

Real-time-first metrics

Aurelijus Banelis



Aurelijus Banelis

Software developer
aurelijus.banelis.lt
aurelijus@banelis.lt

PGP public key rsa2048/ **539B6203**
Key fingerprint = 130D C446 1F1A 2E50 D6E3
 3DA8 3202 05E7 539B 6203





Real-time-first metrics





11
23
X

40

35

30

25

Real time metrics

Real time metrics

in software development

Intro

What are metrics

Compare

**Principles, tradeoffs,
tools and added value**

Demo

**How does it feel to do
Real-time first**

Intro

What are metrics

Compare

**Principles, tradeoffs,
tools and added value**

Demo

**How does it feel to do
Real-time first**

Via example

📖 README.md

Simple downloader from AWS Glacier

Could not find simple utility to download archive from AWS Glacier to my PC, so written my own.

Usage

```
aws configure
aws glacier describe-job --account-id - --vault-name YOUR_VOULT_NAME --job-id YOUR_LONG_JOB_ID > job.json
./gdown job.json OUTPUT_FILE
```

Downloading...

Downloading...

Error: read tcp

**192.168.0.123->54.239.33.110:443:
read connection reset by peer**

```
buf := make([]byte, 1024*1024*100)
_, err = io.CopyBuffer(
    out, resp.Body, buf)
```

network

Total IPv4 Traffic.



mem

Real memory (RAM) used per user. This does not include shared memory.

Users Real Memory (w/o shared) (users.mem)



Users Disk Writes (users.pwrites)



```
buf := make([]byte, 1024*1024*100)
_, err = io.CopyBuffer(
    out, resp.Body, buf)
```

network

Total IPv4 Traffic

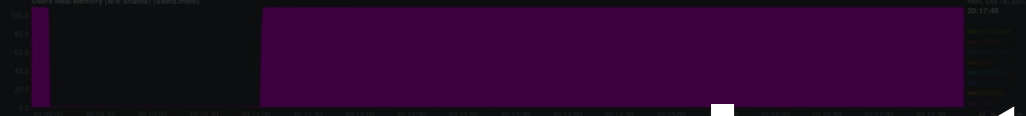
IPv4 Bandwidth (system.ipv4)



Easier debugging

mem

Real memory (RAM) used per user. This does not include shared memory.
Shows Real Memory (not shared) (users.mem)



and actions

```
buf := make([]byte, 1024*1024*100)  
_, err = io.CopyBuffer(  
    out, resp.Body, buf)
```

based on monitoring

Users Disk Writes (users.pwrites)



Data driven development

Real-time

Metrics

Insights

Aggregation

Big data

Data driven development

Correctness

Real-time

Consensus

Validation

Logs

Metrics

Intro

Diagrams to understand
what is really going on

Compare

Principles, tradeoffs,
tools and added value

Demo

How does it feel to do
Real-time first

Intro

What are metrics

Compare

**Principles, trade-offs,
tools and added value**

Demo

How does it feel to do
Real-time first

Decision making

Debugging

Netdata

Decision making

ElasticSearch

CloudWatch

Debugging

Netdata

Decision making


ElasticSearch

CloudWatch

Debugging

Problem

Decision price


<input type="checkbox"/>	▶	MySQL	[redacted]-production	available	<div><div></div></div> 1.67%
<input type="checkbox"/>	▶	MySQL	[redacted]-staging	available	<div><div></div></div> 1.50%
<input type="checkbox"/>	▶	MySQL	[redacted]-mysql-lock	available	<div><div></div></div> 0.67%
	<input type="checkbox"/>	▼	MySQL [redacted] master-db	available	<div><div></div></div> 8.05%


read replication: stopped

[redacted].eu-west-1.rds.amazonaws.com:3306 (**authorized**) 

2 Benefits







netdata alarms

Active All Log

Alarm Log

Search

Event Date

Chart

Alarm

Friendly Value

Status

10:50:39 PM	disk_space_...	out of disk space time	6h	WARNING
10:38:11 PM	net_packets.veth0ec6984	10s received packets storm	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	1m received packets rate	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	inbound packets dropped ratio	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	outbound packets dropped ratio	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	inbound packets dropped ratio	-	UNINITIALIZED
10:38:11 PM	net_packets.veth47f0782	10s received packets storm	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	1m received packets rate	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	outbound packets dropped ratio	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	inbound packets dropped ratio	-	UNINITIALIZED
10:38:09 PM	net_packets.veth1182487	10s received packets storm	-	UNINITIALIZED

Showing 1 to 10 of 1000 rows 10 rows per page

1

2

4

5

...

100

Close

Closer to the source

To automate

netdata alarms

ActiveAllLog

Alarm Log

Search

Refresh

Grid

Print

Event Date	Chart	Alarm	Friendly Value	Status
10:50:39 PM	disk_space_...	out of disk space time	6h	WARNING
10:38:11 PM	net_packets.veth0ec6984	10s received packets storm	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	1m received packets rate	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	outbound packets dropped ratio	-	UNINITIALIZED
10:38:11 PM	net_packets.veth0ec6984	inbound packets dropped ratio	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	10s received packets storm	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	1m received packets rate	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	outbound packets dropped ratio	-	UNINITIALIZED
10:38:10 PM	net_packets.veth47f0782	inbound packets dropped ratio	-	UNINITIALIZED
10:38:09 PM	net_packets.veth1f82487	10s received packets storm	-	UNINITIALIZED

Showing 1 to 10 of 1000 rows10 rows per page

<12345...100>

Close

The average I/O operation size.

The average service time for completed I/O operations. The average service time will be misleading.

The number of merged disk operations. The system is able to handle more than 1000 merged disk operations per second.



Disk space utilization. reserved for root is automatically reserved.




el operations the reporting




server/Local file	Director	Remote file	Size	Priority	Status
sftp://ubuntu@					
/home/aurelijus/Pict...	-->	/MVI_2543.MOV	1,5 GB	Normal	Transferring
00:11:45 elapsed	00:07:27 left	818.397.184 bytes (1,4 MB/s)			
/home/aurelijus/Pict...	-->	/MVI_2592.MOV	1,5 GB	Normal	Transferring
00:11:44 elapsed	00:08:25 left	820.019.200 bytes (1,2 MB/s)			
/home/aurelijus/Pict...	-->	/MVI_2593.MOV	1,2 GB	Normal	Transferring
00:10:41 elapsed	00:06:07 left	744.816.640 bytes (1,2 MB/s)			
/home/aurelijus/Pict...	-->	/MVI_2596.MOV	857,5 MB	Normal	Transferring


2 Trade-offs



Lost data



Relation





Netdata
Trust, fast alarms

Decision making

ElasticSearch
Aggregation,
comparison

CloudWatch
Availability, tooling

Debugging



Netdata

Decision making

ElasticSearch

CloudWatch

Debugging

Problem

Ultimate debugging tool

```
CityServiceImpl.java - spring-boot-sample-data-jpa - IntelliJ IDEA (...)  
spring-boot-sample-data-jpa | src | main | java | sample | data | jpa | service | CityServiceImpl | Sample |  
SampleController.java | CityServiceImpl.java |  
@Override  
public Page<City> findCities(CitySearchCriteria criteria, Pageable pageable) { criteria: CitySearchCriteria@7068 pageable: "Page req  
    Assert.notNull(criteria, "Criteria must not be null");  
    String name = criteria.getName(); name: "Sydney" criteria: CitySearchCriteria@7068  
  
    if (!StringUtils.hasLength(name)) {  
        return this.cityRepository.findAll(null);  
    }  
  
    String country = ""; country: ""  
    int splitPos = name.lastIndexOf(","); splitPos: -1  
  
    if (splitPos >= 0) {  
        country = name.substring(splitPos + 1);  
        name = name.substring(0, splitPos); splitPos: -1  
    }  
  
    Page<City> results = this.cityRepository results: "Page 0 of 1 containing sample.data.jpa.domain.City instances" cityRepository:  
        .findByNameContainingAndCountryContainingAllIgnoringCase(name.trim(), name: "Sydney"  
            country.trim(), pageable); country: "" pageable: "Page request [number: 0, size 10, sort: null]"  
    return results; results: "Page 0 of 1 containing sample.data.jpa.domain.City instances"  
}  
  
@Override  
public City getCity(String name, String country) {  
    Assert.notNull(name, "Name must not be null");  
    Assert.notNull(country, "Country must not be null");  
    return this.cityRepository.findByNameAndCountryAllIgnoringCase(name, country);  
}  
  
@Override  
public Page<HotelSummary> getHotels(City city, Pageable pageable) {  
    Assert.notNull(city, "City must not be null");  
    return this.hotelRepository.findByCity(city, pageable);  
}
```

Evaluate Expression

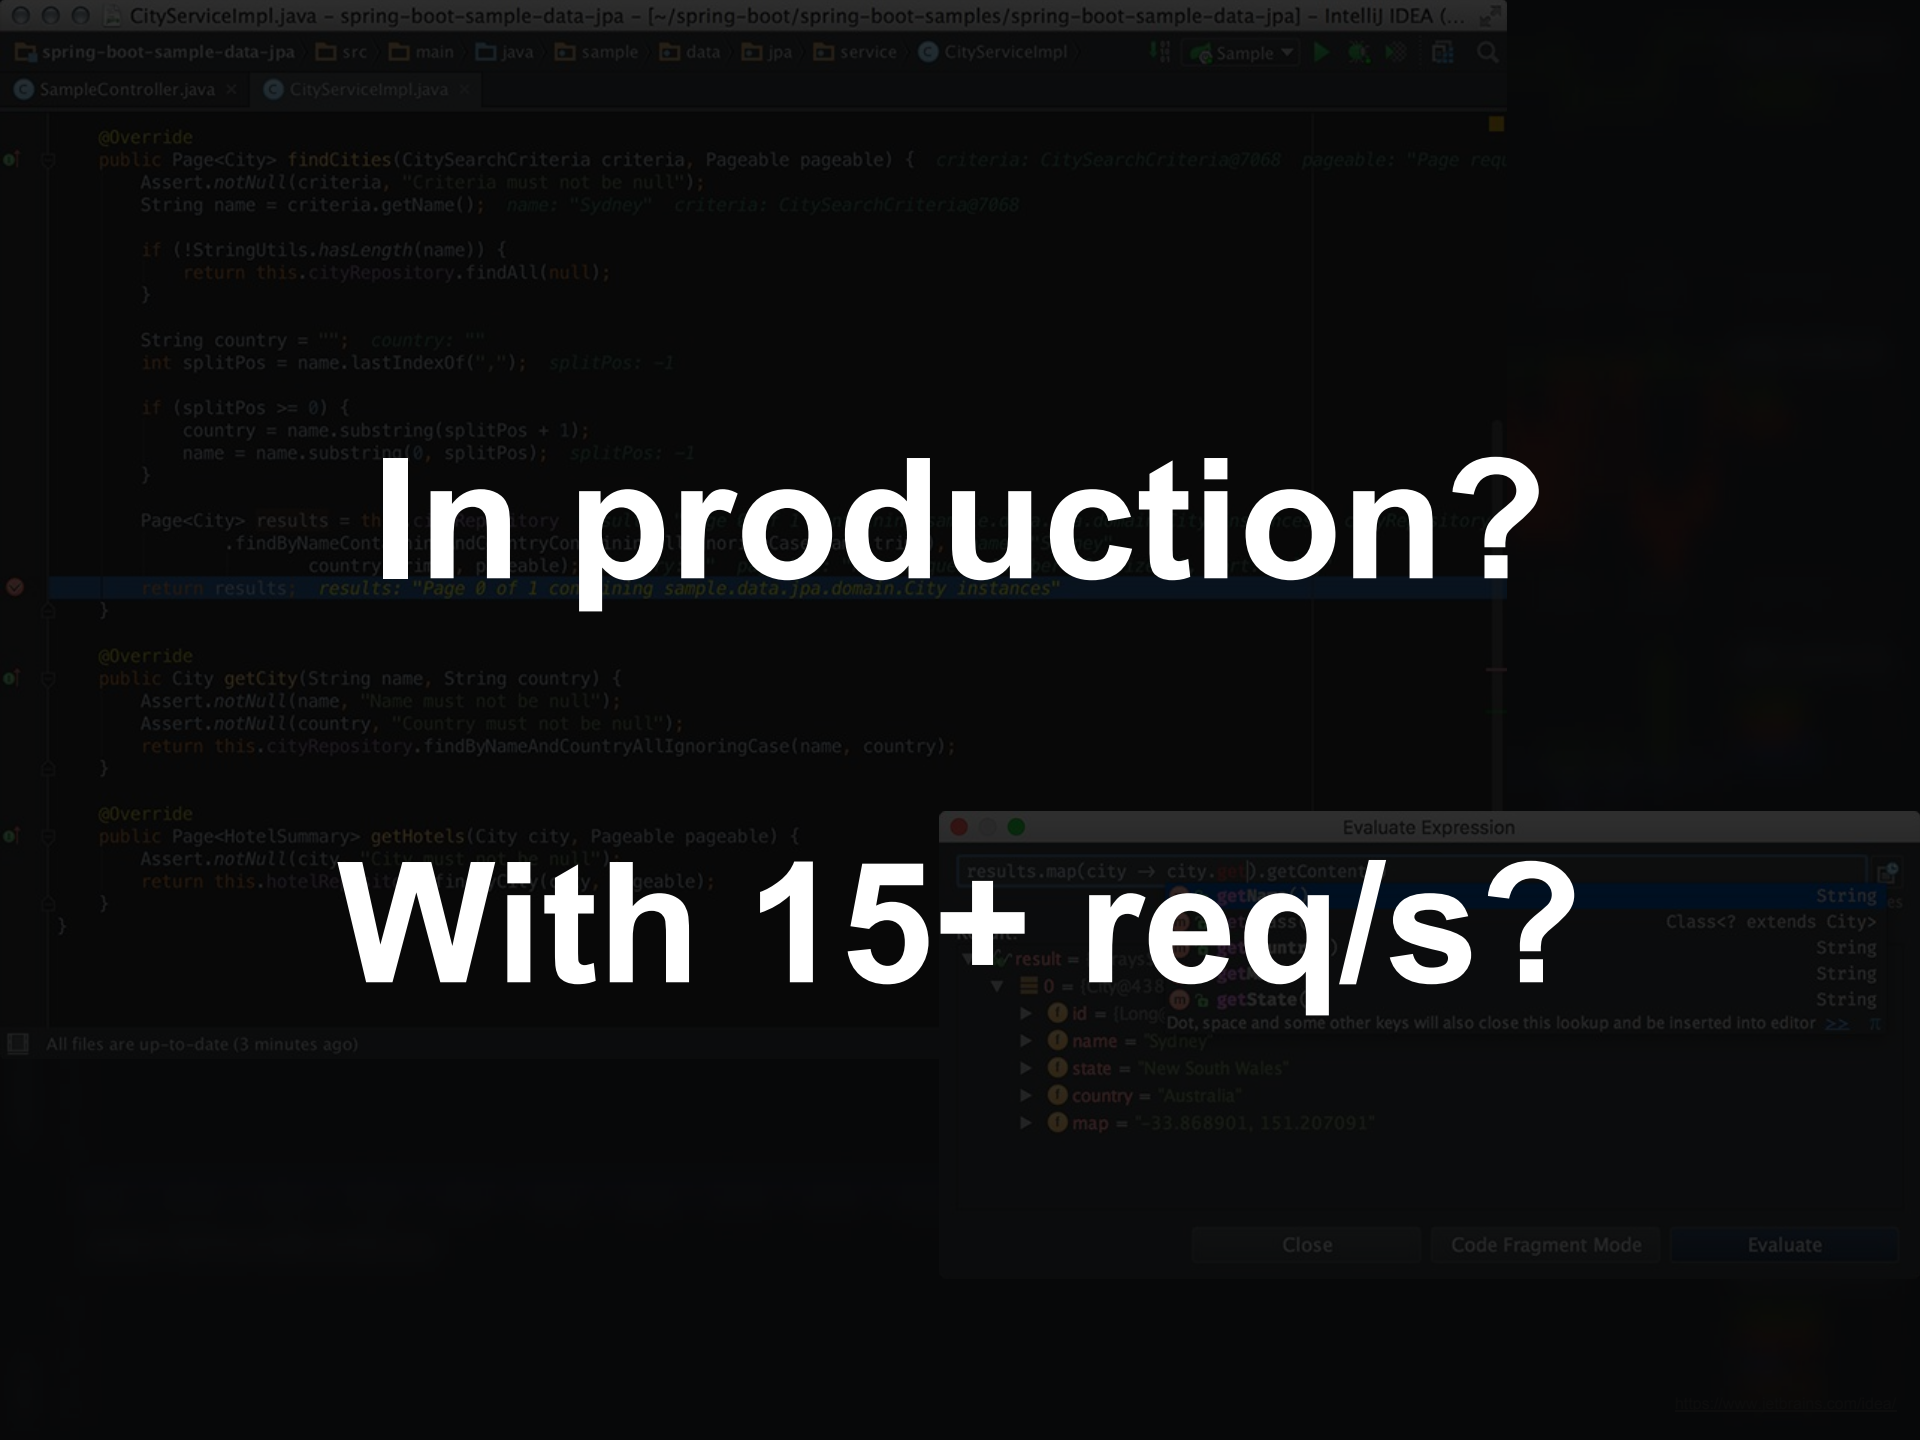
results.map(city -> city.get()).getContent()

Result:

- result = {Arrays\$ArrayList@438}
- 0 = {City@438}
- id = {Long@438}
- name = "Sydney"
- state = "New South Wales"
- country = "Australia"
- map = "-33.868901, 151.207091"

String
Class<? extends City>
String
String
String
String

Close Code Fragment Mode Evaluate



In production?

With 15+ req/s?

[illegible]

And with asynchronous code

Proposal



@timestamp per hour

0

Count



No results found



@timestamp per minute



statsd

counters



cron.hdd-cw (statsd_counter_cron.hdd_cw)



Sun, Nov 12, 2017
21:53:17

events/s
counter 0.00
events 0.00

cron.memory-cw (statsd_counter_cron.memory_cw)



Sun, Nov 12, 2017
21:53:17

events/s
counter 0.00
events 0.00


Monitoring in development environment



Edge cases


TargetResponseTime per AZ

Seconds





TargetResponseTime per AZ

Seconds



TargetResponseTime per AZ




TargetResponseTime per AZ

Seconds

Understanding different scales

NetworkIn NetworkOut



Netdata
Instance level

Decision making

ElasticSearch
Aggregating

Resource starvation

CloudWatch
Cross instance

Debugging

Intro

What are metrics

Compare

**Principles, tradeoffs,
tools and added value**

Demo

How does it feel to do
Real-time first

Intro

What are metrics

Compare

**Principles, tradeoffs,
tools and added value**

Demo

**How does it feel to do
Real-time first**

Netdata

Decision making

DEMO

ElasticSearch Resource starvation

CloudWatch

Debugging

- **Just install**
- **Custom metrics**
- **Publish**
- **Analyse**

DEMO

- **Just install**
- **Custom metrics**
- **Publish**
- **Analyse**

DEMO

**Start
Real-time
Metrics
Now**

Intro

What are metrics

Compare

**Principles, tradeoffs,
tools and added value**

Demo

**How does it feel to do
Real-time first**

Metrics mindset

Real-time-first metrics

Thank you

Aurelijus Banelis



Real-time-first metrics

Questions?

Or bonus topic “Resource starvation”?

Aurelijus Banelis



Slides already at <https://aurelijus.banelis.lt/presentations/build-stuff-2017/Realtime-first-metrics.pdf>

Real-time-first metrics

Aurelijus Banelis



References

- <https://blog.daftcode.pl/hype-driven-development-3469fc2e9b22>
- <https://aws.amazon.com/about-aws/whats-new/2017/07/amazon-cloudwatch-introduces-high-resolution-custom-metrics-and-alarms/>
- <https://www.hometogo.com/media/funding/>
- <https://gist.github.com/jboner/2841832>
- https://en.wikipedia.org/wiki/Circular_buffer
- <https://github.com/firehol/netdata/wiki/monitoring-ephemeral-nodes>
- <https://12factor.net/>
- <https://github.com/saprykin/plibsys>
- <https://github.com/firehol/netdata/issues/217>
- <http://riemann.io/>
- <https://github.com/dspinellis/dgsh>



Bonus

Netdata

Decision making

ElasticSearch

Resource starvation

(synonym: Resource Exhaustion)

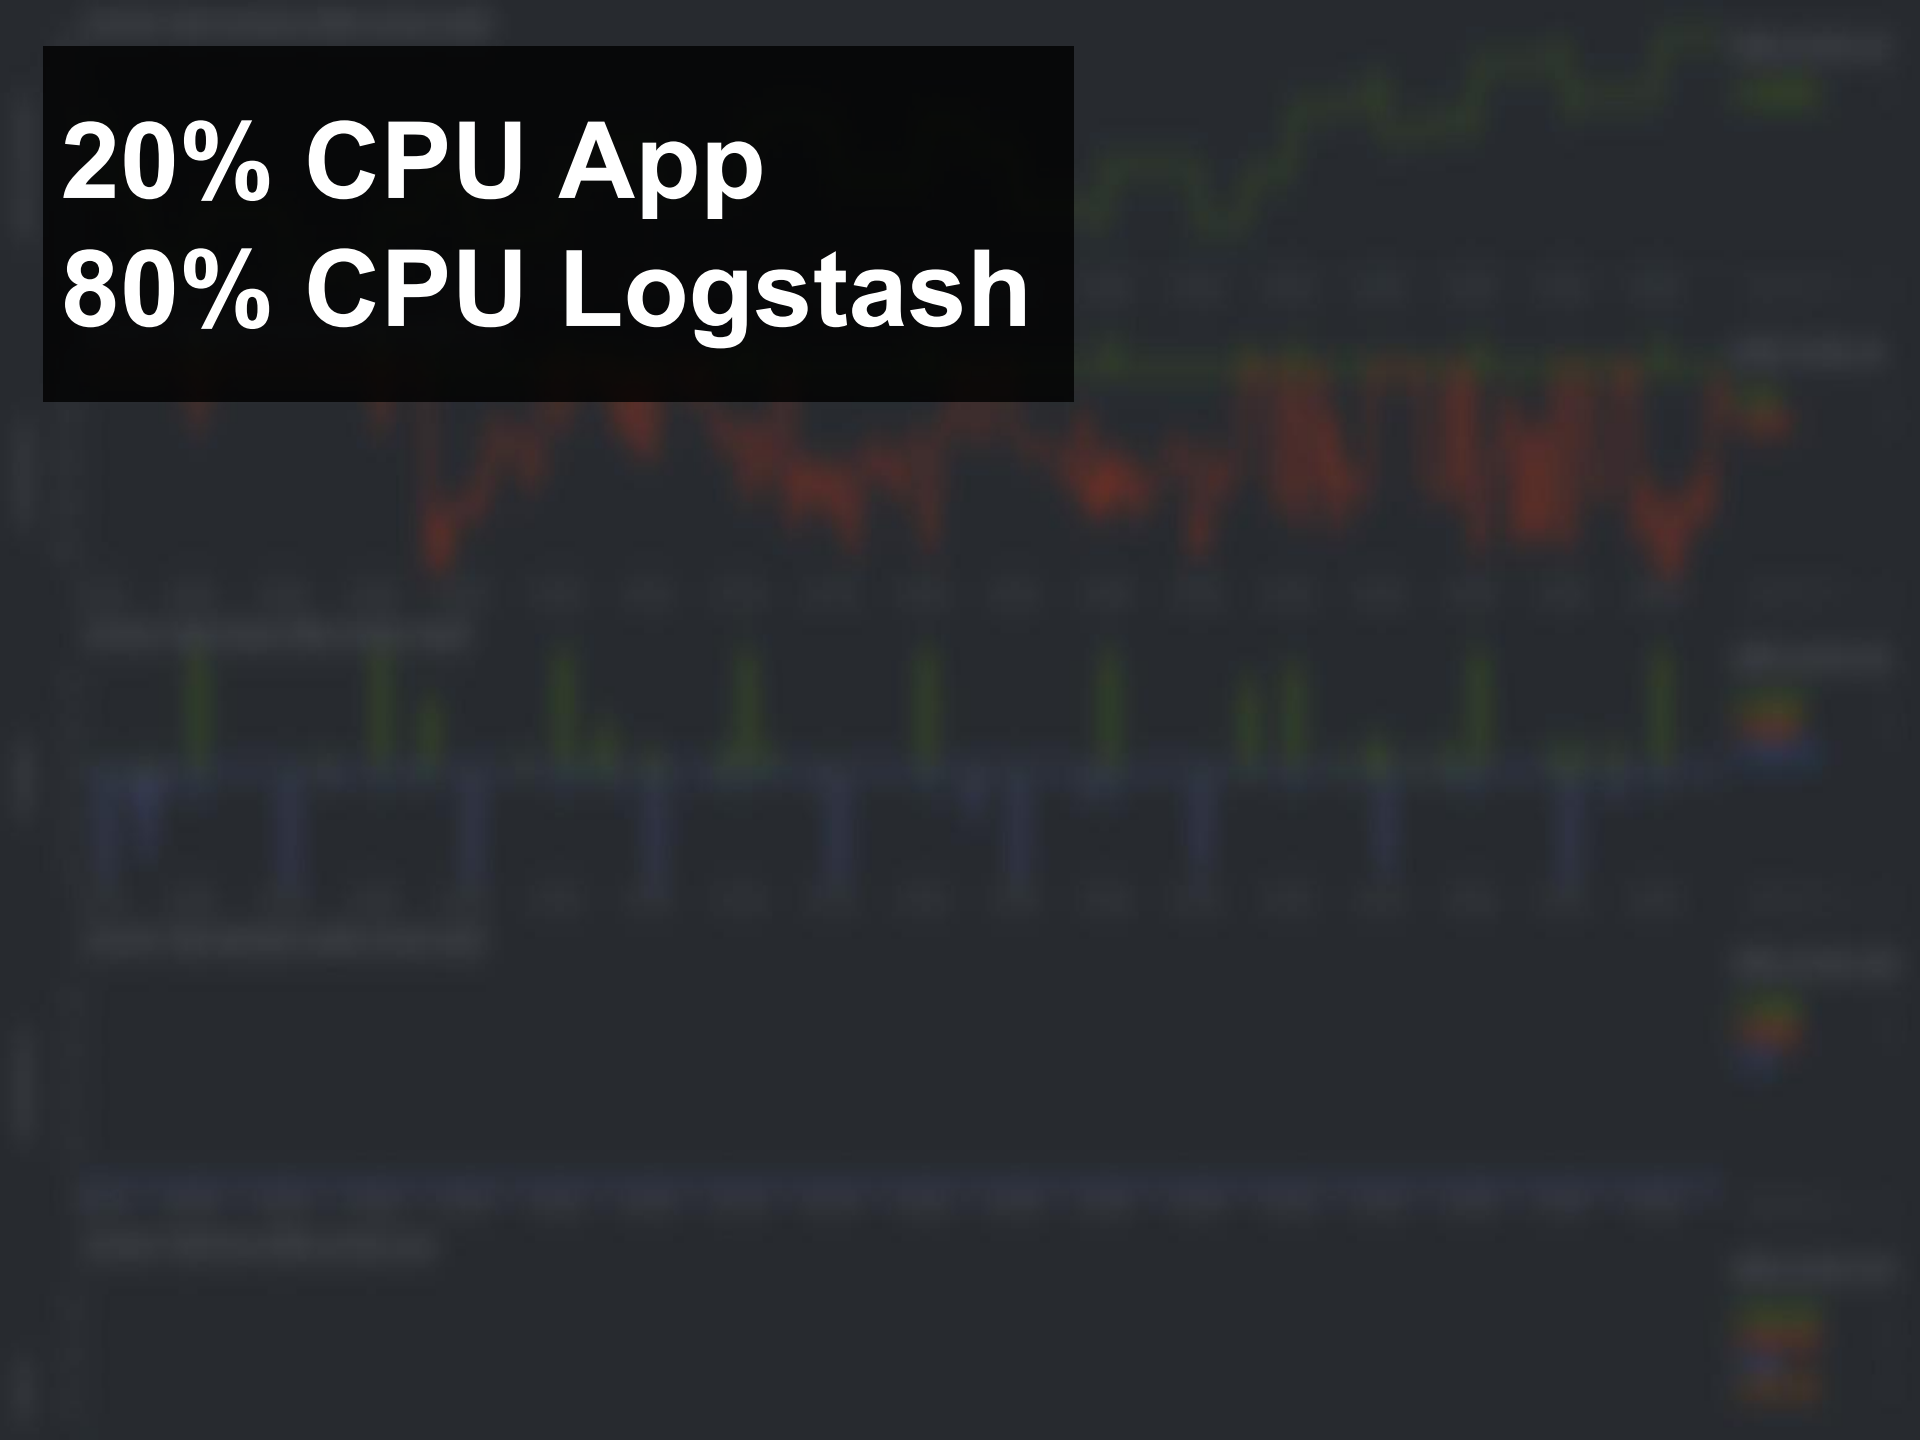
CloudWatch

Debugging

Problem

Log everything

20% CPU App
80% CPU Logstash

The background of the slide is a blurred screenshot of a monitoring dashboard. It features several line graphs with different colored lines (green, red, blue) plotted against a dark background. There are also some bar charts and data points visible, though they are out of focus. The overall aesthetic is technical and data-oriented.

20% CPU App
80% CPU Logstash

1X time with logs
2X time without logs

20% CPU App
80% CPU Logstash

1X time with logs
2X time without logs


zcat

20% CPU App
80% CPU Logstash

1X time with logs
2X time without logs

zcat over 1TB / day

Remember
Mindset of
Real-time
applications



Latency Numbers Every Programmer Should Know

■ 1 ns

■ L1 cache reference: 0.5 ns

■ Branch mispredict: 5 ns

■ L2 cache reference: 7 ns

■ Mutex lock/unlock: 25 ns

■ = 100 ns

■ Main memory reference: 100 ns

■ = 1 μ s

■ Compress 1 KB with Zippy: 3 μ s

■ = 10 μ s

■ Send 1 KB over 1 Gbps network: 10 μ s

■ SSD random read (1 Gb/s SSD): 150 μ s

■ Read 1 MB sequentially from memory: 250 μ s

■ Round trip in same datacenter: 500 μ s

■ = 1 ns

■ Read 1 MB sequentially from SSD: 1 ms


■ Disk seek: 10 ms

■ Read 1 MB sequentially from disk: 20 ms

■ Packet roundtrip CA to Netherlands: 150 ms

Source: <https://gist.github.com/2841832>

3 Architectural decisions




- **Amazingly fast**

responds to all queries in less than 0.5 ms per metric, even on low-end hardware

- **Highly efficient**

collects thousands of metrics per server per second, with just 1% CPU utilization of a single core, a few MB of RAM and no disk I/O at all

Fixed memory (Circular buffer)





- **Amazingly fast**

responds to all queries in less than 0.5 ms per metric, even on low-end hardware

- **Highly efficient**

collects thousands of metrics per server per second, with just 1% CPU utilization of a single core, a few MB of RAM and no disk I/O at all



The image shows a dark-themed dashboard interface. At the top, there are two dropdown menus: 'Time range' set to 'Last 24 hours' and 'Period' set to '1 minute'. To the right of these is a circular refresh icon. Below the controls is a line chart on a grid background. The chart shows a blue line that starts at a high level on the left, trends downwards with some fluctuations, and then rises sharply towards the right edge. Overlaid on the chart in large white font is the text 'Aggregate more in the past'.


Aggregate more in the past

Time range

Last 24 hours ▾

Period

1 minute ▾



Commit point





In-memory now

and HDD optimised



In-memory buffer



Netdata
In fixed memory

Decision making

ElasticSearch
Delayed flush

Resource starvation

CloudWatch
Keep aggregated

Debugging

Real-time-first metrics

Aurelijus Banelis

