timers by Paul Downey (Flickr)
- › Zombie Apocalypse by pasukaru76