

Dia dhuit

David McKay
@rawkode

Head of Developer Relations
@InfluxDB | #InfluxDB



DublinPHP



Introduction to Time Series

Before we begin ...




Pop Quiz

“Invented” When?


Encoding

First Used 410 BC



Encoding

“Documented” in *The Lives of the Noble Grecians and Romans*, by Roman historian Plutarch.






Alcibiades suddenly raised the **Athenian ensign** in the admiral ship, and fell upon those galleys of the Peloponnesians ...


Encoding

In the 14th century, things hadn't actually advanced much more. The *Black Book of Admiralty* listed 2 signals: 1 flag or 2 flags




Encoding

By the 15th century there were 15 flags, each with a single meaning.




Encoding

Finally, in the late 17th century; a French system existed (Mahé de la Bourdonnais) with 10 coloured flags, representing 0-9




Sharding

First Used 150 BC



Sharding

First “documented” example was in ~150 AD, invented and described by Polybius.





We take the **alphabet** and
divide it into **five parts**,
each consisting of **five**
letters.



The Early History of Data Networks



Gerard J. Holzmann • Ryan Whison



History of Time Series

The Romans Did It

The earliest form of a company which issued public shares was the case of the publicani during the Roman Republic.





Like modern joint-stock companies, the publicani were **legal bodies** independent of their members whose ownership was divided into shares, or partes. There is evidence that these shares were **sold to public investors and traded** in a type of over-the-counter market in the Forum, near the Temple of Castor and Pollux. The **shares fluctuated in value**, encouraging the activity of speculators, or quaestors.

In 1602 ...

First IPO: Dutch East India Company

In 1873 ...

First US IPO: Bank of North America

In 1884 ...

What was the price of wheat?

First Documented Time Series

A Comparison of the
Fluctuations in the Price
of Wheat and in the
Cotton and Silk Imports
into Great Britain


J. H. Poynting

Journal of the Statistical
Society of London

Vol. 47, No. 1 (Mar.,
1884), pp. 34-74

What is all this?

This is the first (or one of)
paper that added the
dimension of time to
statistical mathematics





Most data is best understood in the dimension of time

@pauldix, CTO





Introduction to Time Series

What Will We Cover?

- Time Series Data
- Time Series Databases
- Getting to Know InfluxDB
- Value of Time Series Data
- Advancing Monitoring to Time Series

Time Series Data

What is it?

Time Series Data

Data with a timestamp



CPU = 88%

Healthcheck
Failed

DB Migration Run

Scotland Qualify
for World
Cup

Pod Killed by OOM

Pod Restarted

Git Commit

v1.1.3 Deployed

CPU = 12%

CI Passed

CI Started

CPU = 88%

Healthcheck
Failed

DB Migration Run

Scotland Qualify
for World
Cup

Pod Killed by OOM

Pod Restarted

Git Commit

v1.1.3 Deployed

CPU = 12%

CI Passed

CI Started

CPU = 88%

Healthcheck
Failed

DB Migration Run

Scotland Qualify
for World
Cup

Pod Killed by OOM

Pod Restarted

Git Commit

v1.1.3 Deployed

CPU = 12%

CI Passed

CI Started

CPU = 88%

Healthcheck
Failed

DB Migration Run

Scotland Qualify
for World
Cup

Pod Killed by OOM

Pod Restarted

Git Commit

v1.1.3 Deployed

CPU = 12%

CI Passed

CI Started

CPU = 88%

Healthcheck
Failed

DB Migration Run

Scotland Qualify
for World
Cup

Pod Killed by OOM

Pod Restarted


Git Commit

v1.1.3 Deployed

CPU = 12%

CI Passed

CI Started



What is Time Series Data?



What is Time Series Data?

Regular (Metrics)


Irregular (Events)

- Predictable
- Evenly Distributed

- Unpredictable
- Inconsistent Intervals


Regular / Metrics

- ★ CPU Usage
- ★ Memory Usage
- ★ Ping Time for Google.com
- ★ Number of Processes



Irregular / Events

- ★ User Clicked Login
- ★ Authentication Failed
- ★ CI Published v1.3.1
- ★ Network Cable Unplugged



```
e.quicklook
Apr 10 16:17:33 p4x-639 syslogd[41]: ASL Sender Statistics
Apr 10 16:28:05 p4x-639 com.apple.xpc.launchd[1] (com.apple.quicklook[2592]): Endpoint has been activated through legacy launch(3) APIs. Please switch to XPC or bootstrap_check_in(): com.apple.quicklook
Apr 10 16:28:05 p4x-639 syslogd[41]: ASL Sender Statistics
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.802 ksinstall[[2618/0x109ba05c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer started.
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.807 ksinstall[[2618/0x109ba05c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer starting Installation.
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.807 ksinstall[[2618/0x109ba05c0] [lvl=2] -[KeystoneInstallBackend install] Google Software Update attempting to install.
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.954 ksinstall[[2618/0x109ba05c0] [lvl=2] +[KSCodeSigningVerification verifyBundle:applicationId:error:] KSCodeSigningVerification verifying code signing for '/Users/rawkode/Library/Google/GoogleSoftwareUpdate/GoogleSoftwareUpdate.bundle' with the requirement 'anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] exists and certificate leaf[field.1.2.840.113635.100.6.1.13] exists and certificate leaf[subject.OU]="EQHXZ8M8AV" and (identifier="com.google.Keystone)''
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.976 ksinstall[[2618/0x109ba05c0] [lvl=2] -[KeystoneInstallBackend shouldInstallWithVersion:error:] Google Software Update found version that is the same or newer: 1.2.13.37.
Apr 10 16:34:50 p4x-639 ksinstall[2618]: 2019-04-10 16:34:50.976 ksinstall[[2618/0x109ba05c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer ran successfully.
Apr 10 16:36:00 p4x-639 zoom.us[2624]: objc[2624]: Class ZPLogHelperImp is implemented in both /Applications/zoom.us.app/Contents/Frameworks/zChatApp.bundle/Contents/MacOS/zChatApp (0x10cbe7e98) and /Applications/zoom.us.app/Contents/Frameworks/zChatUI.bundle/Contents/MacOS/zChatUI (0x10d43ec68). One of the two will be used. Which one is undefined.
Apr 10 16:36:00 p4x-639 zoom.us[2624]: objc[2624]: Class ZPLogHelperNull is implemented in both /Applications/zoom.us.app/Contents/Frameworks/zChatApp.bundle/Contents/MacOS/zChatApp (0x10cbe7ee8) and /Applications/zoom.us.app/Contents/Frameworks/zChatUI.bundle/Contents/MacOS/zChatUI (0x10d43ecb8). One of the two will be used. Which one is undefined.
Apr 10 16:36:01 p4x-639 zoom.us[2624]: objc[2624]: Class ZSearchTextField is implemented in both /Applications/zoom.us.app/Contents/Frameworks/zChatUI.bundle/Contents/MacOS/zChatUI (0x10d441238) and /Applications/zoom.us.app/Contents/Frameworks/zVideoUI.bundle/Contents/MacOS/zVideoUI (0x117fa37b0). One of the two will be used. Which one is undefined.
Apr 10 16:36:01 p4x-639 zoom.us[2624]: objc[2624]: Class ZPLogHelperImp is implemented in both /Applications/zoom.us.app/Contents/Frameworks/zChatApp.bundle/Contents/MacOS/zChatApp (0x10cbe7e98) and /Applications/zoom.us.app/Contents/Frameworks/zVideoUI.bundle/Contents/MacOS/zVideoUI (0x117fa3b98). One of the two will be used. Which one is undefined.
Apr 10 16:36:01 p4x-639 zoom.us[2624]: objc[2624]: Class ZPLogHelperNull is implemented in both /Applications/zoom.us.app/Contents/Frameworks/zChatApp.bundle/Contents/MacOS/zChatApp (0x10cbe7ee8) and /Applications/zoom.us.app/Contents/Frameworks/zVideoUI.bundle/Contents/MacOS/zVideoUI (0x117fa3be8). One of the two will be used. Which one is undefined.
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.738 ksinstall[[2653/0x1093385c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer started.
Apr 10 16:39:29 p4x-639 syslogd[41]: ASL Sender Statistics
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.744 ksinstall[[2653/0x1093385c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer starting Installation.
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.745 ksinstall[[2653/0x1093385c0] [lvl=2] -[KeystoneInstallBackend install] Google Software Update attempting to install.
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.923 ksinstall[[2653/0x1093385c0] [lvl=2] +[KSCodeSigningVerification verifyBundle:applicationId:error:] KSCodeSigningVerification verifying code signing for '/Users/rawkode/Library/Google/GoogleSoftwareUpdate/GoogleSoftwareUpdate.bundle' with the requirement 'anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] exists and certificate leaf[field.1.2.840.113635.100.6.1.13] exists and certificate leaf[subject.OU]="EQHXZ8M8AV" and (identifier="com.google.Keystone)''
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.954 ksinstall[[2653/0x1093385c0] [lvl=2] -[KeystoneInstallBackend shouldInstallWithVersion:error:] Google Software Update found version that is the same or newer: 1.2.13.37.
Apr 10 16:39:29 p4x-639 ksinstall[2653]: 2019-04-10 16:39:29.955 ksinstall[[2653/0x1093385c0] [lvl=2] -[KeystoneInstallTool main] Google Software Update installer ran successfully.
Apr 10 16:39:47 p4x-639 com.apple.xpc.launchd[1] (com.google.GoogleDrive.FinderSyncAPIExtension.5E4838BB-E54C-45E9-90EB-566F399D890A[2657]): Caller wanted oneshot behavior for pre-existing in stance: caller = Finder
Apr 10 16:39:47 p4x-639 com.apple.xpc.launchd[1] (com.google.GoogleDrive.FinderSyncAPIExtension.71C22EDF-227B-4190-AC56-A9C70CA7F936[2659]): Caller wanted oneshot behavior for pre-existing in stance: caller = Slack
Apr 10 16:39:47 p4x-639 com.apple.xpc.launchd[1] (com.google.GoogleDrive.FinderSyncAPIExtension.00555BD1-8478-4EED-BDEA-DEB68950CE46[2658]): Caller wanted oneshot behavior for pre-existing in stance: caller = Google Chrome
Apr 10 16:39:49 p4x-639 com.apple.xpc.launchd[1] (com.apple.FolderActionsDispatcher): This service is defined to be constantly running and is inherently inefficient.
Apr 10 16:51:51 p4x-639 syslogd[41]: ASL Sender Statistics
Apr 10 16:57:02 p4x-639 com.apple.xpc.launchd[1] (com.apple.quicklook[2716]): Endpoint has been activated through legacy launch(3) APIs. Please switch to XPC or bootstrap_check_in(): com.apple.quicklook
Apr 10 16:57:02 p4x-639 com.apple.xpc.launchd[1] (com.google.GoogleDrive.FinderSyncAPIExtension.AF7987D0-4122-489F-8263-9247374F5EBF[2714]): Caller wanted oneshot behavior for pre-existing in stance: caller = Google Chrome
Apr 10 16:57:02 p4x-639 com.apple.xpc.launchd[1] (com.google.GoogleDrive.FinderSyncAPIExtension.0774F172-95E9-44BB-91E4-B5861BB573FE[2715]): Caller wanted oneshot behavior for pre-existing in stance: caller = Finder
Apr 10 17:02:53 p4x-639 syslogd[41]: ASL Sender Statistics
```



ken UEFA CHAMPIONS LEAGUE Heineken Enjoy Responsibl Heineken



54:51

WIJNALDUM 54'
ORIGI 7'



LIVERPOOL

2 | 0

BARCELONA



AGGREGATE 2-3

Collecting Metrics & Events

With Prometheus Exporters or Telegraf

Collecting Metrics & Events

Inputs:

- CloudWatch
- Elasticsearch
- Kafka
- Jenkins
- Kubernetes
- Linux
- Puppet
- Windows
- x509

Outputs:

- CloudWatch
- Kafka
- DataDog
- Elasticsearch
- Graphite
- Prometheus

Exporters:

- Atlassian
- Ceph
- Consul
- Kubernetes
- Memcached
- MySQL

PUSH



PULL



Push AND Pull

Metrics are pulled at a regular interval

Events NEED to be pushed as they happen

Consistent and reliable intervals

Inconsistent intervals



Time Series Data

Use Cases

Use Cases for Time Series

Monitoring

- Infrastructure
- Applications
- Third Party Services

IoT / Sensor

- Thermostats
- Electric Engines
- Smart Things
- GPS
- Fitbits

Real Time Analytics

- Website Tracking
- Stock Prices
- Currency Exchange Rates

Time Series Databases

TSDB's


Time Series Databases

Time Series databases are optimized for collecting, storing, retrieving, and processing of Time Series data.

Time Series Databases


- High Write Frequency
- Reads are range scans
- TTL / Lifecycle Management
- Time Sensitive

Trend of the last 24 months



12%

Are you in the 88%?



[adult swim]

Analyst vetted. Customer approved.

We're proud to be recognized as a Leader in Gartner's 2019 Magic Quadrant for APM Suites and a 2019 Gartner Customers' Choice as the highest peer-rated APM vendor. Get your complimentary Gartner Magic Quadrant report now.

[Download report](#)[Request demo](#)

A Leader in Gartner's 2019 APM Magic Quadrant >



New Relic Team



What brings you to New Relic today?

Write a reply...



New Relic
INSIGHTS

NRQL > SELECT

Query

Data Explorer

Data Explorer

 Full-stack Visibility

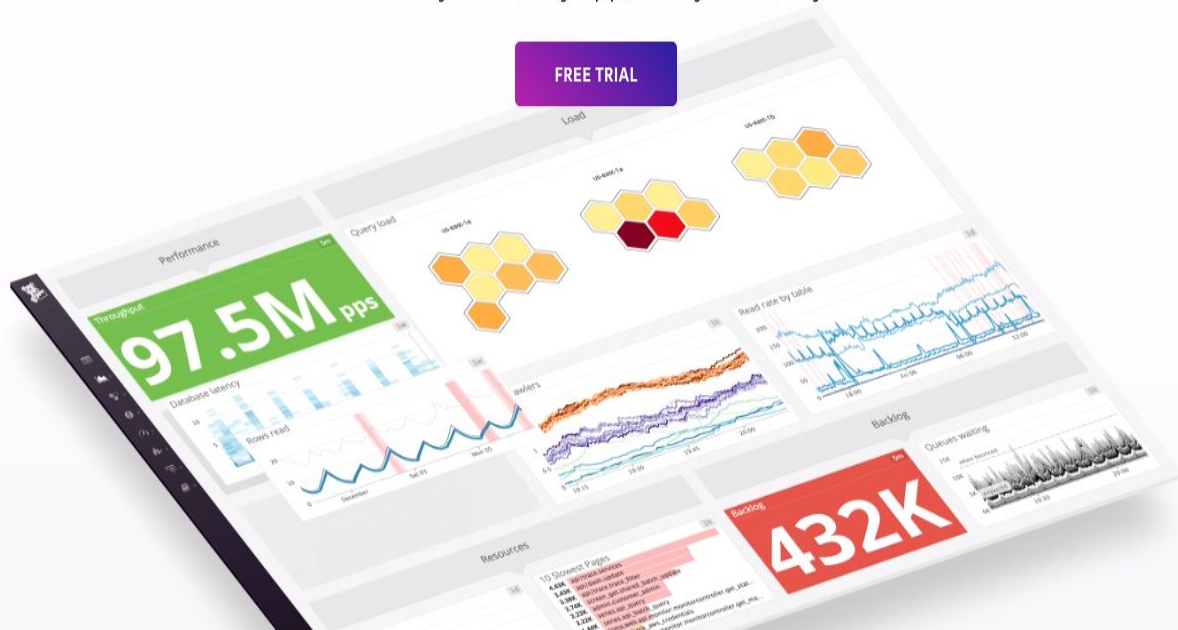
Created by alarson@newrelic.com
Last edited 3/28/18



Modern monitoring & analytics

See inside any stack, any app, at any scale, anywhere.

FREE TRIAL







David McKay
@rawkode



David McKay
@rawkode



I run **#Kubernetes** in production and I monitor it with ...



902 votes • Final results

12:13 PM - 3 Feb 2019

30 Retweets 28 Likes



Edit profile

13%

It's Not Too Late!



influxdb

Disclaimer

Most of this isn't unique to InfluxDB




InfluxDB

Introductions


InfluxDB

- TSDB
- Open-Source
- FullStack (Telegraf, InfluxDB, Chronograf, and Kapacitor)
- v2 ...



Points

At any point in time, this
value was N



Point

- Series
- Fields
- Timestamp

```
load,host=vm1 1m=6.32,5m=8.20,15m=9.55 123456789
```

Series

- `load,host=vm1`
- `stock_price,market=NASDAQ,ticker=G00G`
- `users,service=comments`

- Name
- Tag Keys
- Tag Values

Tags & Fields

Tags

- Indexed
- String Types

Fields

- Not Indexed
- Multiple Data Types

Value of Time Series Data

Isn't It Valuable Forever?

Resolution

The predictable interval at which we will collect our time series data

Value of Time Series Data

The value of all time series data is directly correlated with the resolution that the data is available

Cost of Time Series Data

Wait, Isn't It Free?!

Example

```
cpu,machine=abc1 usage=1.66 timestamp
```

Resolution

- 1 Measurement
- 1 Series
- 1s Resolution

86400
Points
Per Day

Resolution

- 1 Measurement
- 2 Series
- 1s Resolution

172800
Points
Per Day

Resolution

- 5 Measurement
- 10 Series
- 1s Resolution

4320000
Points
Per Day

Nasdaq

- 1 Measurement
- 3300 Series
- 1ms Resolution

28512000
0000
Points
Per Day

Nasdaq

- 1 Measurement
- 3300 Series
- 1m Resolution

4752000
Points
Per Day

Nasdaq

- 1 Measurement
- 3300 Series
- 1h Resolution

79200
Points
Per Day

Nasdaq

- 1 Measurement
- 3300 Series
- 6h Resolution

13200
Points
Per Day

Rollups

Lowering the Resolution

Rollups with Continuous Queries

```
CREATE CONTINUOUS QUERY "rollup_1h" ON "nasdaq"  
  
BEGIN  
    SELECT mean(price) INTO yearly FROM weekly  
        GROUP BY time(1h)  
END
```

Events?


Outlier / Anomaly Detection



InfluxDB Anomaly Detection

Advancing Monitoring to Time Series

Taking Small Steps for Giant Leaps



CPU > 80%

Application

MEM > 80%

Response
Time >
300ms

Database

Black
Friday

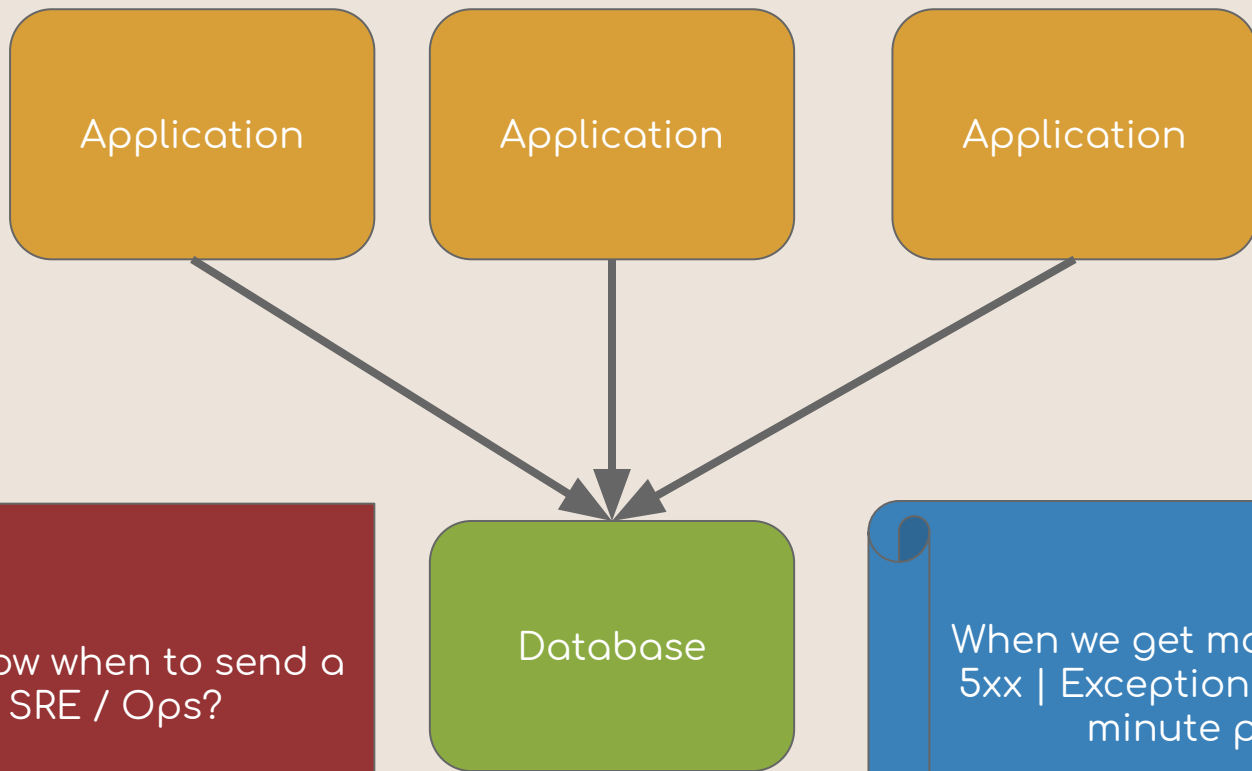
How do we know when to send a page to SRE / Ops?

Application



Database

When the application fails the health-check



How do we know when to send a page to SRE / Ops?

When we get more than 100 [5xx | Exceptions] within a 5 minute period

Service A

Service B

Service B

Service C

Canary

Virtual Network

Service Mesh

Ummm?

Database A

Database B

Database C

Cloud Native Architectures

Convenience Vs. Cost

You can treat the symptoms for a while ...

Upgrade Your Monitoring



Causality

Treating the Disease

Causality

- Look at last weeks, months, and years of data
- Use tags to build correlation
- Get Statistical
 - ◆ INTEGRAL()
 - ◆ LINEAR_PREDICTION()
 - ◆ DERIVATIVE()
 - ◆ MOVING_AVERAGE()
 - ◆ HOLT_WINTERS()

Causality

Have you ever been paged at 4am because the disk usage of a machine went above 85%?

Could this have been determined during office hours?
(Linear Growth)

Can we use correlations to determine the cause during anomalies?

Causality

In our distributed application, our p99 reports that our users are being served healthy responses in under 2ms. Our pager is going off because we're getting too many exceptions in the code

```
SELECT mode(*) FROM logs;
```

Proactive Ops

We run Big News Corp and we need to reduce our cloud costs. Instead of running at 30% utilisation, can we run at 80% utilisation?

HOLT_WINTERS

Build Automation

Through Causality, Historical Data, Prediction, and ML

Summary

- Use a TSDB
- Understand Cost / Select Tags Wisely
- Understand the resolution you need for 1m, 6m, > 12m
- Rollup metrics
- Perform outlier detection on events
- Build automation, dashboarding, and reporting around your data (past, present, and future)

Cheers!

David McKay
@rawkode

Head of Developer Relations
@InfluxDB | #InfluxDB

That's All
Folks!

