Next-generation **alerting and fault detection**

Dieter Plaetinck, raintank

SRECon16 Europe – Dublin, Ireland   July 12, 2016

Alerting, fault & anomaly detection through:

Machine learning

event & stream processing

Alerting IDE's

UNIVERSITEIT
GENT

vimeo

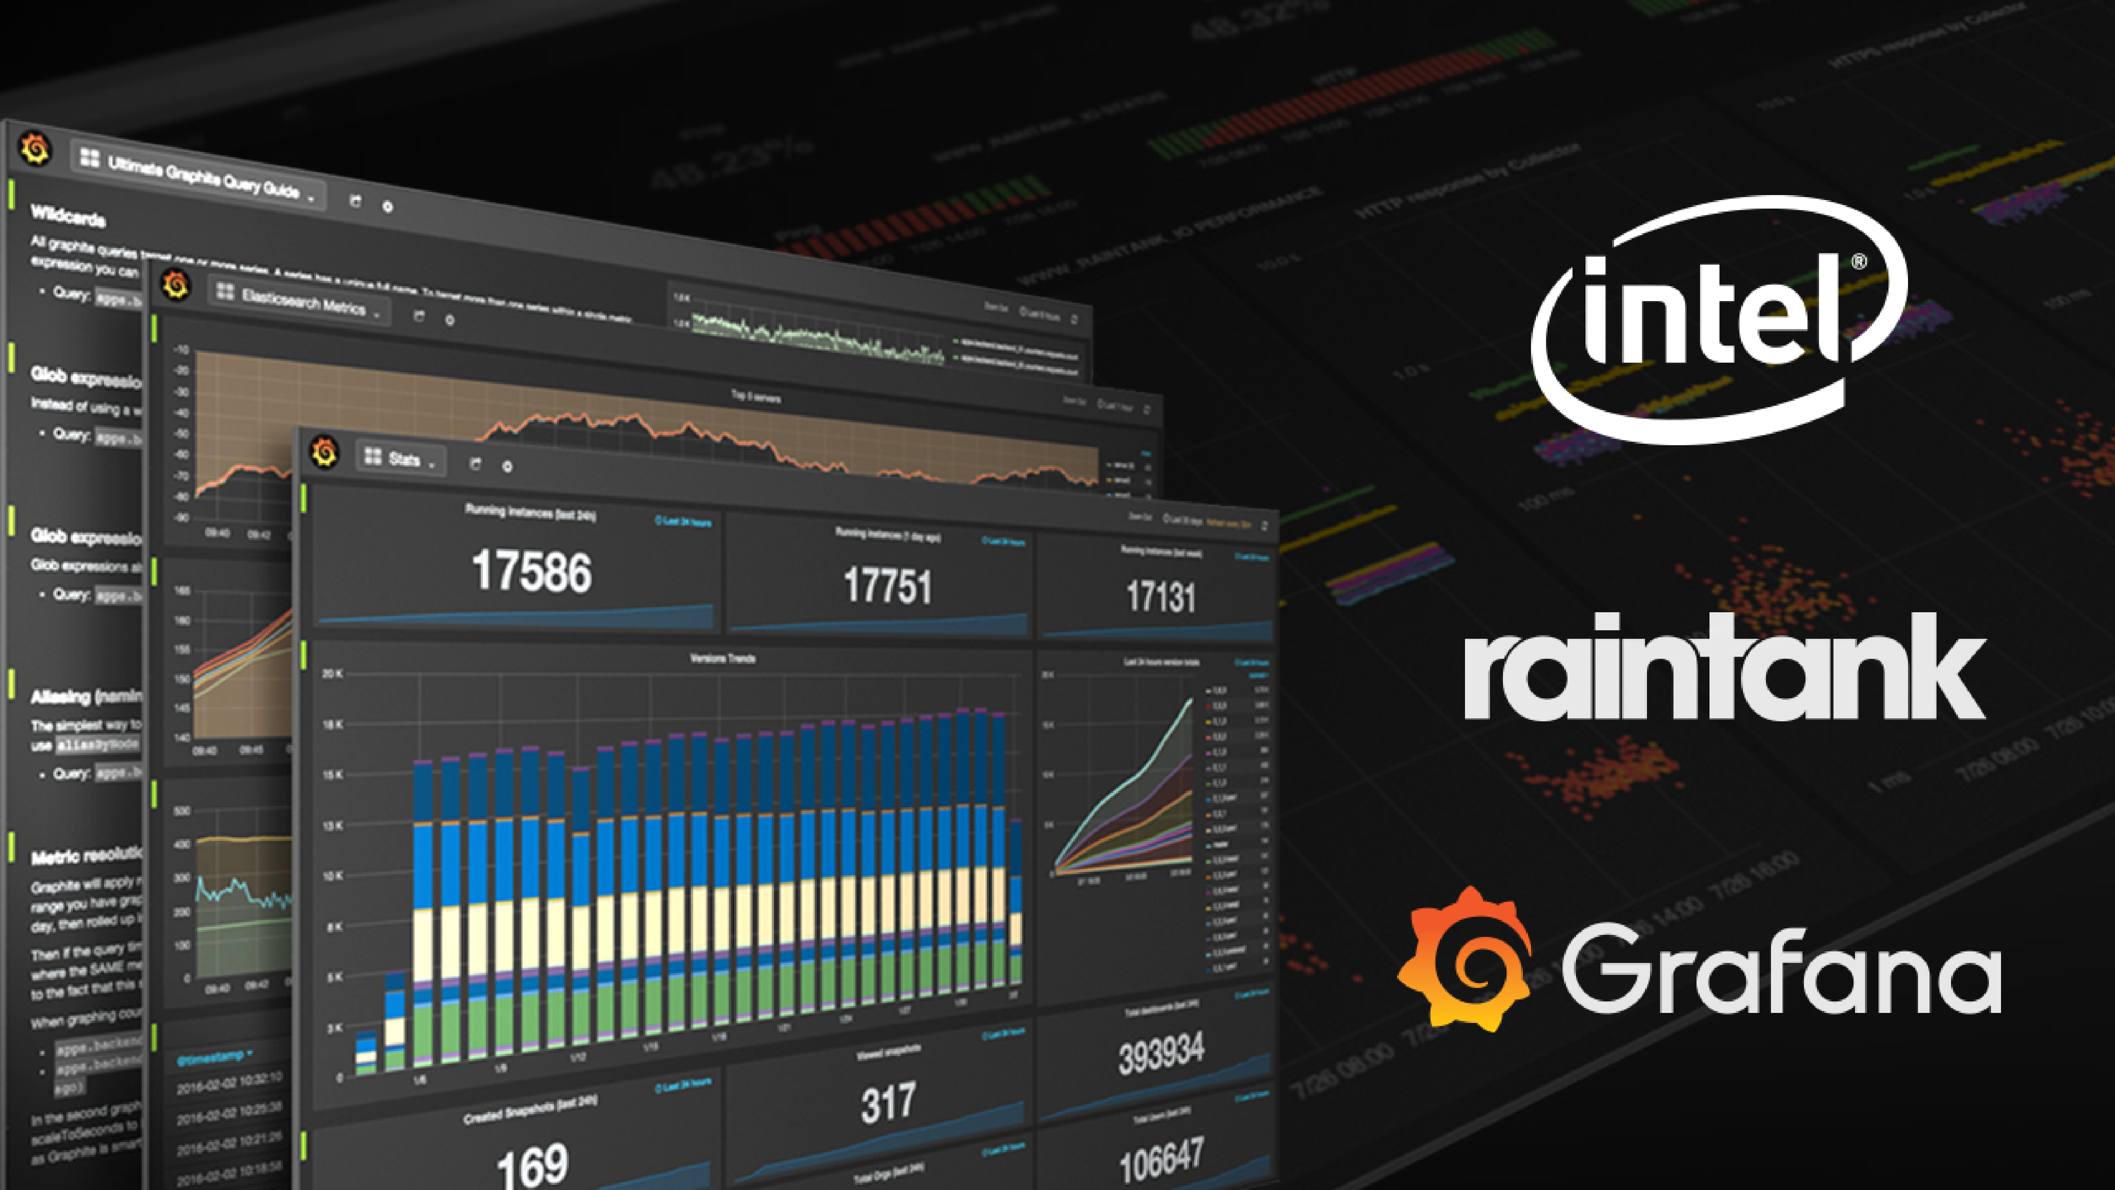
Ultimate Graphite Query Guide
Wildcards
Elasticsearch Metrics
Stats
Running instances (last 24h)
17586
17751
17131
Versions Trends
Viewed snapshots
317
393934
Created Snapshots (last 24h)
169
106647
intel
raintank
Grafana

Grafana.net

Dashboards.

Filter by:

Data Source

All

Panel Type

All

Category

All

Search within this list

Share your dashboards

Sign up for a free Grafana.net account and share your creations with the community.

Sign Up

Carbon C Relay by matejz
Dashboard showing Carbon C Relay internal metrics
Last Updated: July 4th 2016, 6:33 am
Downloads: 45

CollectD Server Metrics by Torkel Ödegaard
CollectD & Graphite Server Metrics Dashboard with CPU, Memory, IO & Disk Stats
Last Updated: June 24th 2016, 4:32 am
Downloads: 218

Dynamic Dashboard by brianl
InfluxDB dashboards for telegraf metrics.
Last Updated: June 30th 2016, 4:17 pm
Downloads: 65

Front HTTP by brianl
Simple HTTP logs Dashboard using Elasticsearch Stack.
Last Updated: June 30th 2016, 4:22 pm
Downloads: 75

Graphite Carbon Metrics by Torkel Ödegaard
Last Updated: June 23rd 2016, 6:26 am
Downloads: 36

Icinga2 by Icinga Project
Icinga2 & Graphite Dashboard
Last Updated: June 27th 2016, 5:34 am
Downloads: 53

Internal Grafana Stats by Grafana Project
Data Proxy request timings (percentiles), dashboard loads, logins etc, Graphite version.
Last Updated: June 24th 2016, 10:42 am
Downloads: 54

Memcached Server Metrics by Torkel Ödegaard
Memcached Server Metrics, requires CollectD memcached plugin and Graphite
Last Updated: June 24th 2016, 4:33 am
Downloads: 7

Nginx Stats by Torkel Ödegaard
Nginx Stats Dashboard using CollectD and Graphite
Last Updated: June 24th 2016, 4:34 am
Downloads: 32

Grafana.net

Plugins.

Plugin Type:

- All
- Panels
- Data source
- Apps

APPLICATION

worldPing

by Raintank

worldPing is a plug-in for Grafana that continually tests, stores and alerts on the global...

APPLICATION

Snap

by Raintank

Snap Data Source and API UI

APPLICATION

Example

by Grafana Project

Example app for Grafana

APPLICATION

Kentik Connect Pro

by Kentik

Kentik Connect Pro allows you to quickly and easily add network activity visibility metrics to...

APPLICATION

NS1 for Grafana

by NS1

NS1 for Grafana allows for the collection and graphing of NS1 data over time.

APPLICATION

Voxter VoIP Platform Metrics

by voxter

Voxter for Grafana allows for the collection and graphing of Voxter data over time.

PANEL

Worldmap Panel

by Grafana Project

World Map panel for grafana. Displays time series data or geohash data from Elasticsearch...

APPLICATION

Zabbix

by Alexander Zobnin

Zabbix plugin for Grafana

PANEL

Clock

by Grafana Project

Clock panel for grafana

PANEL

Pie Chart

APPLICATION

Bosun

DATA SOURCE

Open-Falcon

Also on Grafana.net

- support for Graphite and Grafana

Also on Grafana.net

- support for Graphite and Grafana
- hosted Graphite and Grafana

Presumptions

- Monitoring using metrics in place

Presumptions

- Monitoring using metrics in place
- Alerting on metrics

Presumptions

- Monitoring using metrics in place
- Alerting on metrics
- Alerts need high signal/noise ratio

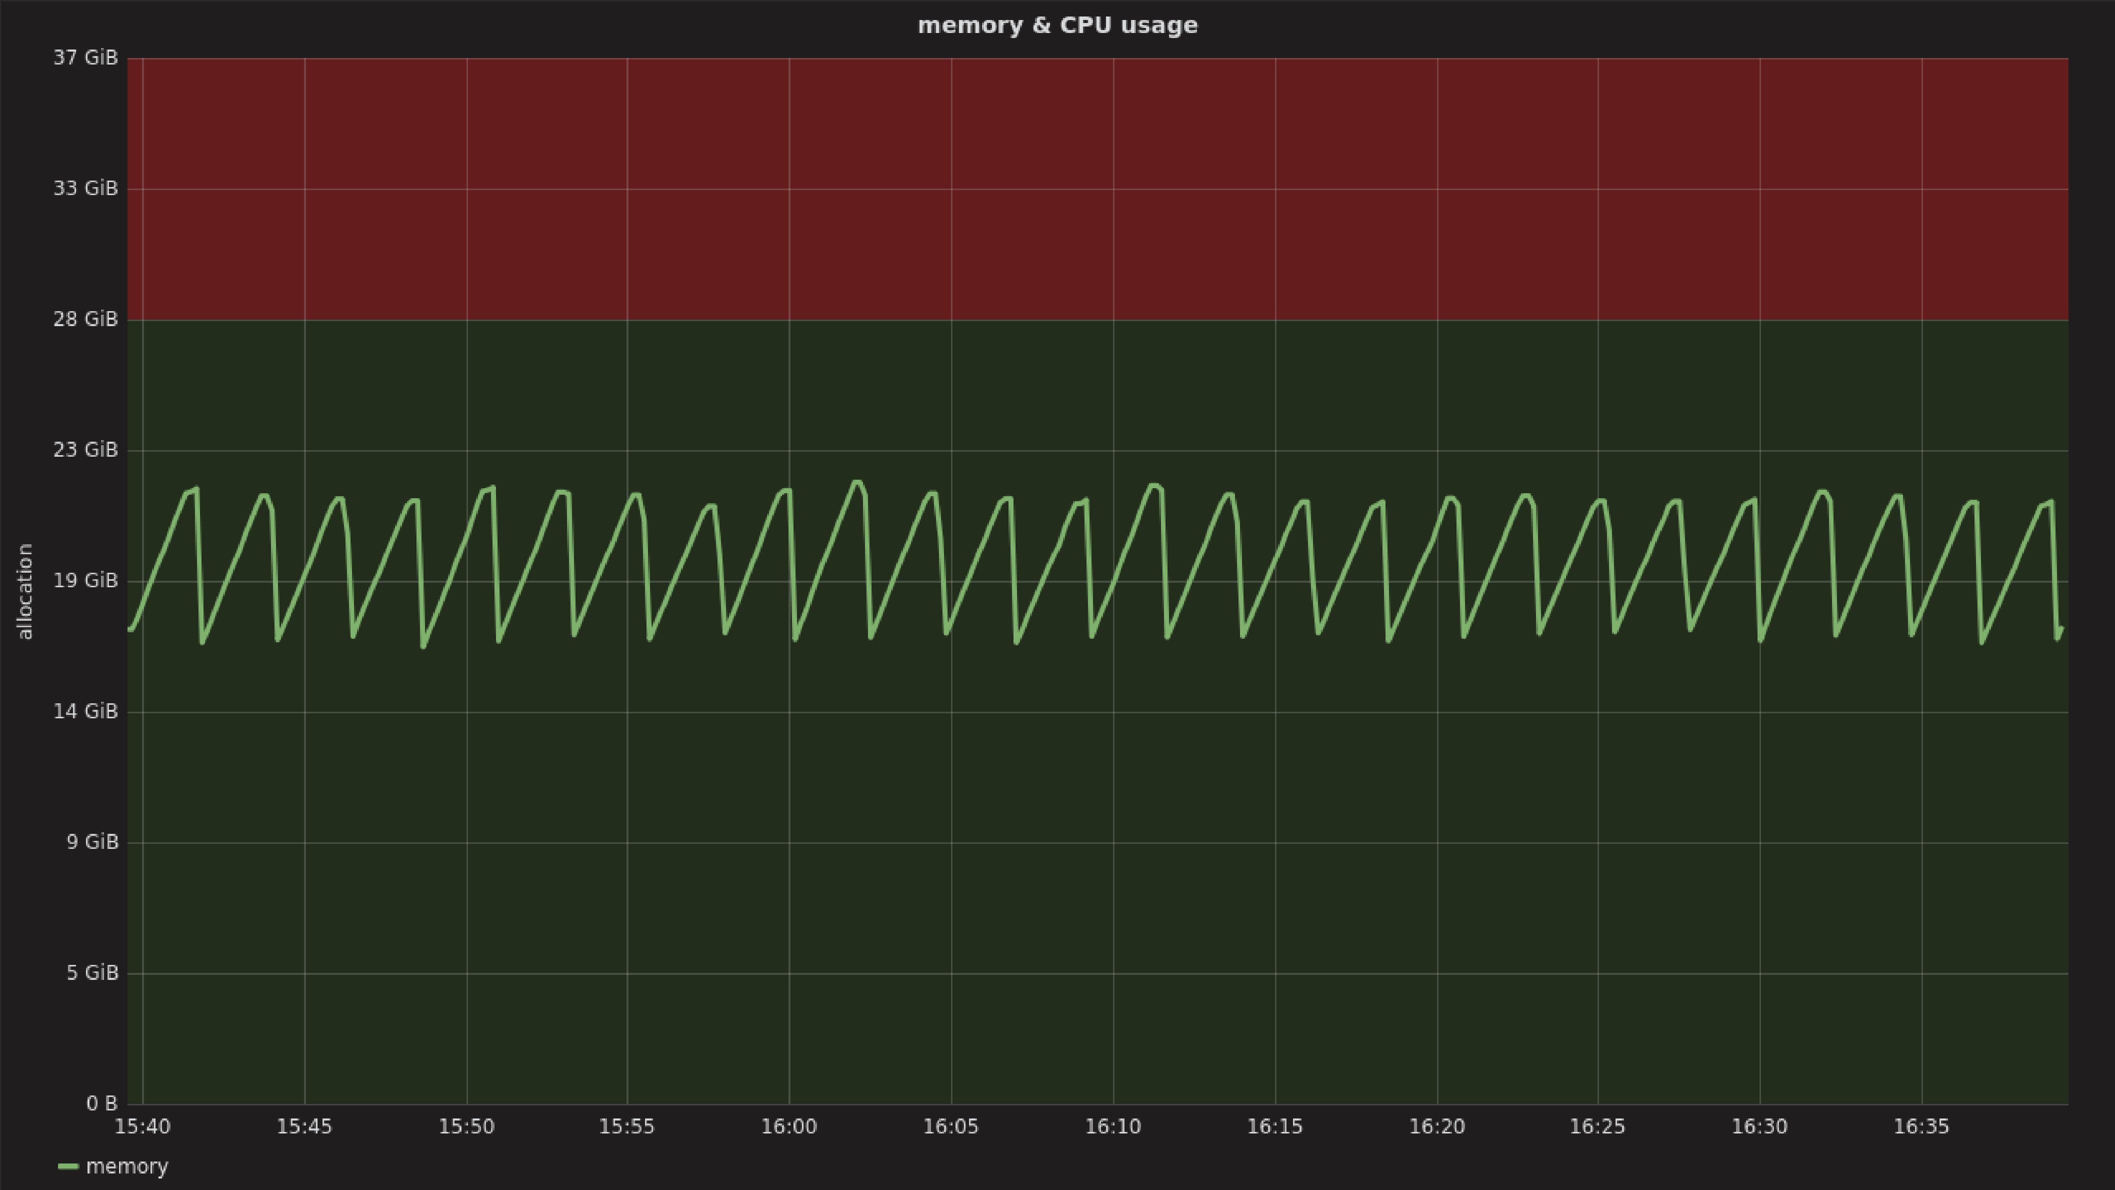
memory & CPU usage
37 GiB
33 GiB
28 GiB
23 GiB
19 GiB
14 GiB
9 GiB
5 GiB
0 B
allocation
15:40
15:45
15:50
15:55
16:00
16:05
16:10
16:15
16:20
16:25
16:30
16:35
memory

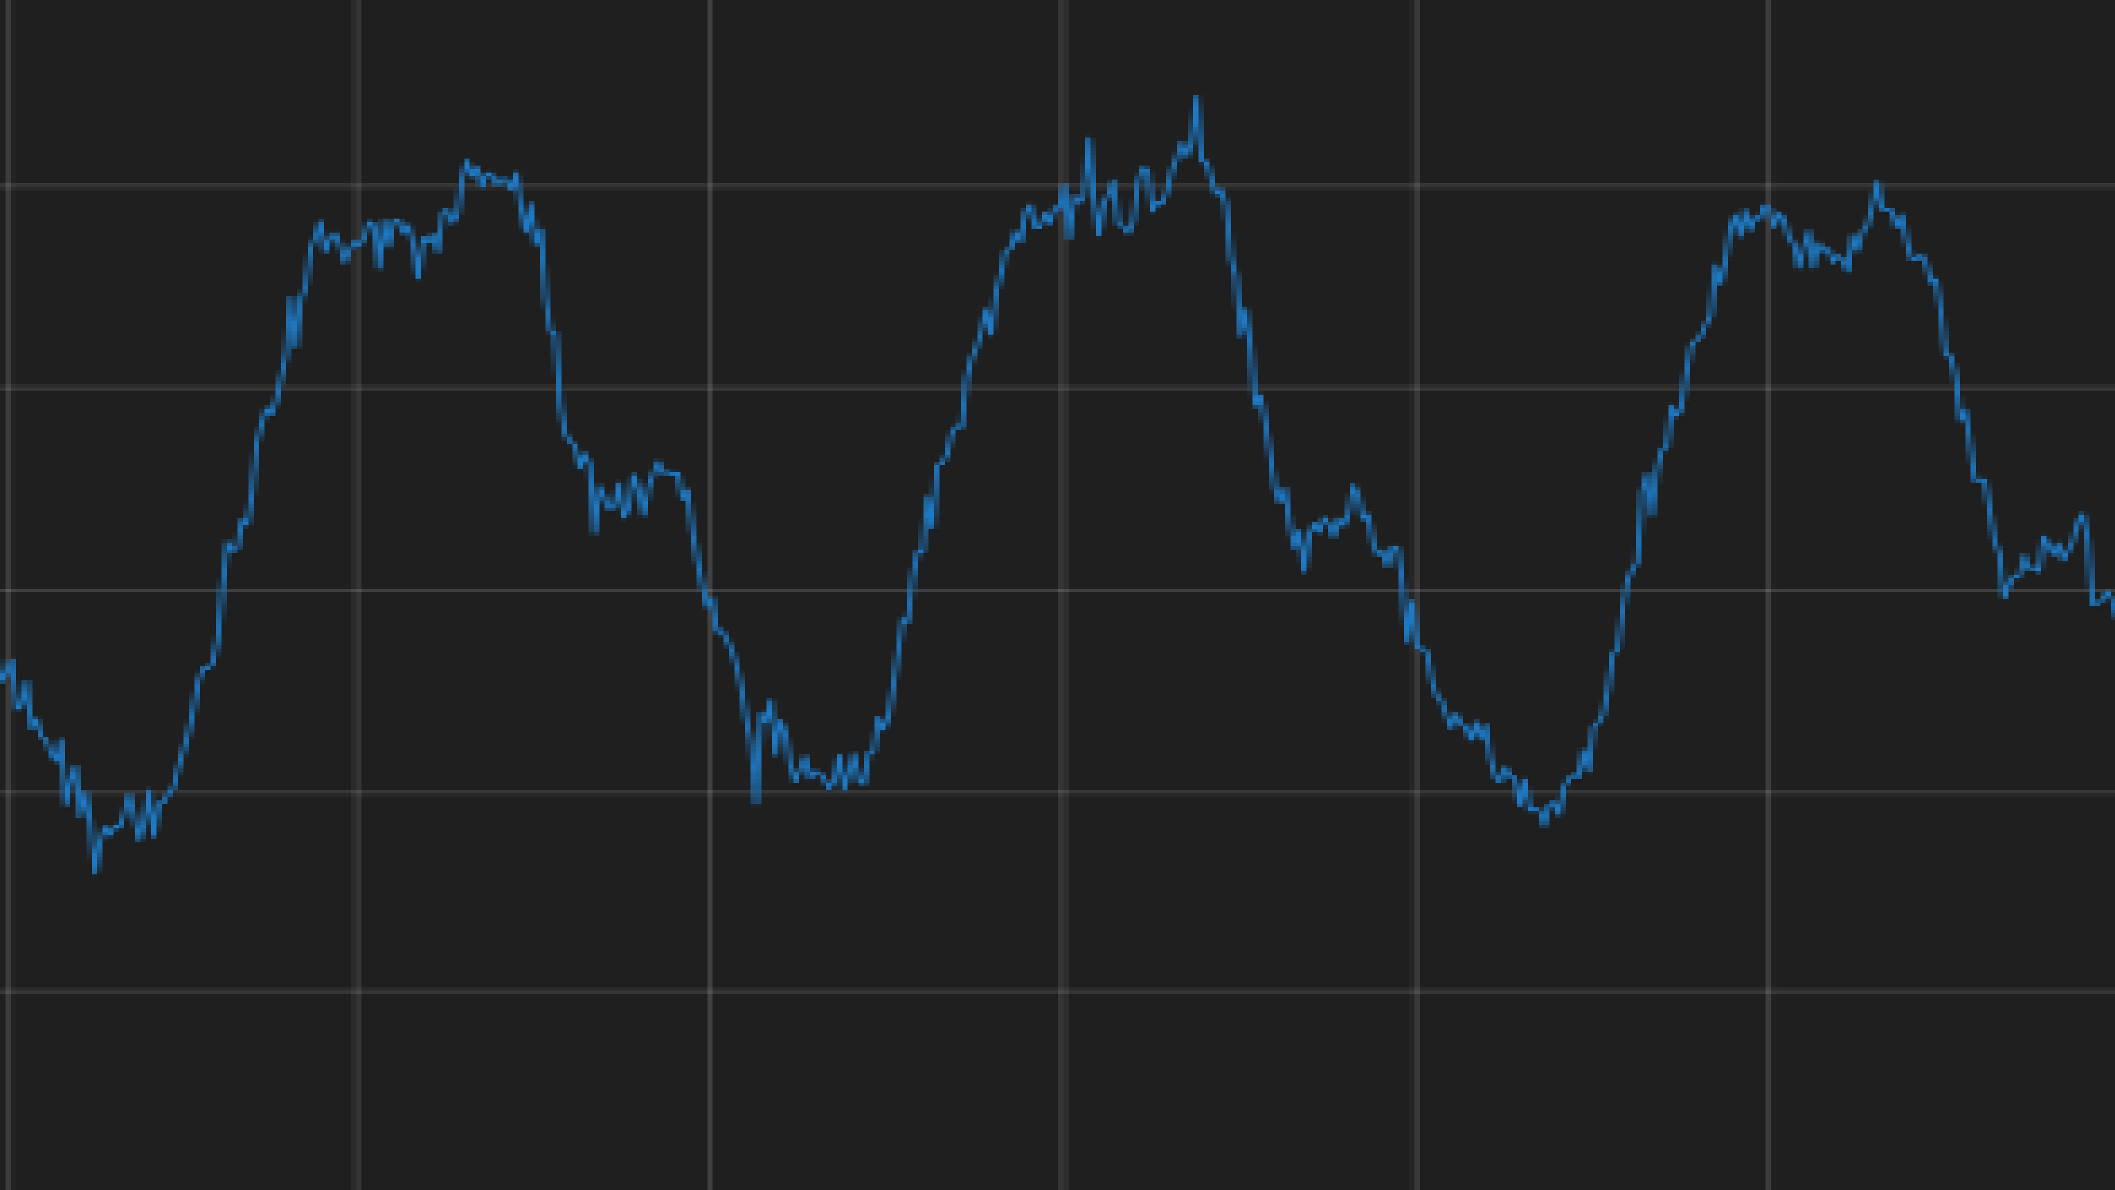

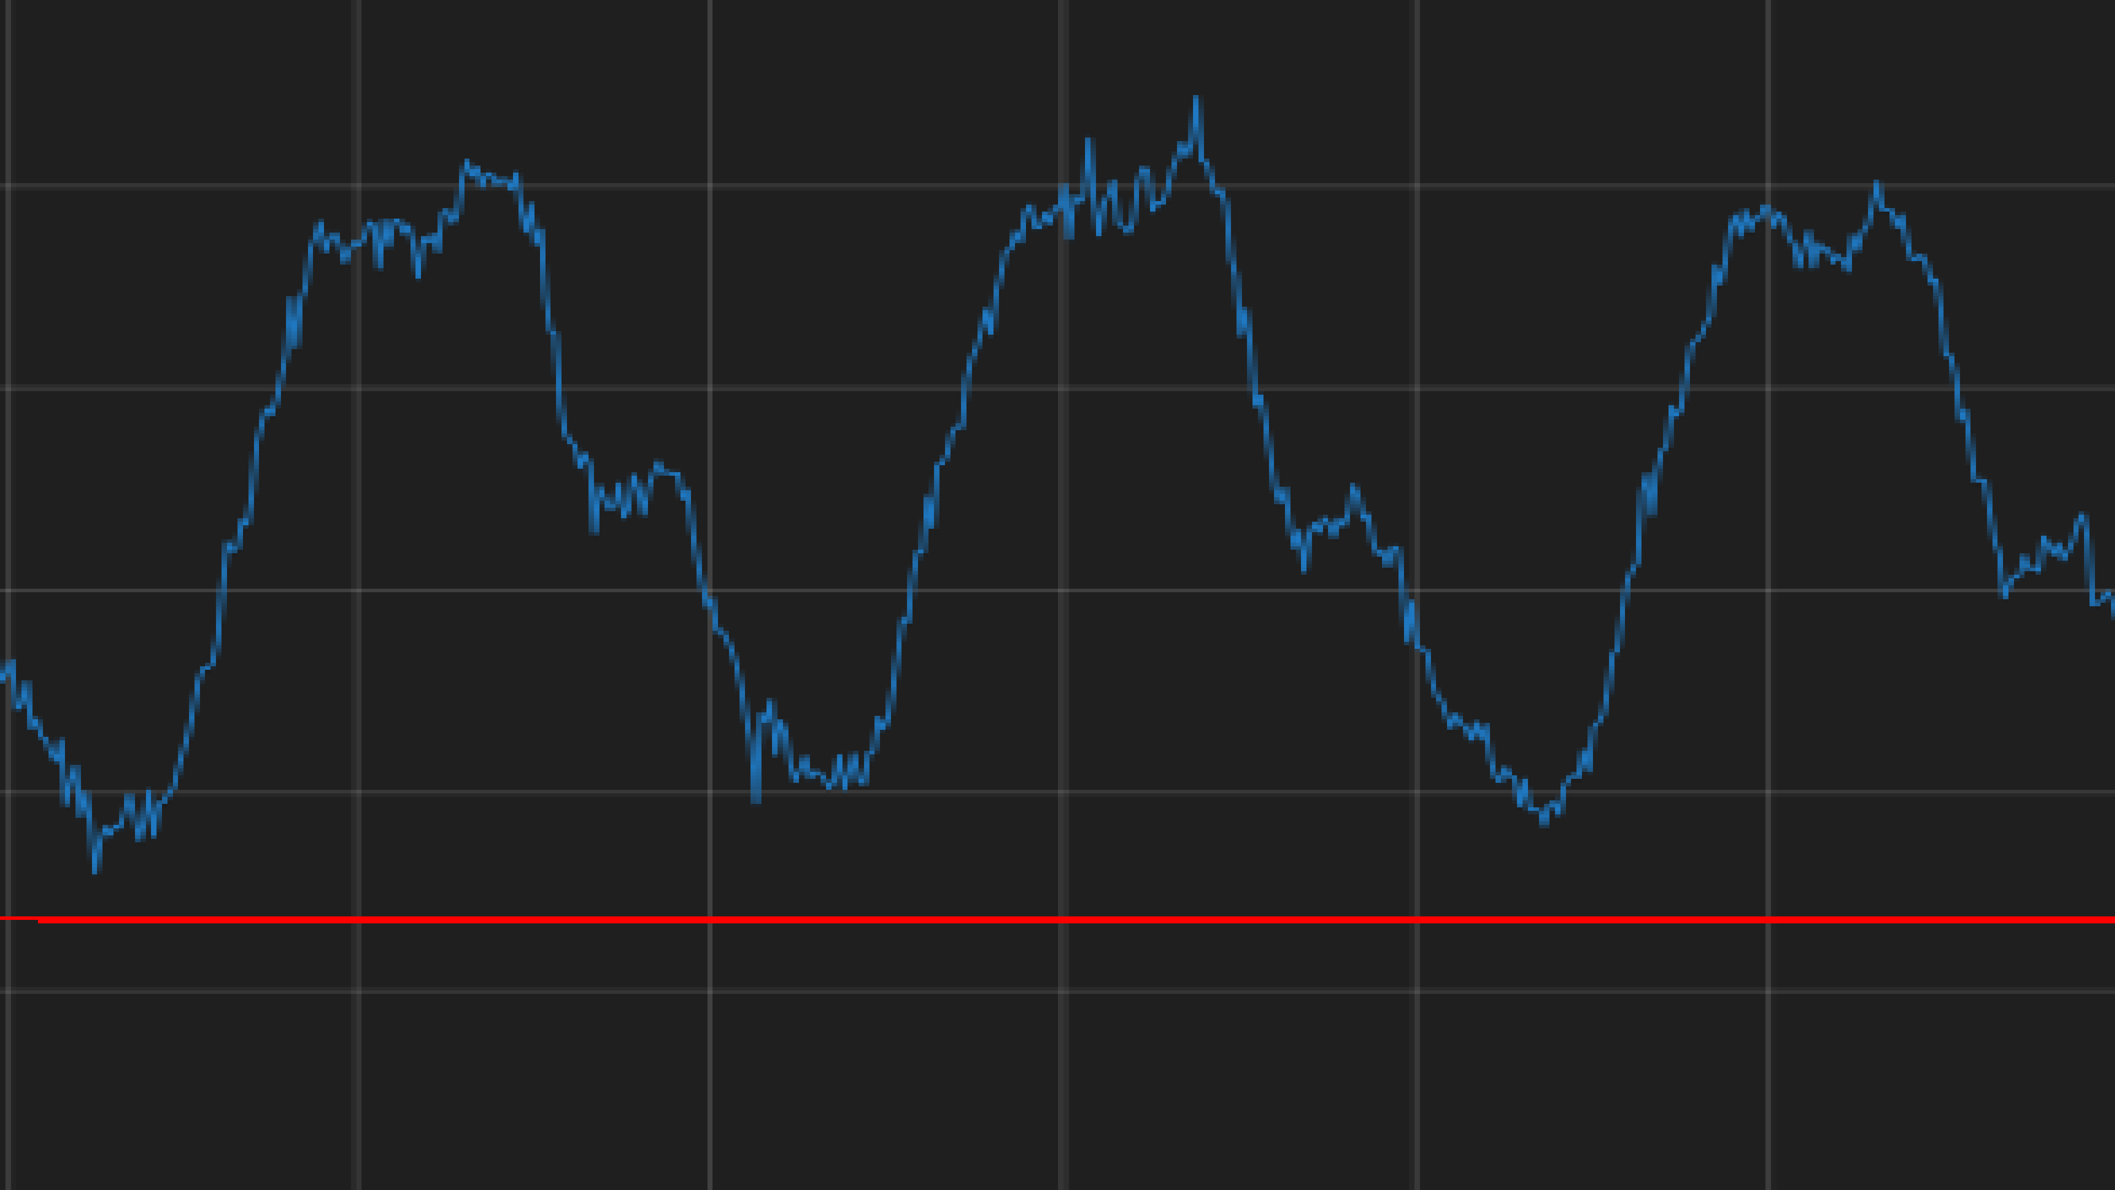

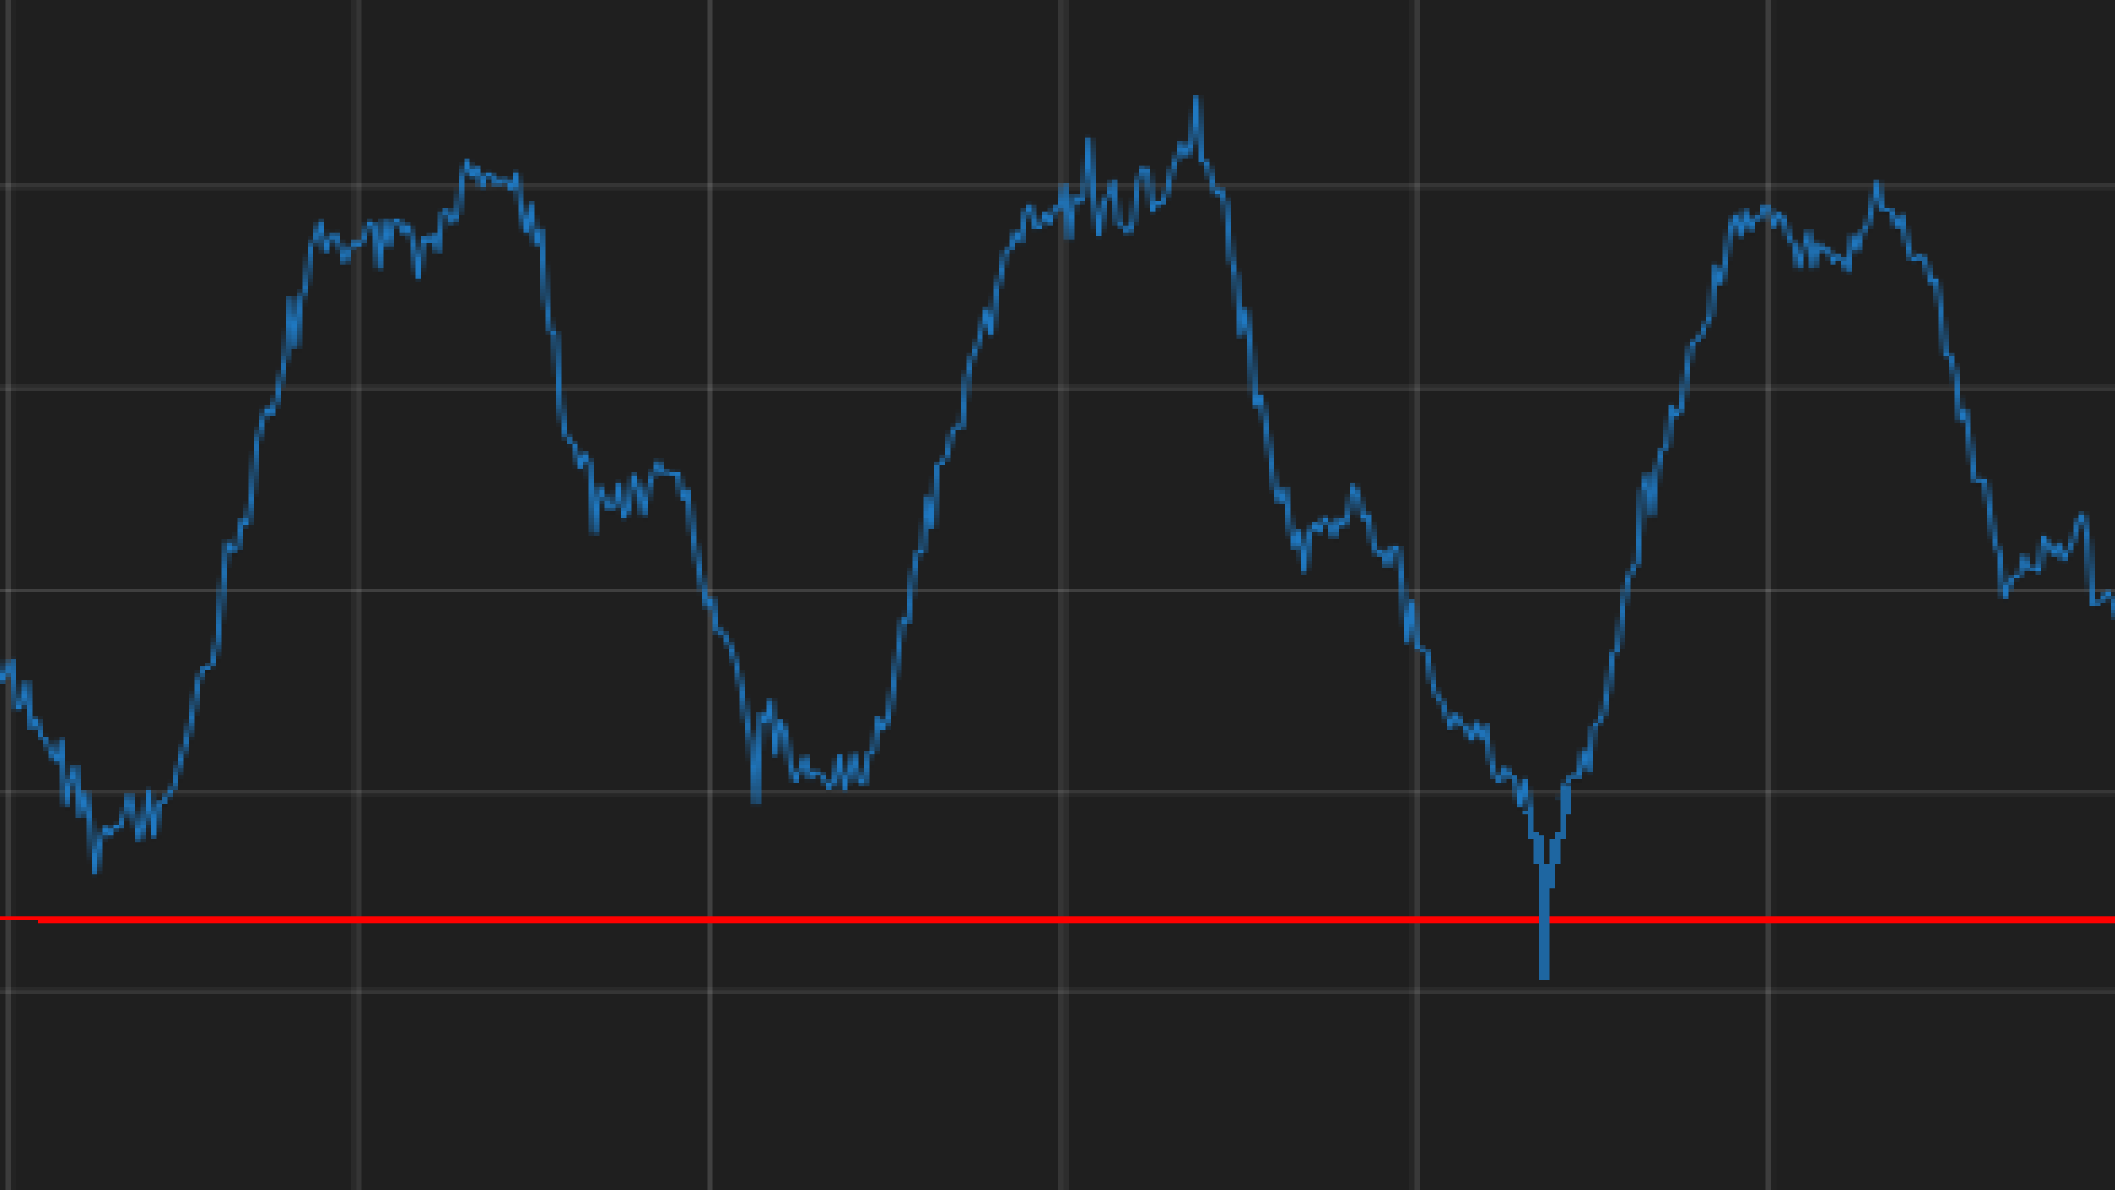

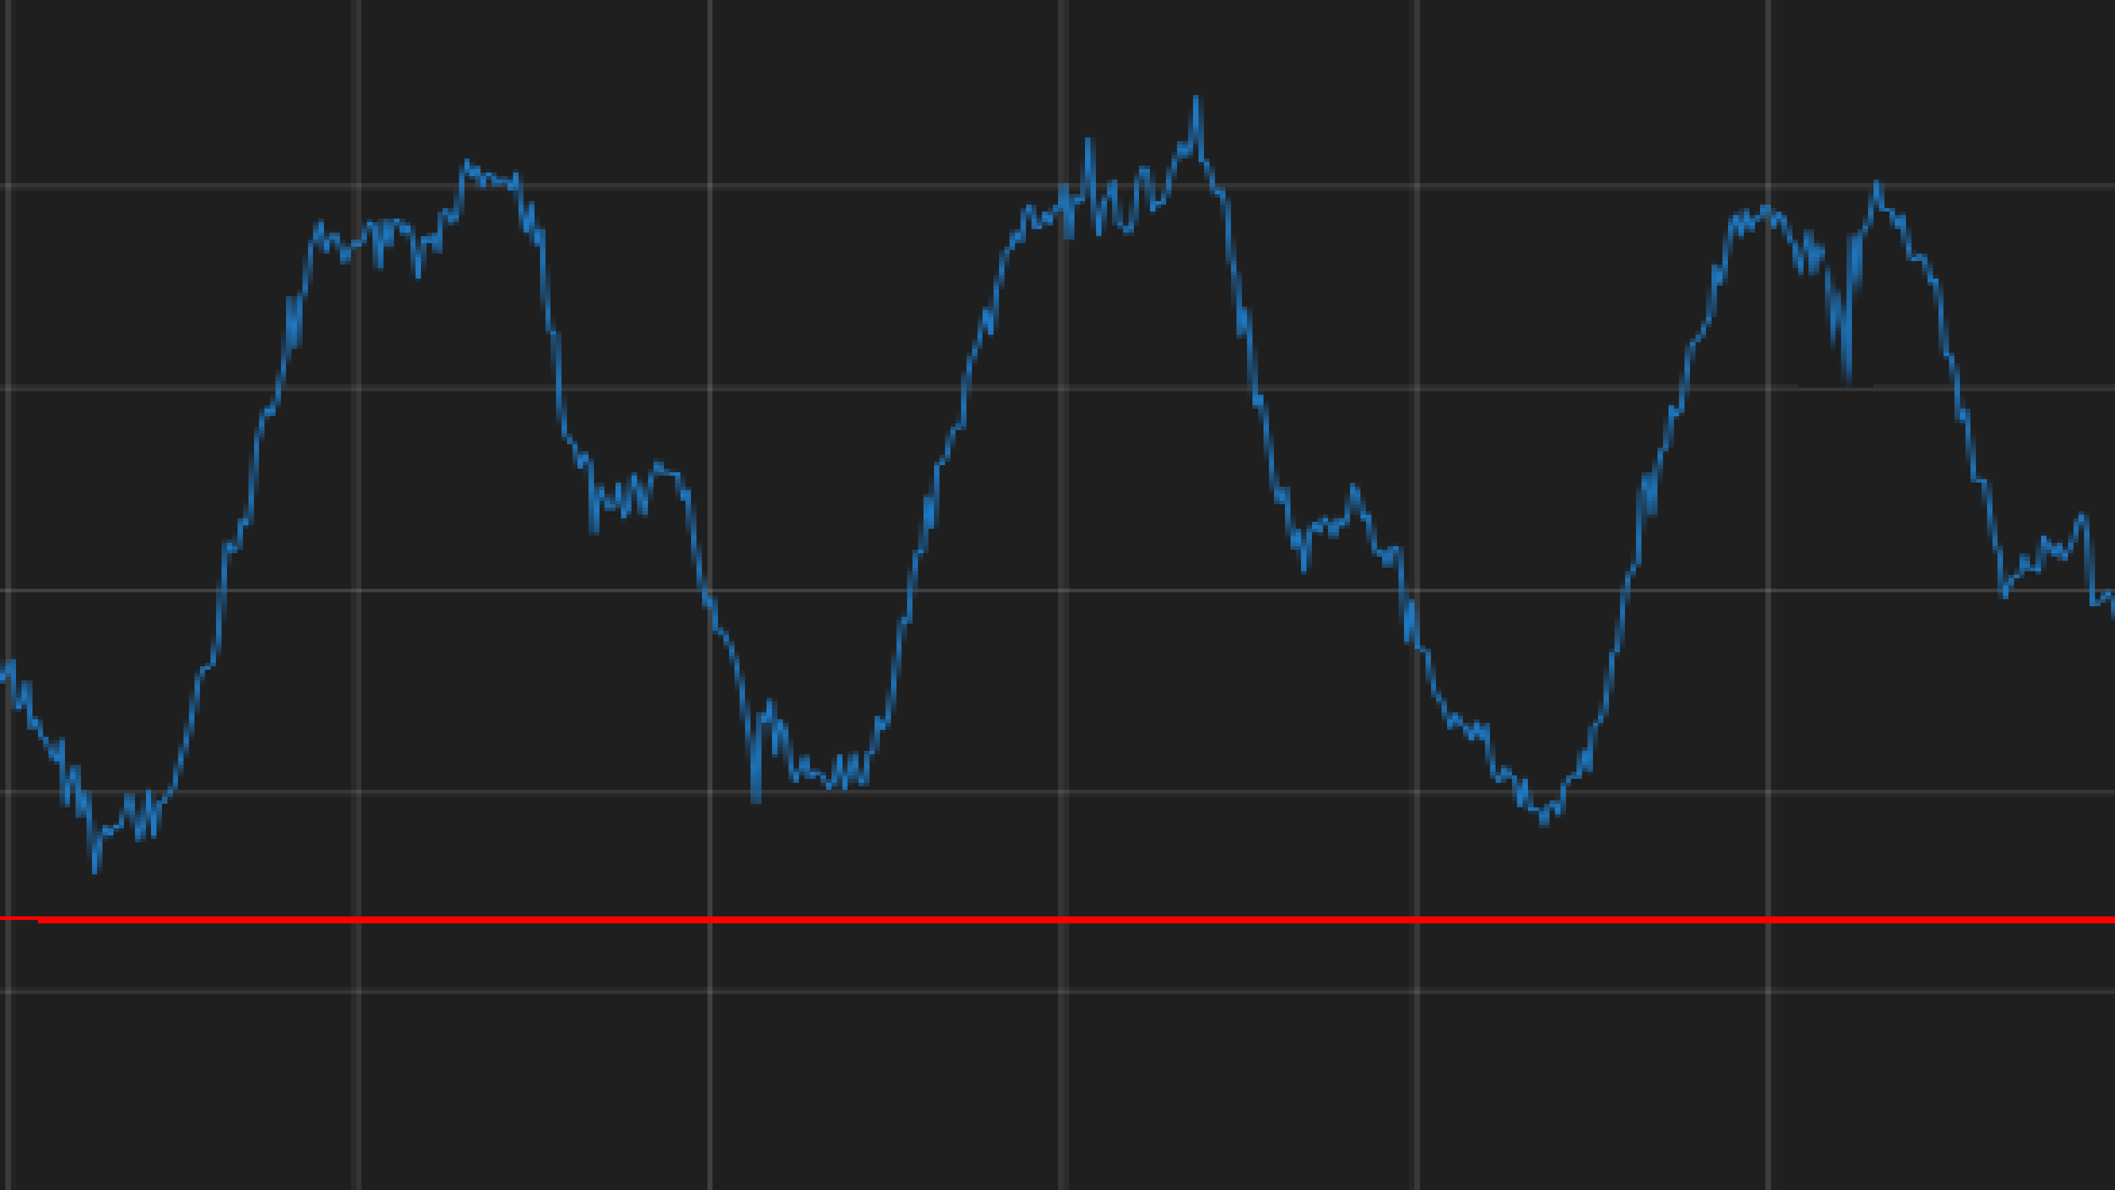

Are you facing a situation where your metric monitoring and alerting that uses static threshold rules is not scaling and you need an automated anomaly detection solution?

I would be happy to know your use case. We are building a solution in this space and researching use cases faced by real world companies.

Request ▾ Follow 6 Comment 1 Share Downvote •••

Add a comment... Comment

Jonah Kowall

There are a lot of these (at least 30 options) commercially and open source already.

Reply... Upvote • Downvote • Report • Feb 22

www.quora.com/unanswered

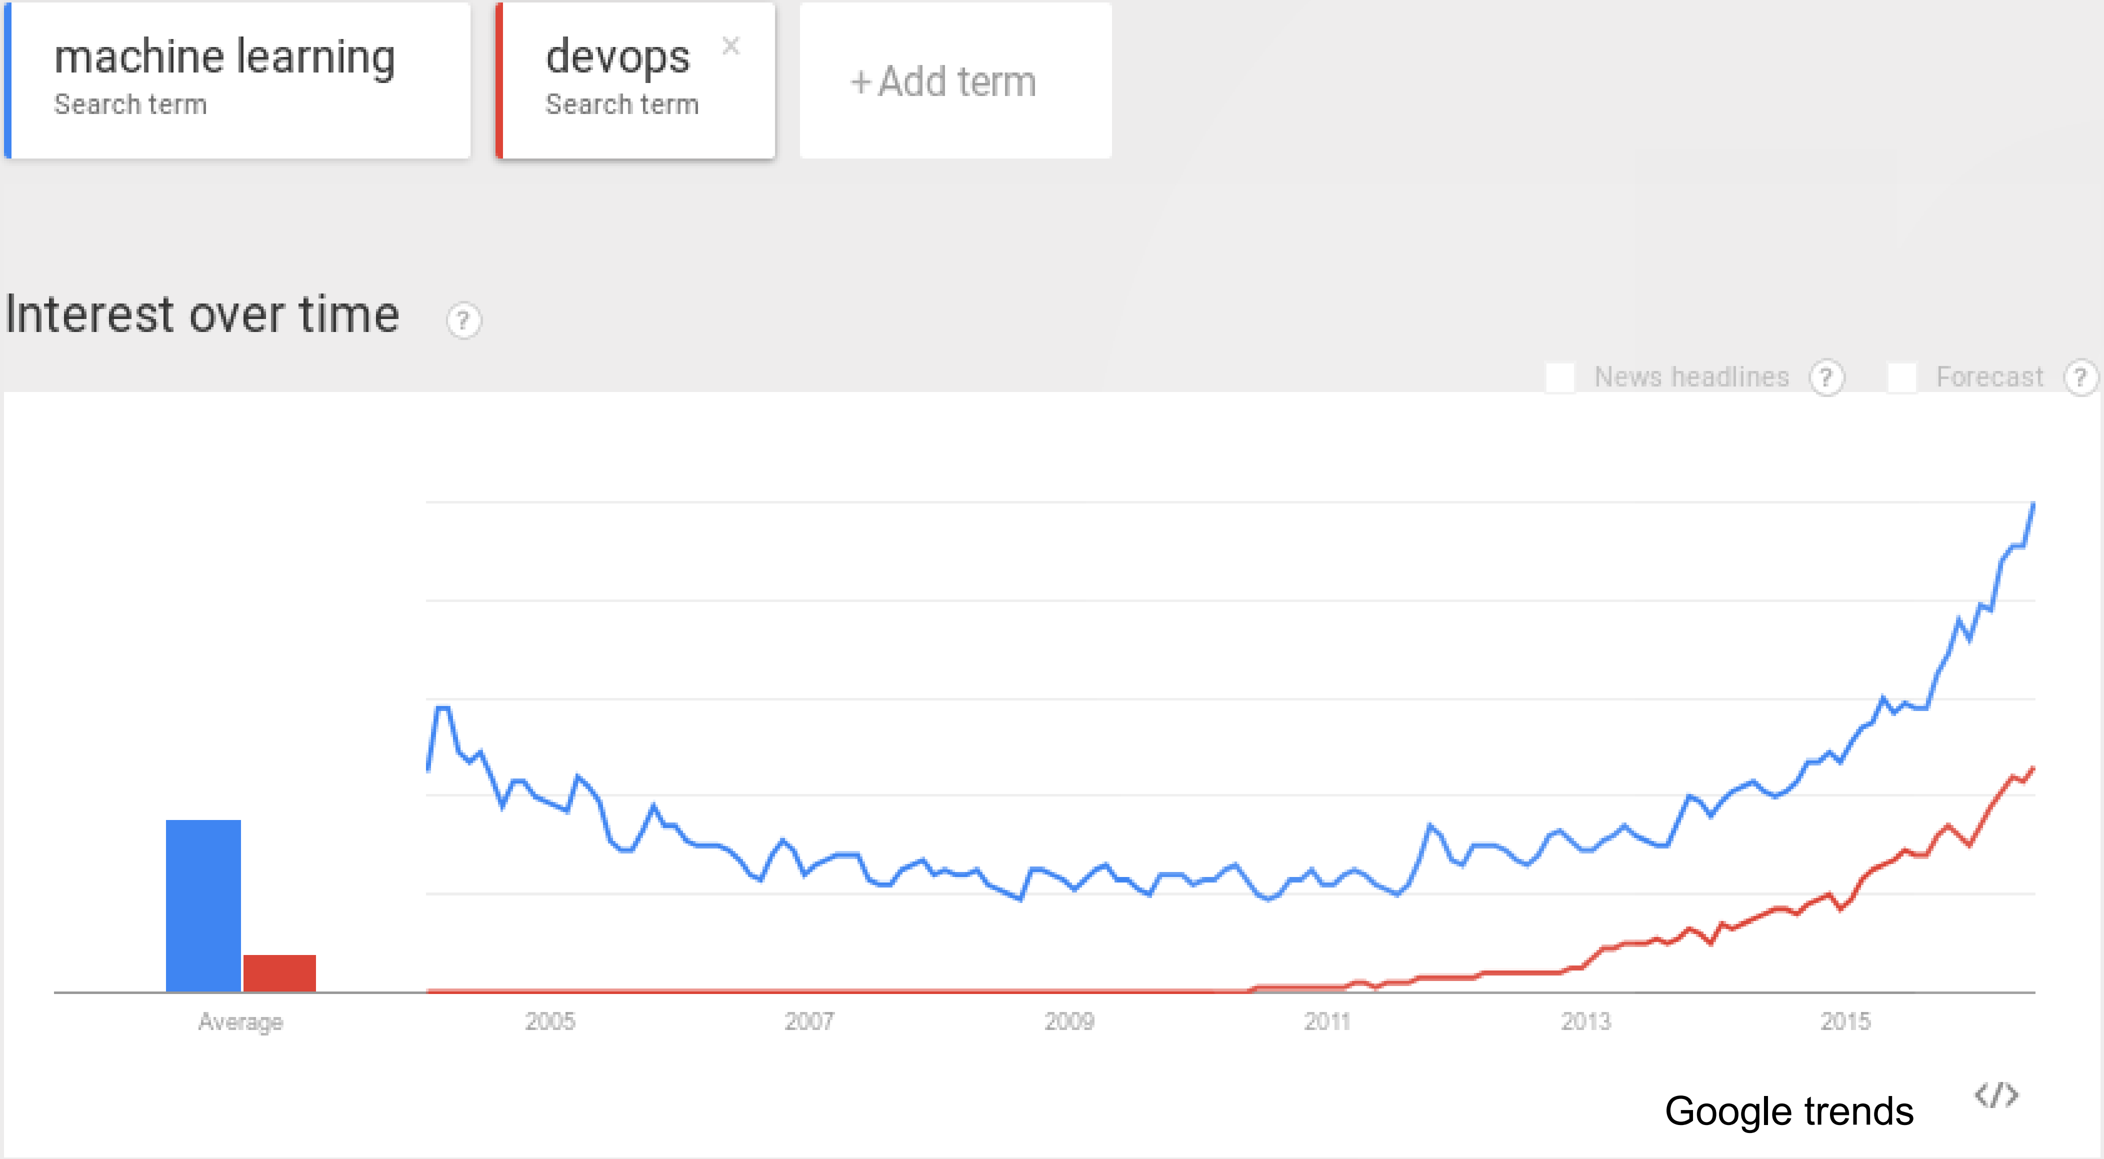
machine learning
Search term
devops
Search term
+ Add term
Interest over time
News headlines
Forecast
Average
2005
2007
2009
2011
2013
2015
Google trends

Static thresholds → automated anomaly detection

- Not scaling / too much data

Static thresholds → automated anomaly detection

- Not scaling / too much data
- Infrastructure complexity

Static thresholds → automated anomaly detection

- Not scaling / too much data
- Infrastructure complexity
- Alerting on Patterns

Machine learning is a subfield of computer science that evolved from the study of pattern recognition and computational learning theory in artificial intelligence. In 1959, Arthur Samuel defined machine learning as a field of study that

gives computers the ability to learn without being explicitly programmed.

Machine learning explores the study and construction of

algorithms that can learn from and make predictions on data.

Such algorithms operate by building a **model** from an example training set of input observations in order to make data-driven predictions or decisions expressed as outputs, rather than following strictly static program instructions.”

https://en.wikipedia.org/wiki/Machine_learning

http://www.extremetech.com/extreme/224445-its-2-0-how-googles-deepmind-is-beating-the-best-in-go-and-why-that-matters

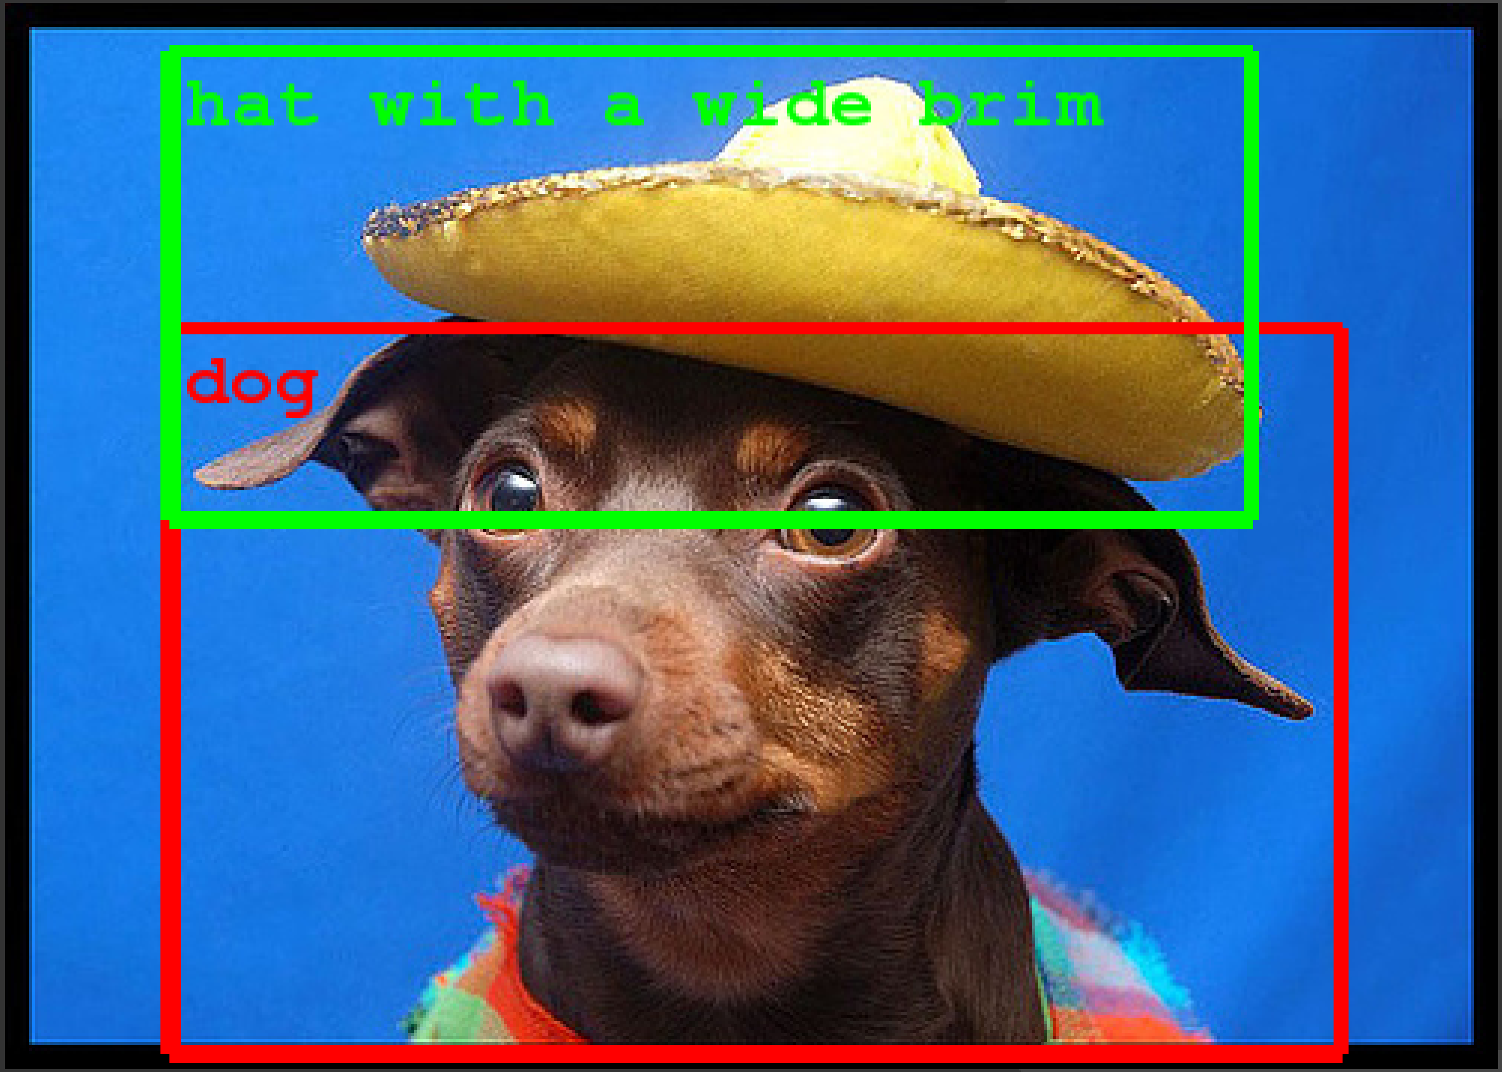

https://research.googleblog.com/2014/09/building-deeper-understanding-of-images.html

Using **machine learning**

for **automated anomaly detection**

a·nom·a·ly

/əˈnäməlē/

noun

1. something that deviates from what is standard, normal, or expected.
   "there are a number of anomalies in the present system"
   synonyms: oddity, peculiarity, abnormality, irregularity, inconsistency, incongruity, aberration, quirk, rarity
   "the growth on the duck's bill is a harmless anomaly"
2. ASTRONOMY
   the angular distance of a planet or satellite from its last perihelion or perigee.

Translations, word origin, and more definitions

fault

/fôlt/

noun

1. an unattractive or unsatisfactory feature, especially in a piece of work or in a person's character.
   "my worst fault is impatience"
2. responsibility for an accident or misfortune.
   "an ordinary man thrust into peril through no fault of his own"
   synonyms: responsibility, liability, culpability, blameworthiness, guilt
   "it was not my fault"

verb

1. criticize for inadequacy or mistakes.
   "her colleagues and superiors could not fault her dedication to the job"
   synonyms: find fault with, criticize, attack, censure, condemn, reproach; More
2. GEOLOGY
   (of a rock formation) be broken by a fault or faults.
   "rift valleys where the crust has been stretched and faulted"

Translations, word origin, and more definitions

Feedback

Challenges.

Daniel Kibblesmith
@kibblesmith

Follow

Amazon is a $250 billion dollar company that reacts to you buying a vacuum by going THIS GUY LOVES BUYING VACUUMS HERE ARE SOME MORE VACUUMS

RETWEETS 7,124

LIKES 11,417

Challenges

1
context

- e.g. Amazon, Facebook, LinkedIn

Challenges

1 context

- e.g. Amazon, Facebook, LinkedIn
- e.g. infrastructure change

Challenges

2
Changing rules

- Games vs your infra

Challenges

2
Changing rules

- Games vs your infra
- Trained model doesn’t work on new scenarios

Challenges

3
Signal strength

Image recognition, security
vs
ops metrics

Challenges

4
relevancy

- e.g. super fast to fast

Challenges

4
relevancy

- e.g. super fast to fast
- e.g. redundancy failover

Challenges

4
relevancy

- e.g. super fast to fast
- e.g. redundancy failover

operator knows best

Challenges

5
Effort

- data prep: filtering, selection, cleaning
- statistical modeling, model selection
- training, testing
- track performance & maintenance
- operate infrastructure
- fitting UX/UI

Challenges

6
Complexity

- Intrinsic

Challenges

6
Complexity

- Intrinsic
- Incidental
  https://engineering.quora.com/Avoiding-Complexity-of-Machine-Learning-Systems

ML / AD for operations

has merits

BUT:

ML / AD for operations

has merits

BUT:

- Anomalies != Faults. Signal/noise trap

ML / AD for operations

has merits

BUT:

- Anomalies != Faults. Signal/noise trap
- Significant effort & complexity

ML / AD for operations

has merits

BUT:

- Anomalies != Faults. Signal/noise trap
- Significant effort & complexity
- Limited use cases

What might help

- Enrich metric metadata (metrics20.org)

What might help

- Enrich metric metadata (metrics20.org)

    clustered, stronger signals with more context

What might help

- Enrich metric metadata (metrics20.org)

    clustered, stronger signals with more context

    classification for model selection

What might help

- Enrich metric metadata (metrics20.org)

    clustered, stronger signals with more context

    classification for model selection

    derive relevancy

What might help

- Enrich metric metadata (metrics20.org)
  - clustered, stronger signals with more context
  - classification for model selection
  - derive relevancy
- integration with CM, PaaS.

What might help

- Enrich metric metadata (metrics20.org)
  - clustered, stronger signals with more context
  - classification for model selection
  - derive relevancy
- integration with CM, PaaS.
  - awareness of infrastructure

What might help

- Enrich metric metadata (metrics20.org)
    - clustered, stronger signals with more context
    - classification for model selection
    - derive relevancy
- integration with CM, PaaS.
    - awareness of infrastructure
    - awareness of infrastructure change

HOW do THEY do it?

https://codeascraft.com/2013/06/11/introducing-kale/

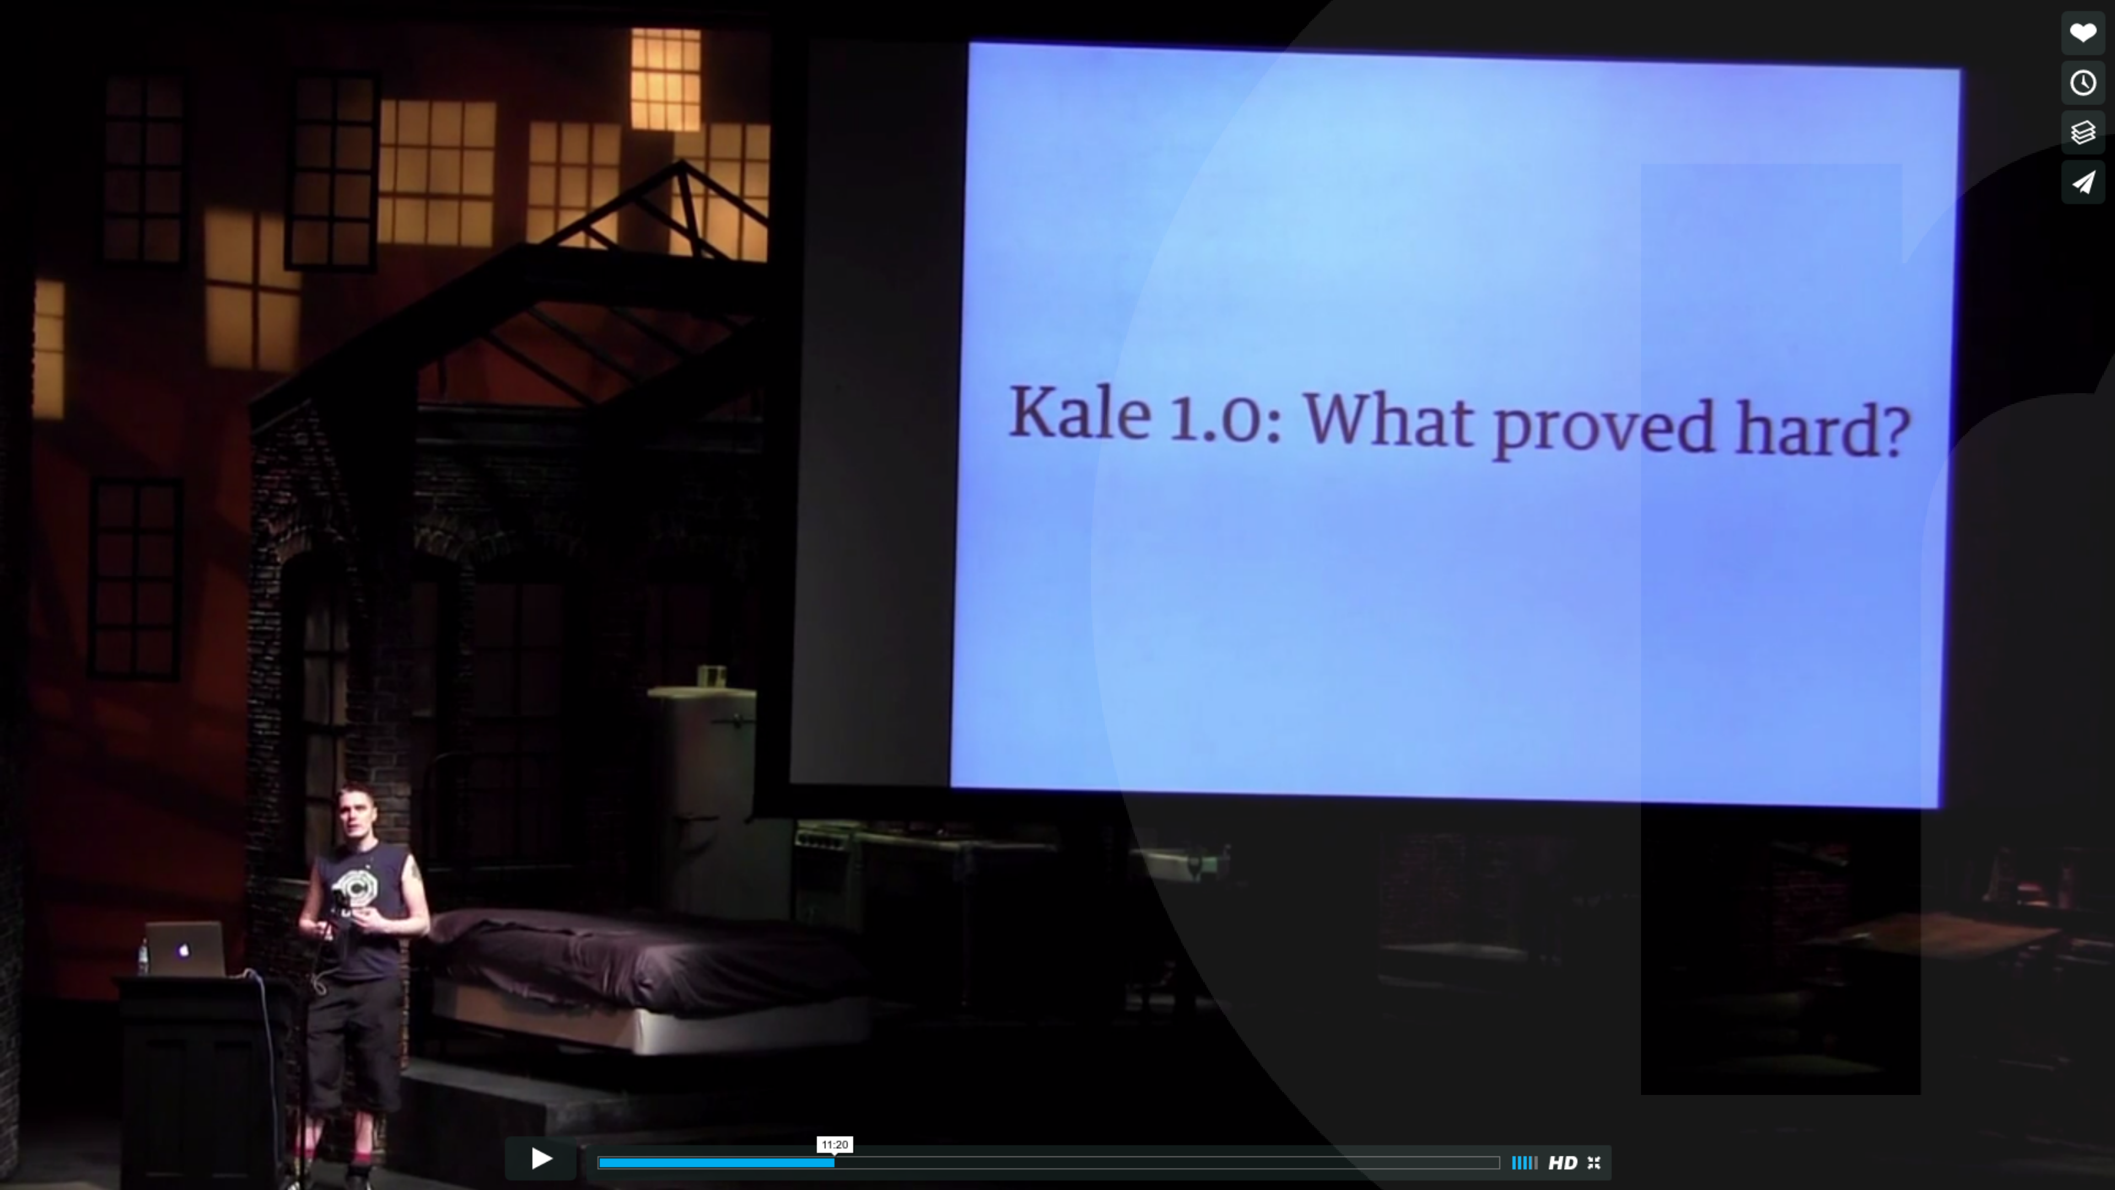
Kale 1.0: What proved hard?
11:20
HD

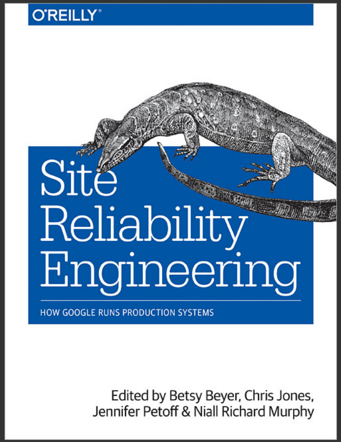

"it's important that monitoring systems - especially the critical path from the onset of a production problem, through a page to a human, through basic triage and deep debugging - be kept **simple** and **comprehensible by everyone on the team**."

"Similarly, to keep noise low and signal high, the elements of your monitoring system that direct to a pager need to be **very simple and robust**. Rules that generate alerts for humans should be **simple to understand and represent a clear failure**."

https://www.oreilly.com/ideas/monitoring-distributed-systems

Conclusion


WYNN
CLASSIC
1000
NO CASH VALUE

CEP

& Stream processing

CEP & stream processing

e.g. storm, riemann.io, spark streaming

in → logic → out

```clojure
(where (service "thumbnailer rate")
  ; Convert build numbers from strings to longs
  (adjust [:build #(Long. %)]
    ; Compute a throughput for each specific build
    (by :build
      (smap #(assoc % :service (str (:service %) " build " (:build %)))
        (rate 5 index)))

    ; Or maybe an old version reported numbers that were 2x larger than they
    ; should have been
    (where (< (:build event) 1055)
      (scale 1/2 index)
      (else index))))
```

CEP & stream processing

Compared to query-based alerting systems:

CEP & stream processing

Compared to query-based alerting systems:

- Good scheduling guarantees/execution timeliness

CEP & stream processing

Compared to query-based alerting systems:

- Good scheduling guarantees/execution timeliness
- Unfamiliar paradigm (maybe)

CEP & stream processing

Compared to query-based alerting systems:

- Good scheduling guarantees/execution timeliness
- Unfamiliar paradigm (maybe)
- Performance/scalability (maybe)

CEP & stream processing

Compared to query-based alerting systems:

- Good scheduling guarantees/execution timeliness
- Unfamiliar paradigm (maybe)
- Performance/scalability (maybe)
- operational complexity (maybe)

Conclusion

Not a bad idea…

But doesn't get to the root of the alerting problems.

Aha!

Bosun

Picture by Matt Simmons

IDE for alerting

Support **programmers** building and maintaining **software**

IDE for alerting

Support **programmers**	building and maintaining **software**
Support **operators**	building and maintaining **alerting**

Bosun Items Expression Rule Editor Silence Short Link Help

graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host")

Date yyyy-mm-dd Time HH:MM:SS Test Rule Image

Results Graph

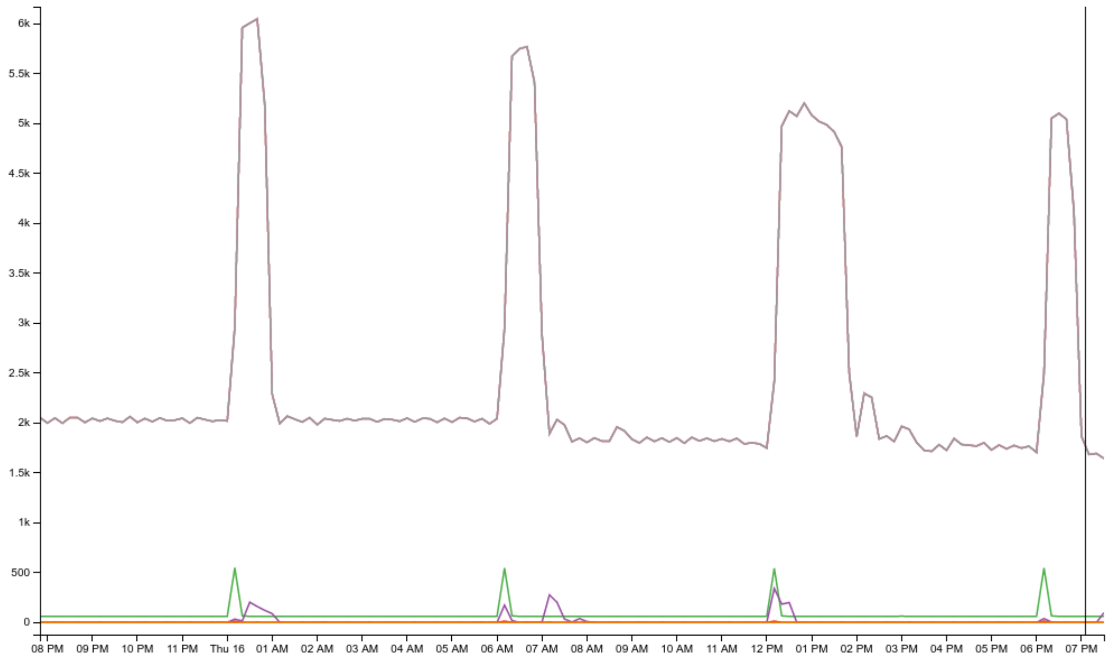

Time: 2016/06/16-19:05:20 (28m26s ago)

{env=development,host=standalone-dev}: 1.50608

{env=poc[redacted]host=[redacted]-poc-tsdb-1}: 0.04667

{env=poc-raintank,host=rt-tsdb-poc-2}: 61.18764

{env=poc[redacted],host=[redacted]-tsdb-test-1}: 0.06003

{env=poc-[redacted],host=[redacted]-poc-tsdb-1}: 0.05333

{env=production,host=cassandra-1_prod}: 1.684k

{env=production,host=cassandra-2_prod}: 1.684k

{env=production,host=cassandra-3_prod}: 1.68356k

Bosun Items Expression Rule Editor Silence Short Link Help

```
$data = graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host")
avg($data)
```

Date yyyy-mm-dd Time HH:MM:SS Test Rule Image

Results Bar Chart

Queries

group	result	computations	
{ env=development, host=standalone-dev }	0.8832880652681122	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	0.88329
{ env=poc-[redacted] host=[redacted]poc-tsdb-1 }	0.049391608391606306	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	0.04939
{ env=poc-raintank, host=rt-tsdb-poc-2 }	75.05446911421897	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	75.05447
{ env=poc-[redacted], host=[redacted]tsdb-test-1 }	16.122761853146944	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	16.12276
{ env=poc-[redacted] host=[redacted]-poc-tsdb-1 }	0.0552913752913754	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	0.05529
{ env=production, host=cassandra-1_prod }	2450.5714628904534	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	2.45057k
{ env=production, host=cassandra-2_prod }	2450.583547400918	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	2.45058k
{ env=production, host=cassandra-3_prod }	2450.5867731118865	avg(graphite("aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)", "24h", "", "env.host"))	2.45059k

Bosun Items Expression Rule Editor Silence Short Link Help

Jump to: alert template lookup notification macro template cassandra_writes Download Toggle Syntax Highlighting Validate

```
        warn = 10000000
        crit = 13500000
    }
    entry env=* {
        warn = 75000
        crit = 100000
    }
}

alert cassandra_writes {
    # alerts if cluster as a whole is doing significantly less writes than same hour of day, on past days
    # also warns if the write load is not well balanced across the nodes
    template = cassandra_writes
    $metric = aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)
    $graphData4d = graphite("$metric", "4d", "", "env.host")
    $graphData2h = graphite("$metric", "2h", "", "env.host")
    $past = median(graphiteBand("$metric", "1h", "1d", "env.host", 2))
    $now = avg(graphite("$metric", "1h", "", "env.host"))
    $cluster_avg_now = avg(t($now, "env"))
    $cluster_avg_past = avg(t($past, "env"))
    $cluster_frac = $cluster_avg_now / $cluster_avg_past
    $dev_from_cluster_avg = $now / $cluster_avg_now
    $devs = t($dev_from_cluster_avg, "env")

```

From yyyy-mm-dd HH:MM To yyyy-mm-dd HH:MM Intervals 5 Step Duration (m)

Email Template Group env=production Test cassandra_writes

Results Template Timeline

Subject

[production] warning: Incident #0 cassandra_writes: cluster write load is 0.73x of what it was in the past. -- nodes in cluster are balanced

Body

Acknowledge alert

Cluster average writes now vs past days

this hour, today: 3961.1274
same hour, past few days: 5426.7667

Deviations from the cluster average

- cassandra-1_prod 1.000088
- cassandra-2_prod 1.000011
- cassandra-3_prod 0.999901

Writes past 4 days

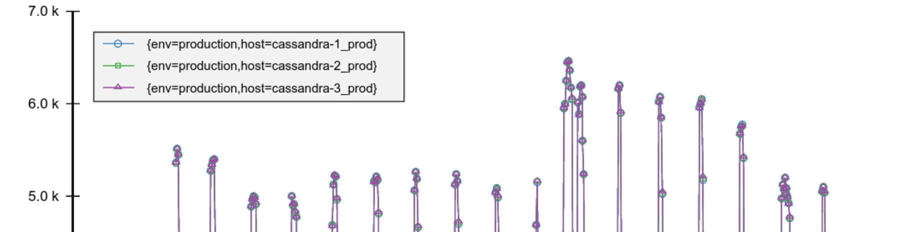

Bosun Items Expression Rule Editor Silence Short Link Help

Jump to: alert template lookup notification macro template cassandra_writes Download Toggle Syntax Highlighting Validate

```
        warn = 75000
        crit = 100000
    }
}

alert cassandra_writes {
    # alerts if cluster as a whole is doing significantly less writes than same hour of day, on past days
    # also warns if the write load is not well balanced across the nodes
    template = cassandra_writes
    $metric = aliasByNode(perSecond(collectd.*.*.*.GenericJMX.cassandra_columnfamilies_stats.counter.WriteOperations),1,3)
    $graphData4d = graphite("$metric", "4d", "", "env.host")
    $graphData2h = graphite("$metric", "2h", "", "env.host")
    $past = median(graphiteBand("$metric", "1h", "1d", "env.host", 2))
    $now = avg(graphite("$metric", "1h", "", "env.host"))
    $cluster_avg_now = avg(t($now, "env"))
    $cluster_avg_past = avg(t($past, "env"))
    $cluster_frac = $cluster_avg_now / $cluster_avg_past
    $dev_from_cluster_avg = $now / $cluster_avg_now
    $devs = t($dev_from_cluster_avg, "env")

    $traffic_warn = $cluster_frac < 0.8
    $traffic_crit = $cluster_frac < 0.6
    $unbalanced_warn = min($devs) < 0.67 || max($devs) > 1.3
```

From 2016-06-12 HH:MM To 2016-06-16 HH:MM Intervals 50 Step Duration (m) 115

Email Template Group env=production Test cassandra_writes

Results Template Timeline

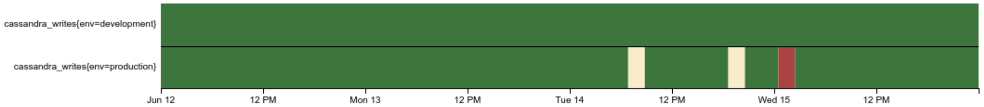

2016/06/15-14:53:27
cassandra_writes{env=production}

cassandra_writes{env=development}	1 events
cassandra_writes{env=production}	7 events
cassandra_writes{env=production}: normal	2016/06/12-00:00:00 (4d19h19m27s ago) for 2d06h51m25s
cassandra_writes{env=production}: warning	2016/06/14-06:51:25 (2d12h28m01s ago) for 1h57m33s
cassandra_writes{env=production}: normal	2016/06/14-08:48:58 (2d10h30m28s ago) for 9h47m45s
cassandra_writes{env=production}: warning	2016/06/14-18:36:44 (2d00h42m43s ago) for 1h57m33s
cassandra_writes{env=production}: normal	2016/06/14-20:34:17 (1d22h45m09s ago) for 3h55m06s
cassandra_writes{env=production}: critical	2016/06/15-00:29:23 (1d18h50m03s ago) for 1h57m33s
cassandra_writes{env=production}: normal	2016/06/15-02:26:56 (1d16h52m30s ago) for 21h33m03s

Key features

1
[historical] testing

vs traditional alerting, machine learning

Key features

2
data juggling

Arbitrary scope

Key features

2
data juggling

Arbitrary scope
Arbitrary data

Key features

3

dependencies

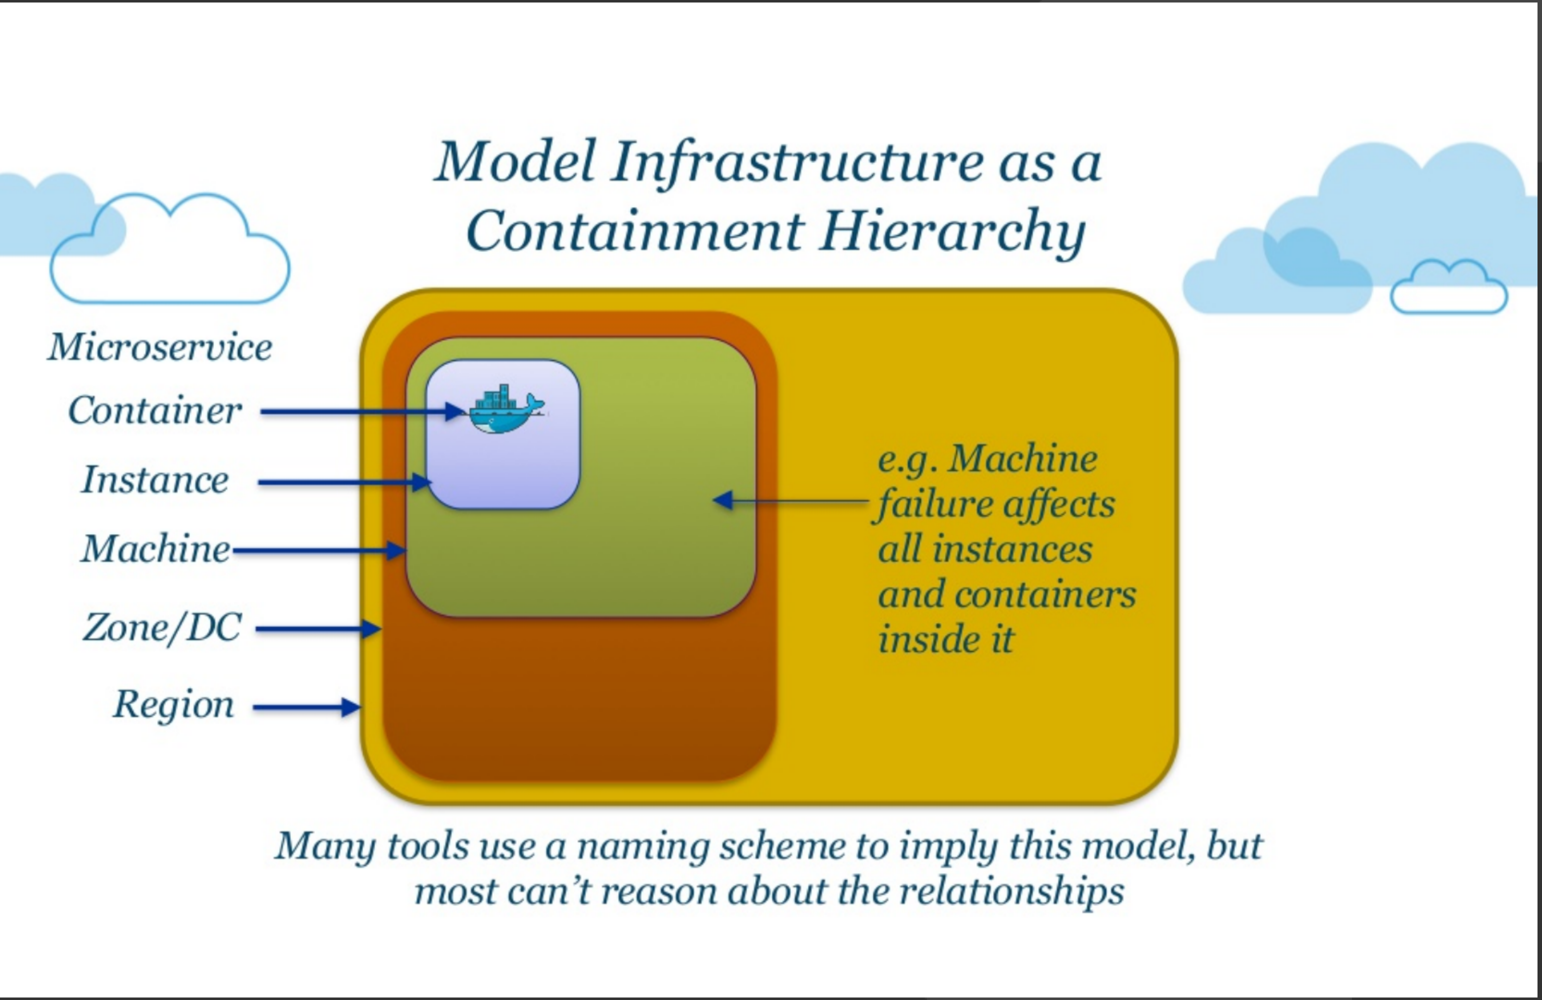

http://www.slideshare.net/adriancockcroft/gluecon-monitoring-microservices-and-containers-a-challenge

Key features

4
transcience

Key features

5

DRY

Key insights.

Key insights

1

remove hassle wrt improving signal/noise

Key insights

1

remove hassle wrt improving signal/noise

- ongoing maintenance & tuning is critical

Key insights

1

remove hassle wrt improving signal/noise

- ongoing maintenance & tuning is critical
- code for UI and logic > knobs

Key insights

1
remove hassle wrt improving signal/noise

- ongoing maintenance & tuning is critical
- code for UI and logic > knobs
- leveraging additional data

Key insights

2
communication

- Author to recipient

Key insights

2
communication

- Author to recipient
- Alert often primary UI

Key insights

3
Human > computer

Key insights

4

attention is scarce, expensive

Key insights

4

attention is scarce, expensive

"provide monitoring platform that enables operators to efficiently utilize their attention"

fault detection

with **bosun**

Classify series & find KPI’s

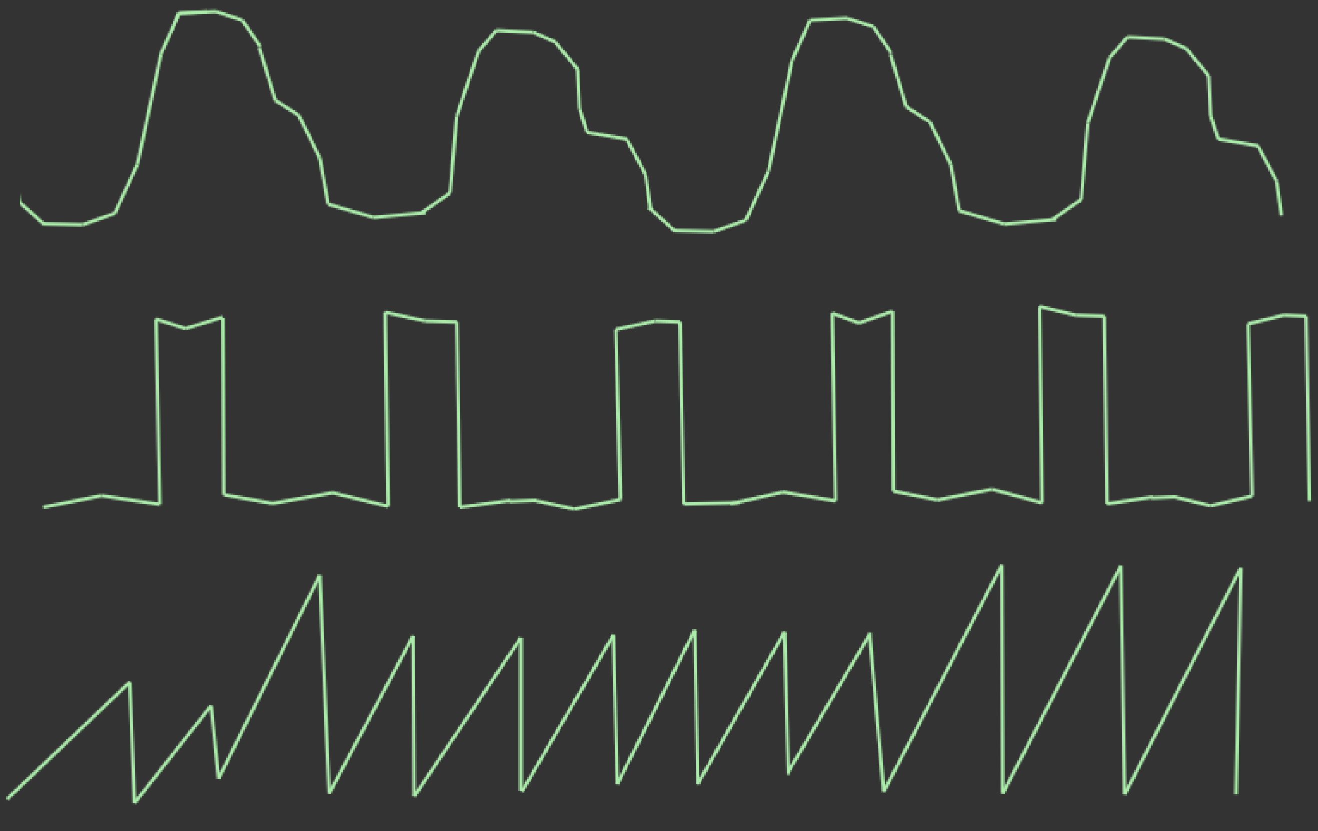

Smoothly seasonal: good

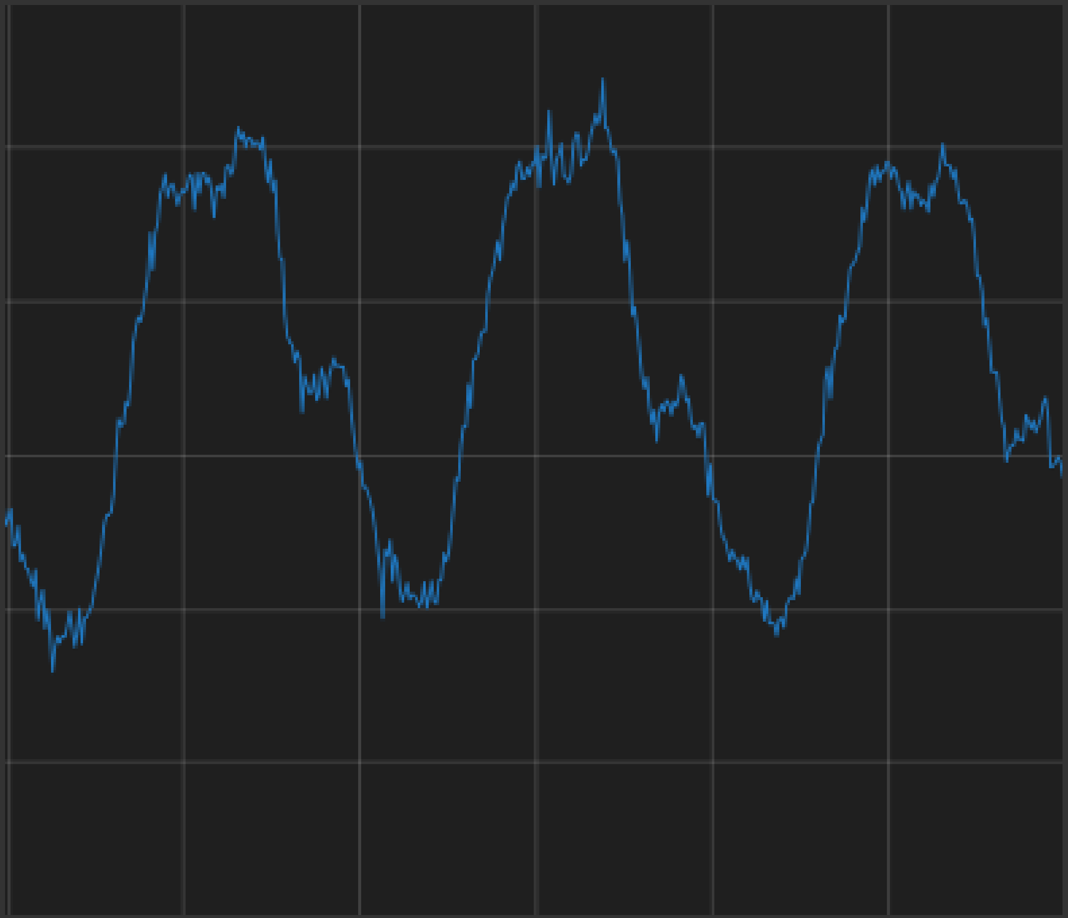

Smoothly seasonal: offset

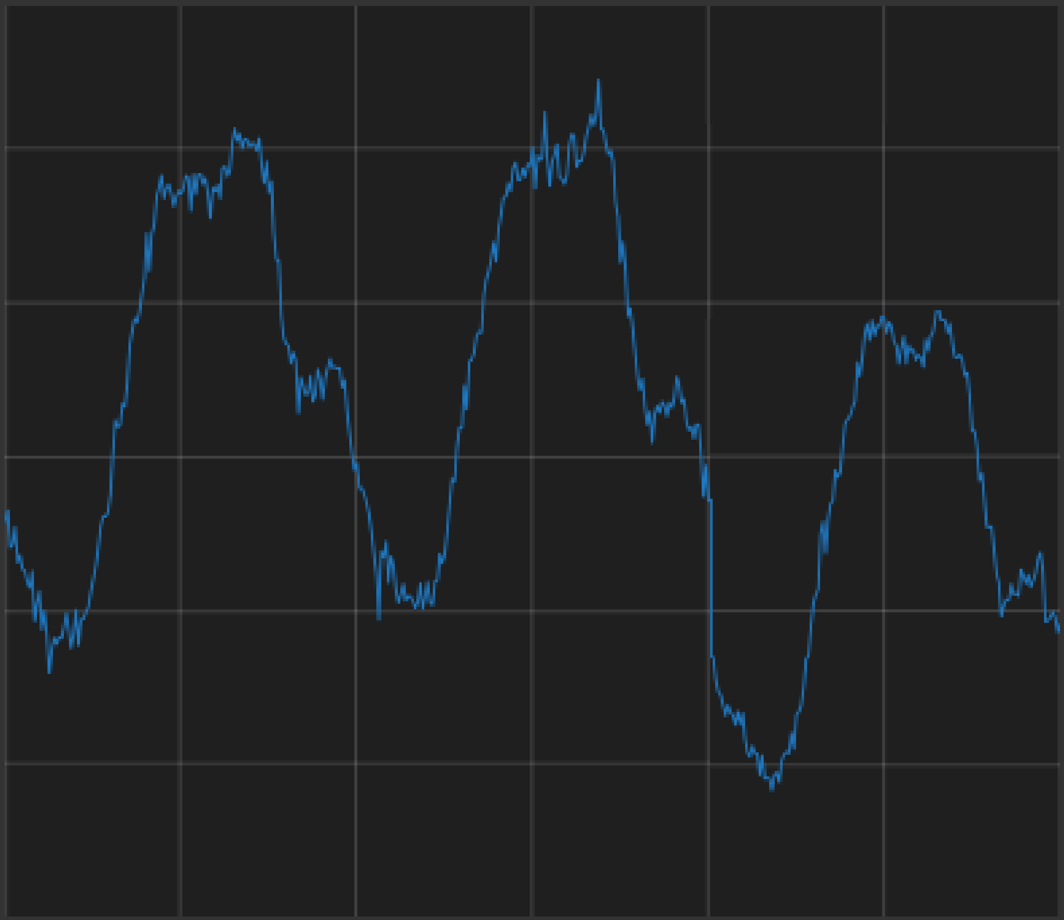

Smoothly seasonal: spikes

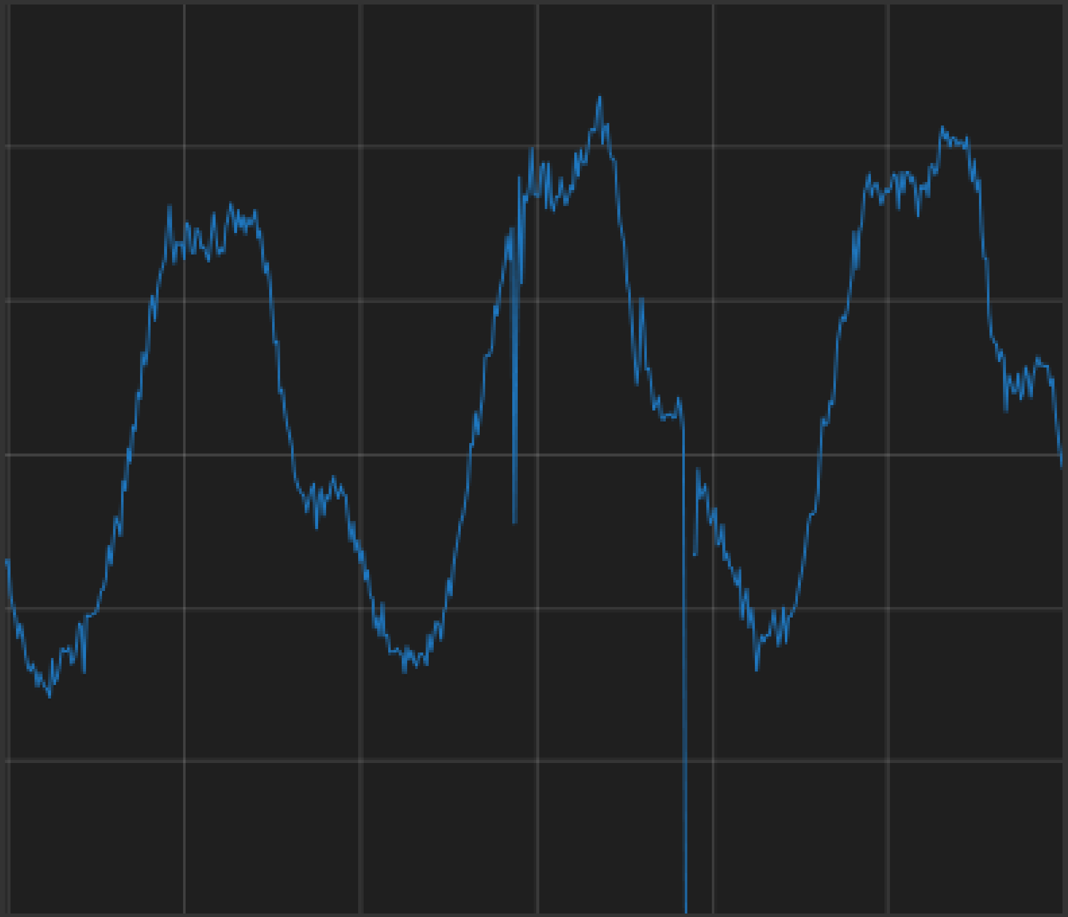

Smoothly seasonal: erratic

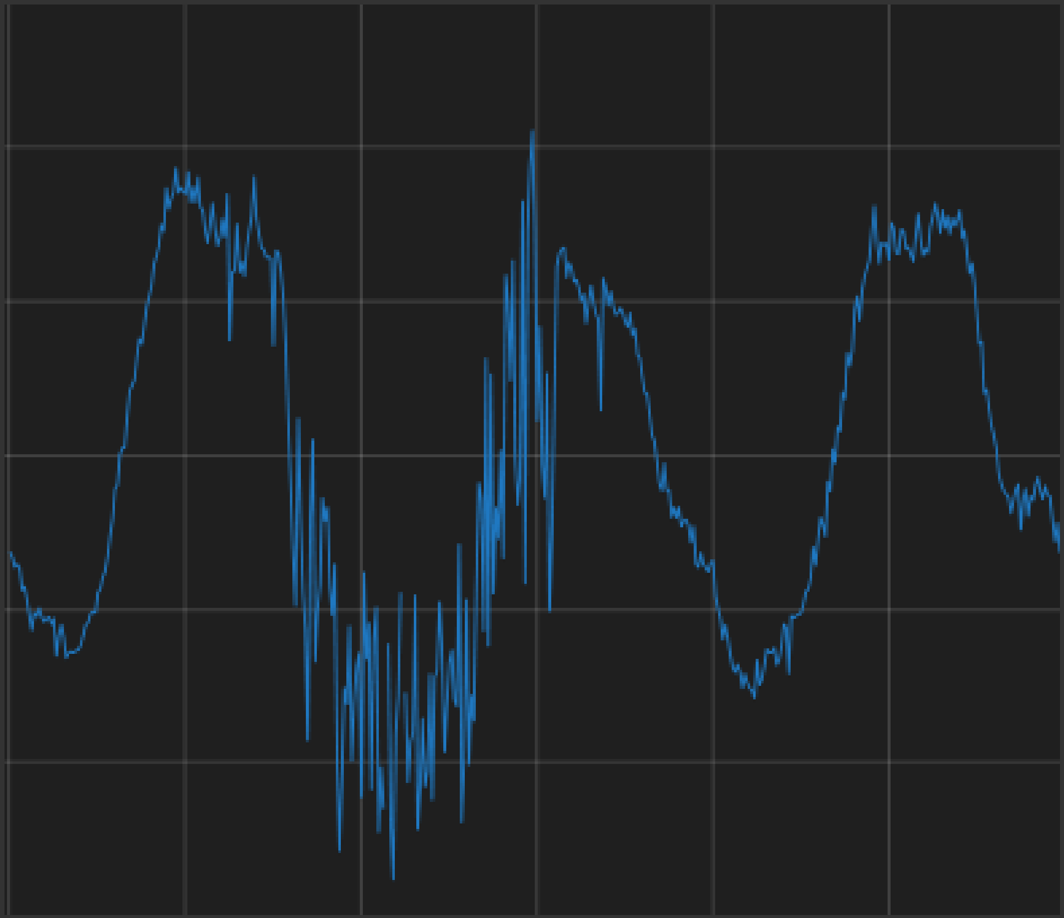

Solution 1/2 : strength

Band(), graphiteBand()

bosun.org/expressions.html

Solution 1/2 : strength

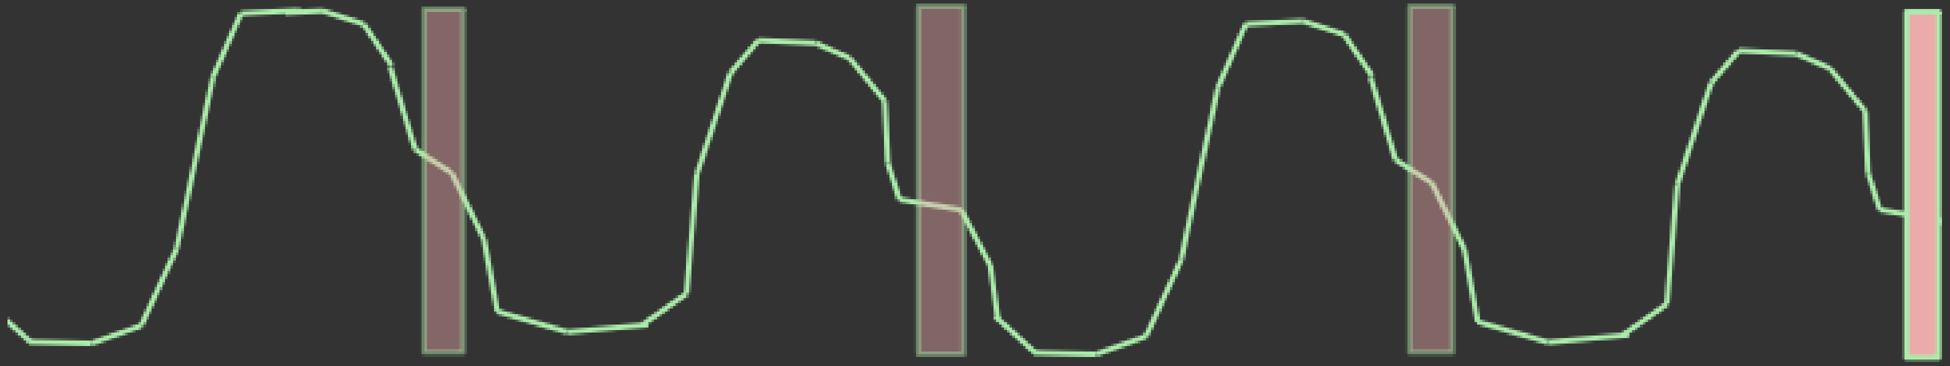

Solution 2/2 : erraticness

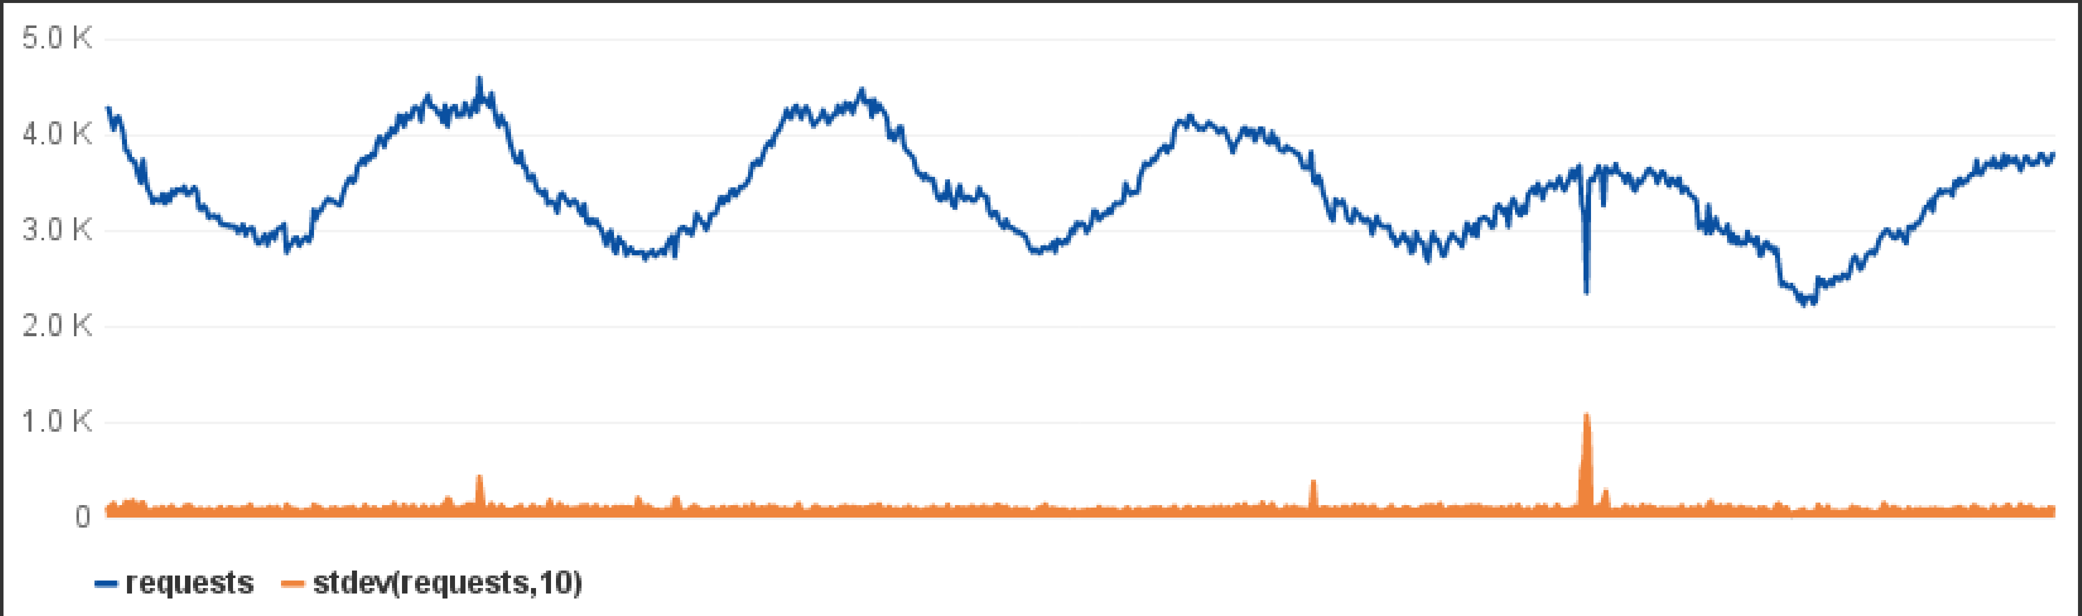

Solution 2/2 : erraticness

$$\textbf{Erraticness now} = \frac{\text{Deviation-now}}{\text{Deviation-historical}}$$

Solution 2/2 : erraticness

$$\textbf{Erraticness now} = \frac{\text{Deviation-now}}{\text{Deviation-historical}} * \frac{\text{median-historical}}{\text{median-now}}$$

Solution 2/2 : erraticness

$$\textbf{Erraticness now} = \frac{\text{deviation-now} * \text{median-historical}}{(\text{deviation-historical} * \text{median-now}) + 0.01}$$

[Acknowledge alert]

Web requests, Global

Total amount now should not be much less than in the past
7d ago **12485**
Now **13102**

Erraticness of Web requests, Global

Erraticness - increased deviation - could be indicative of a spike or drop. Low values up to 6 are ok. 9 are critical
9.410244004576379

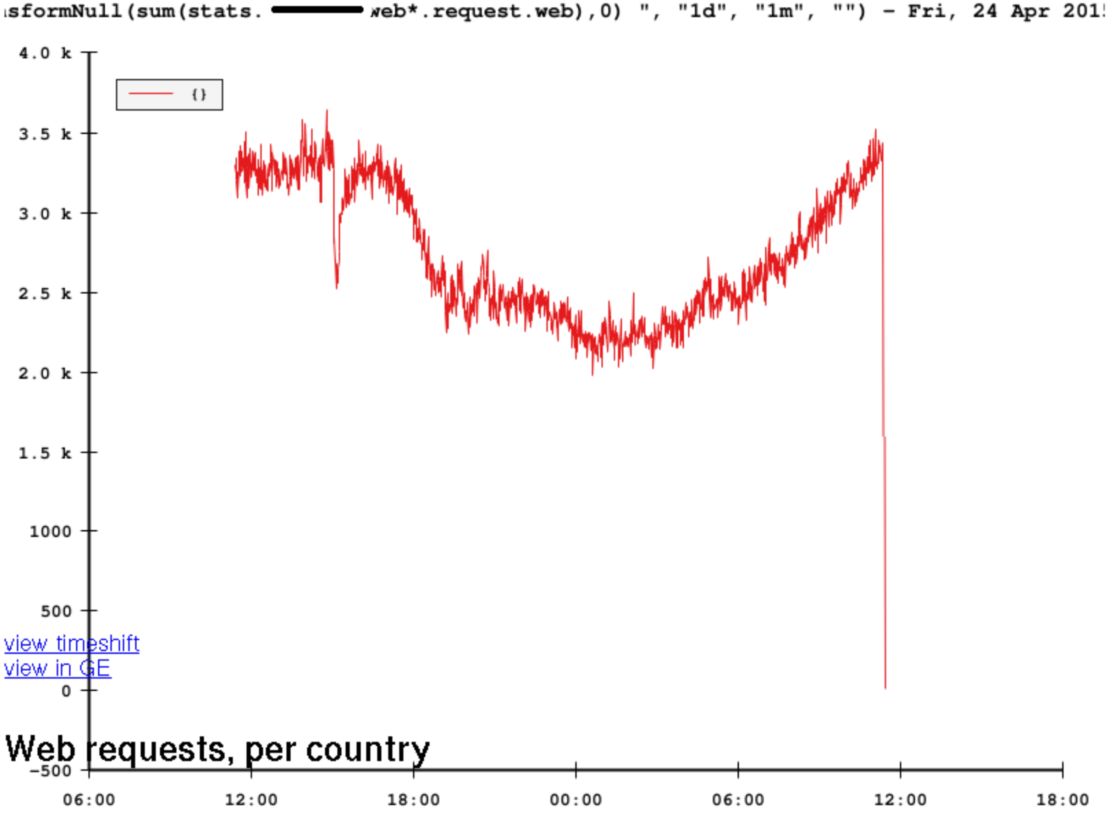

view timeshift
view in GE

Web requests, per country

Web requests, per country

GE graph

median diff lower than -5 is bad (in red).

country	importance	7d ago (med +- dev)	now (med +- dev)	med diff (in devs)		
DE		1455 +- 69	1343 +- 255	-2	view timeshift	view in GE
FR		158 +- 20	173 +- 36	1	view timeshift	view in GE
IT		155 +- 23	148 +- 32	-0	view timeshift	view in GE
NL		117 +- 16	117 +- 25	0	view timeshift	view in GE
US		95 +- 12	97 +- 22	0	view timeshift	view in GE
GB		93 +- 14	92 +- 22	-0	view timeshift	view in GE
TR		90 +- 18	77 +- 20	-1	view timeshift	view in GE
KR		80 +- 14	103 +- 25	2	view timeshift	view in GE
JP		70 +- 12	63 +- 16	-1	view timeshift	view in GE
MX		53 +- 13	48 +- 13	-0	view timeshift	view in GE
JP		53 +- 10	50 +- 13	-0	view timeshift	view in GE
KR		48 +- 10	45 +- 13	-0	view timeshift	view in GE
JP		40 +- 9	33 +- 10	-1	view timeshift	view in GE
MX		35 +- 8	33 +- 9	-0	view timeshift	view in GE

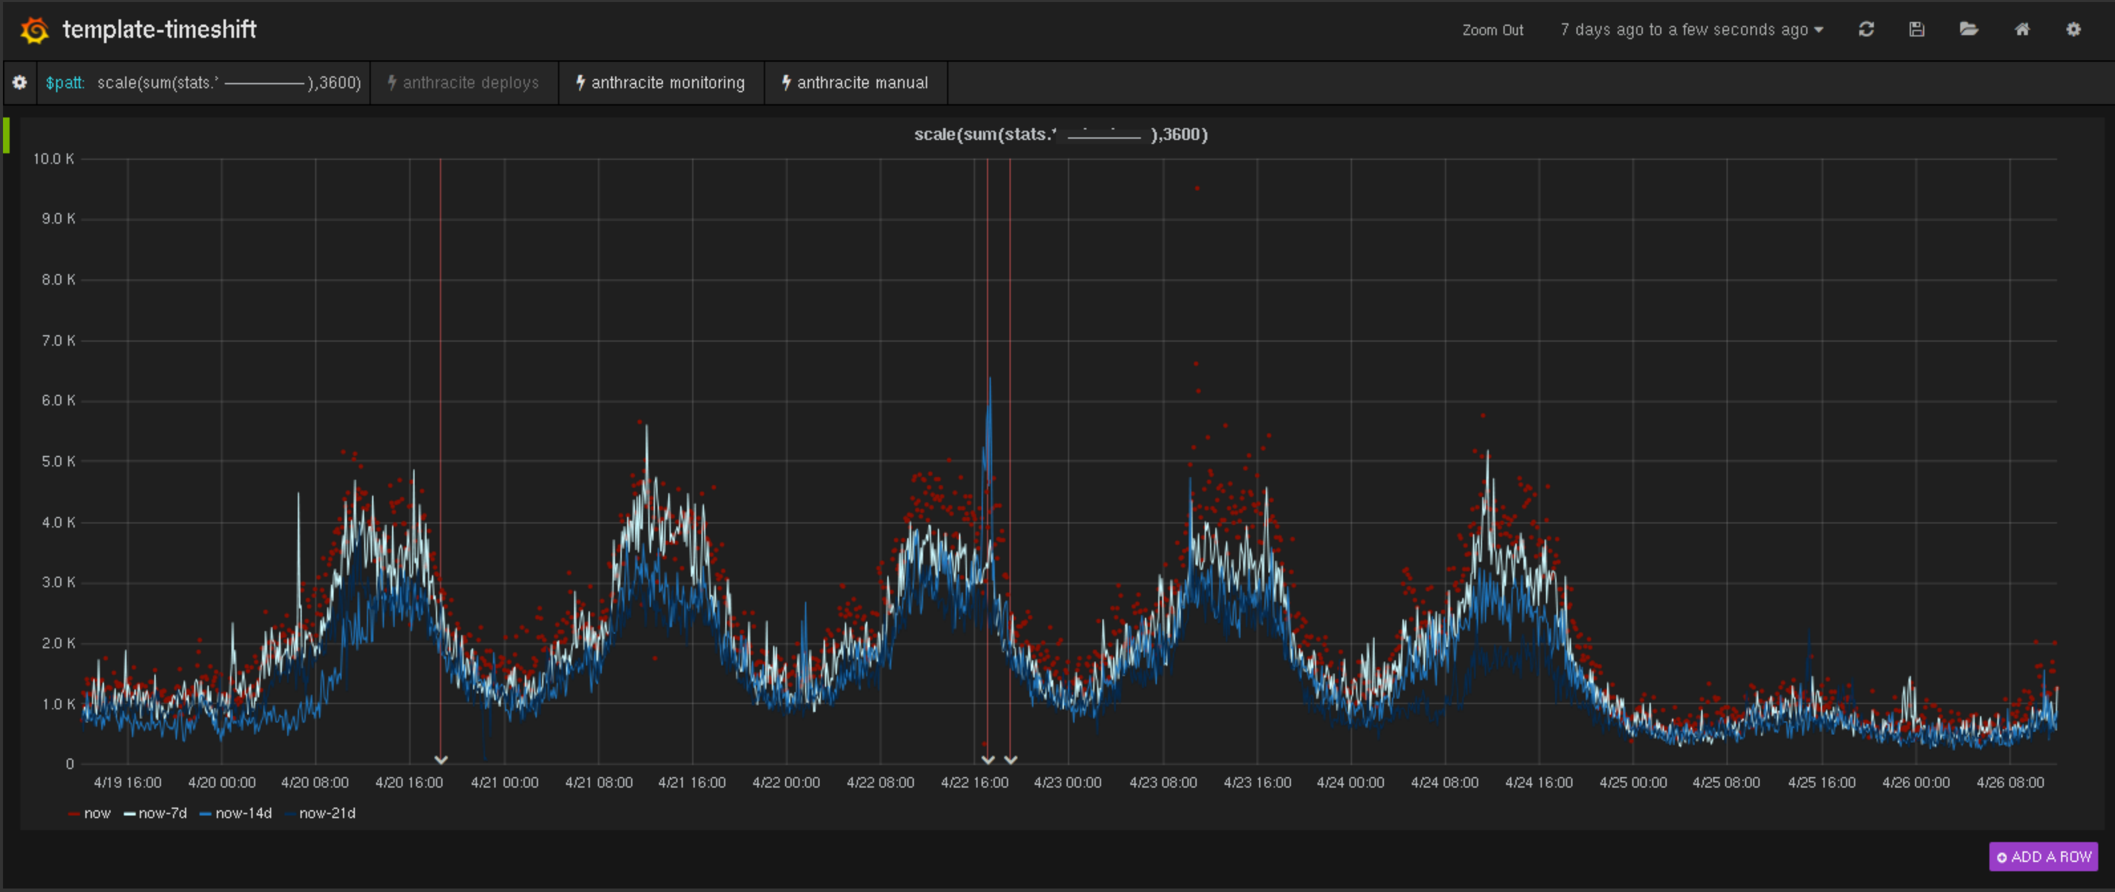
template-timeshift
Zoom Out
7 days ago to a few seconds ago
$patt: scale(sum(stats.'),3600)
anthracite deploys
anthracite monitoring
anthracite manual
scale(sum(stats.'),3600)
10.0 K
9.0 K
8.0 K
7.0 K
6.0 K
5.0 K
4.0 K
3.0 K
2.0 K
1.0 K
0
4/19 16:00
4/20 00:00
4/20 08:00
4/20 16:00
4/21 00:00
4/21 08:00
4/21 16:00
4/22 00:00
4/22 08:00
4/22 16:00
4/23 00:00
4/23 08:00
4/23 16:00
4/24 00:00
4/24 08:00
4/24 16:00
4/25 00:00
4/25 08:00
4/25 16:00
4/26 00:00
4/26 08:00
now
now-7d
now-14d
now-21d
ADD A ROW

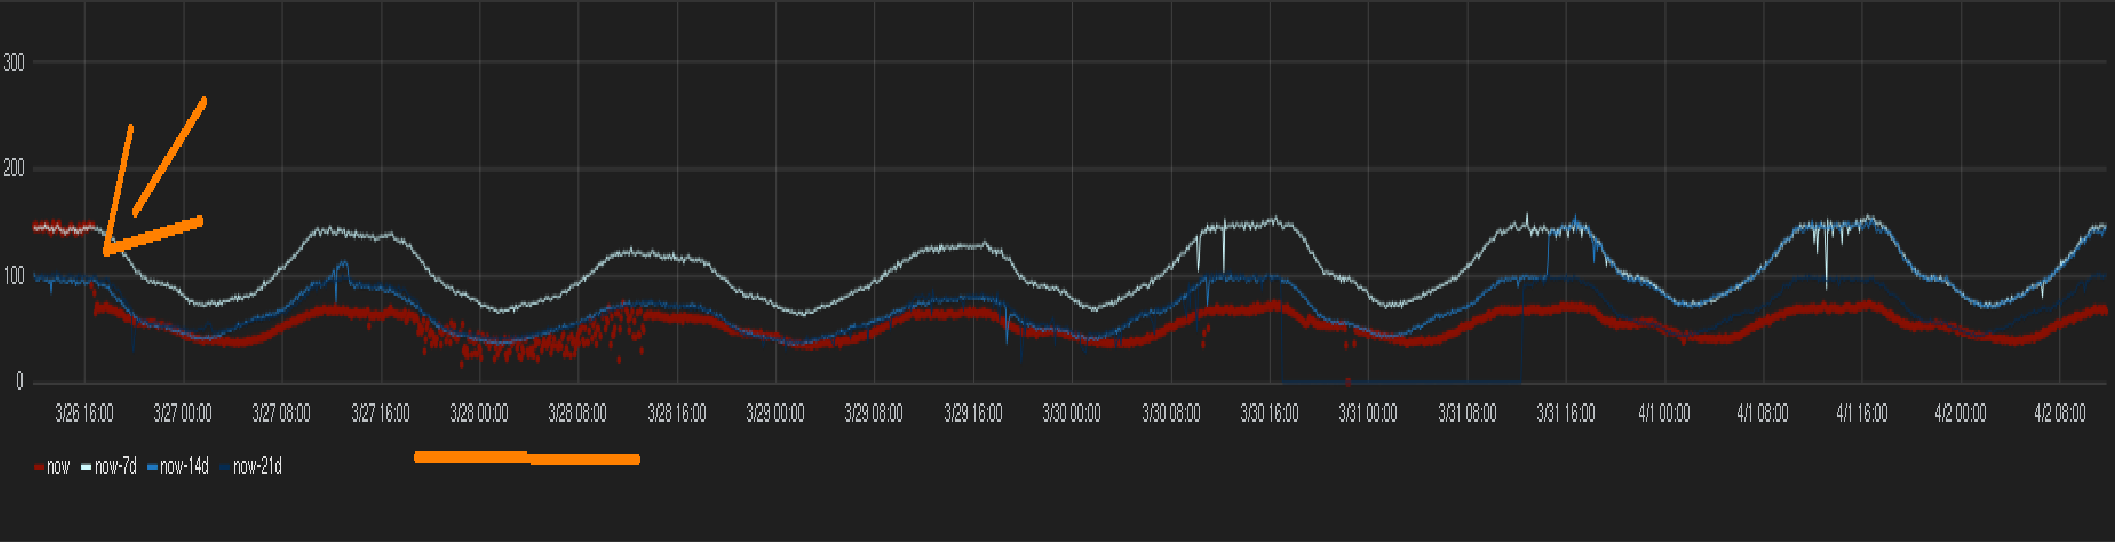

300
200
100
0
3/26 16:00
3/27 00:00
3/27 08:00
3/27 16:00
3/28 00:00
3/28 08:00
3/28 16:00
3/29 00:00
3/29 08:00
3/29 16:00
3/30 00:00
3/30 08:00
3/30 16:00
3/31 00:00
3/31 08:00
3/31 16:00
4/1 00:00
4/1 08:00
4/1 16:00
4/2 00:00
4/2 08:00
now
now-7d
now-14d
now-21d

dieter.plaetinck.be/post/practical-fault-detection-on-timeseries-part-2

More details

Bosun macro, template & example

Grafana dashboard

Static thresholds → automated anomaly detection

- Not scaling / too much data

Static thresholds → automated anomaly detection

- Not scaling / too much data
- Infrastructure complexity

Static thresholds → automated anomaly detection

- Not scaling / too much data
- Infrastructure complexity
- Alerting on patterns

Conclusion

- All about the workflow

- All about the workflow
- An IDE like bosun exponentially boosts ability to maintain high signal/noise alerting

- All about the workflow
- An IDE like bosun exponentially boosts ability to maintain high signal/noise alerting
- Build & share!

Want more ?

- bosun.org/resources presentations by Kyle Brandt (LISA 2014 + Monitorama 2015)
- "my philosophy on alerting" by Rob Ewaschuk
- kitchensoap.com/2015/05/01/openlettertomonitoringproducts
- kitchensoap.com/2013/07/22/owning-attention-considerations-for-alert-design
- "monitoring microservices" by Adrian Cockroft
- (dieter.plaetinck.be/post/practical-fault-detection-alerting-dont-need-to-be-data-scientist)
- dieter.plaetinck.be/post/practical-fault-detection-on-timeseries-part-2
- metrics20.org/media
- mabrek.github.io
- iwringer.wordpress.com