Introduction DevOps / Monitoring / Observabilité

EPITECH, Mercredi 5 février 2020

Who am I?

Lionel Porcheron, CEO & co-founder Bleemeo

- DevOps for +15 years (started my monitoring journey with ~~nagios~~ netsaint)
- Toulouse DevOps Meetup Leader
- Capitole du Libre Leader

Who are you?

Quick Surveys:

- Who is familiar with DevOps?
- Who is familiar with Docker?
- Who is familiar with Kubernetes?
- Who is familiar with Monitoring?

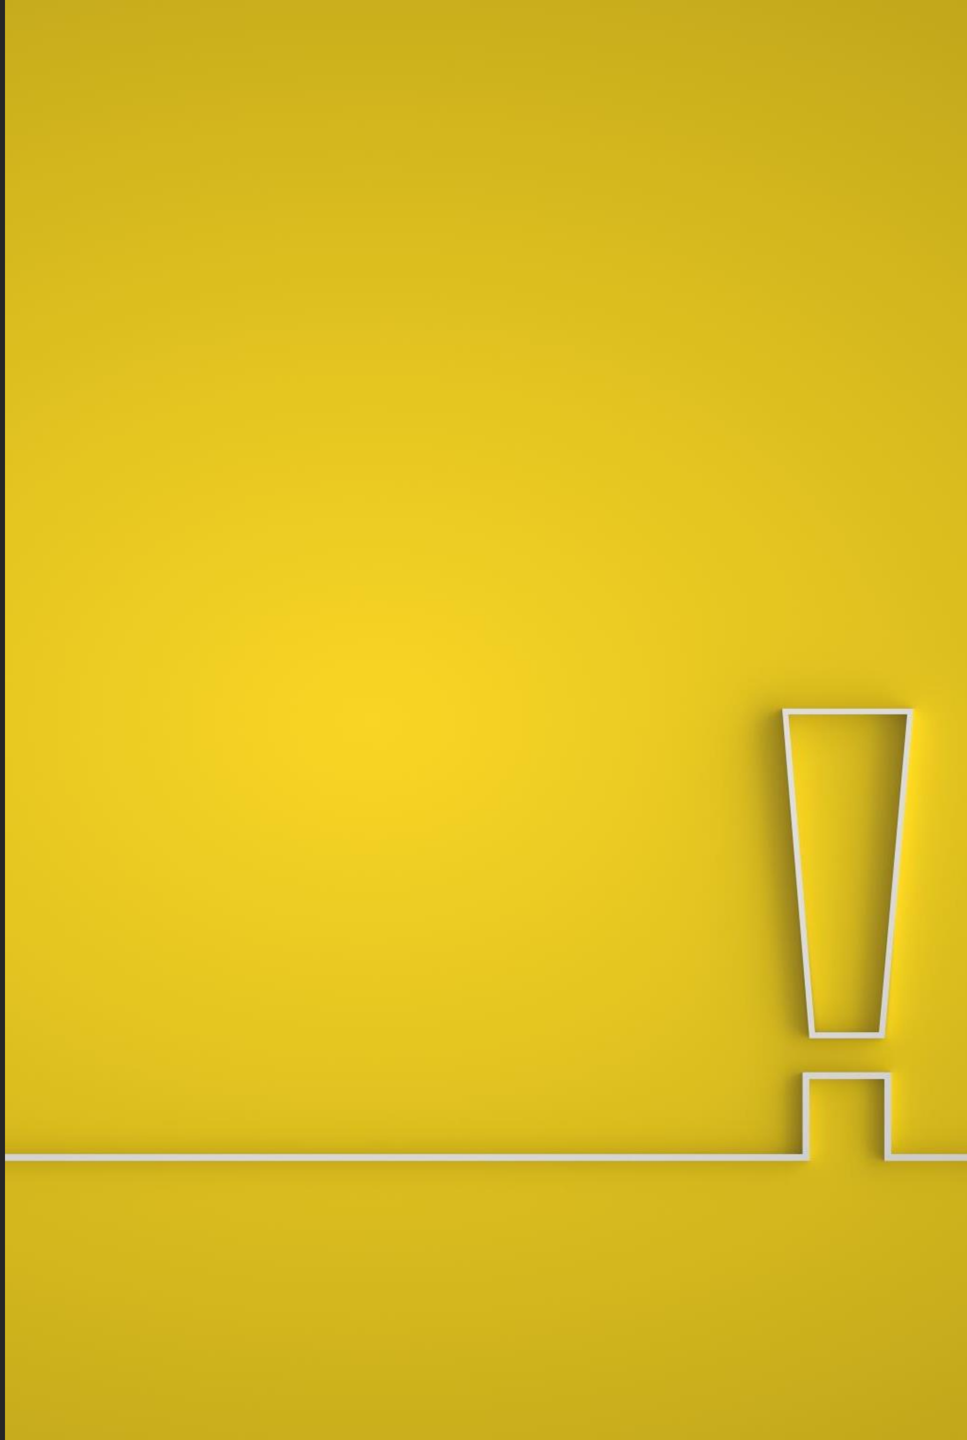

Let's define some (buzz)words (1/2)

Operations (exploitation in French): operate a service, running code in production and make sure the service is available for customers/users.

Because the goal of writing code is not only to write code!

Monitoring: check that a service is operating as expected. Use to be "simple" checks like ping, checks ports, etc.

Let's define some (buzz)words (2/2)

DevOps:

- "agile" movement for the operations
- movement initiated around 2007 to get the development teams closer to the operations and avoid silos. Dev vs Ops was (and is still) a classical patterns in numerous organizations.
- While we often talk of tooling, automation, the initial mindset of this movement was a real culture change with mutual understanding.
- This is also pushed by major Internet actors: Facebook, Amazon, Google as they provide a service, not code
- Every company implement his DevOps strategy in a different way: SRE for Google, "you build it, you run it" for Amazon.
- Huge adoption, as all companies want to transform in services companies (getting recurring revenues) instead of product companies (selling one shot).

DevOps

- Improve overall tooling
- Introduce "continuous deployment"
- Monitoring/Observability are core in the DevOps approach

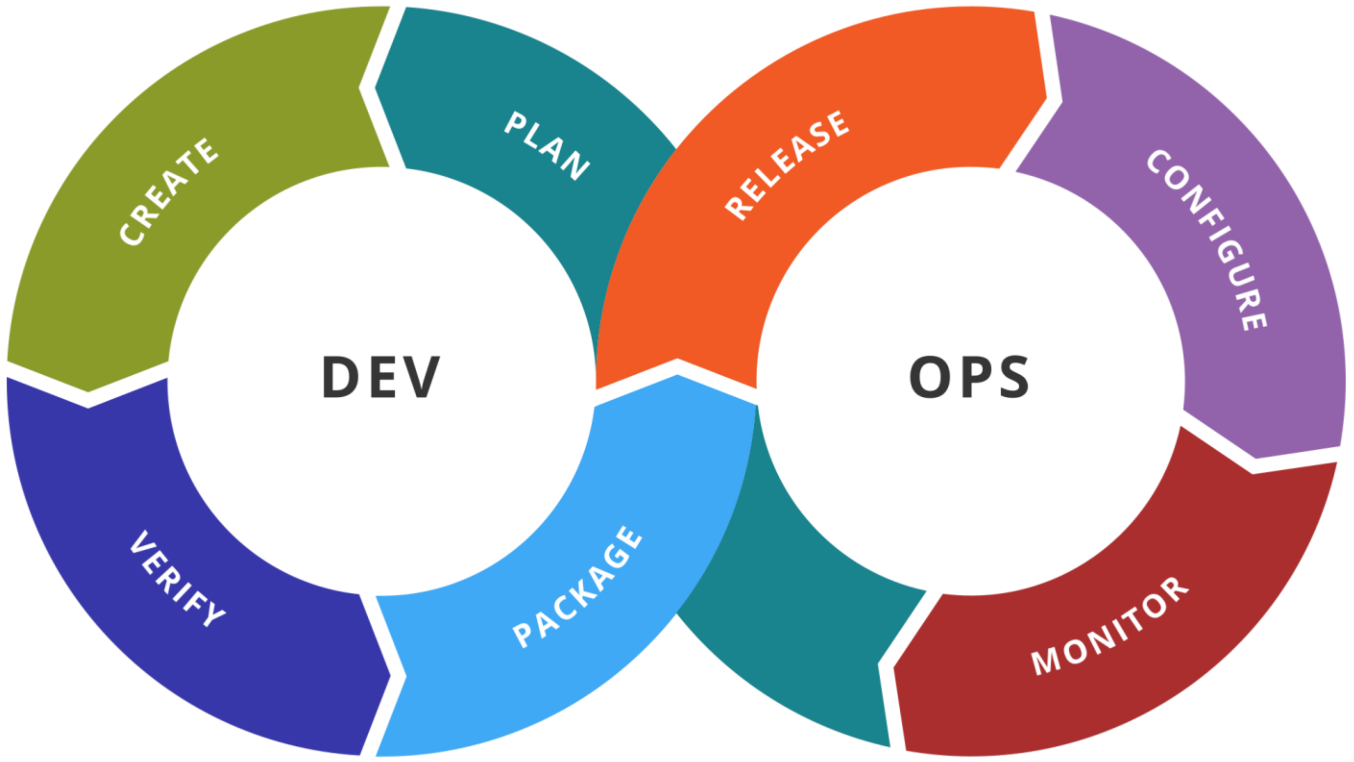

Bleemeo?

Observability & Monitoring as a service solution

Compatible with market standards

Two Open Source projects:

- **Glouton**, universal monitoring agent written in Go with Prometheus, Statsd, Graphite, Nagios, Zabbix compatibility
- **SquirrelDB**, a scalable Prometheus compatible storage backend based on Cassandra

Old Monitoring Days

- Monitor your server as a blackbox
- Only monitor server & services (web server, database)
- Only availability, not metrics
- Nagios & derivates

Services Dashboards

- Up / Down indicators
- Host oriented
- No metrics

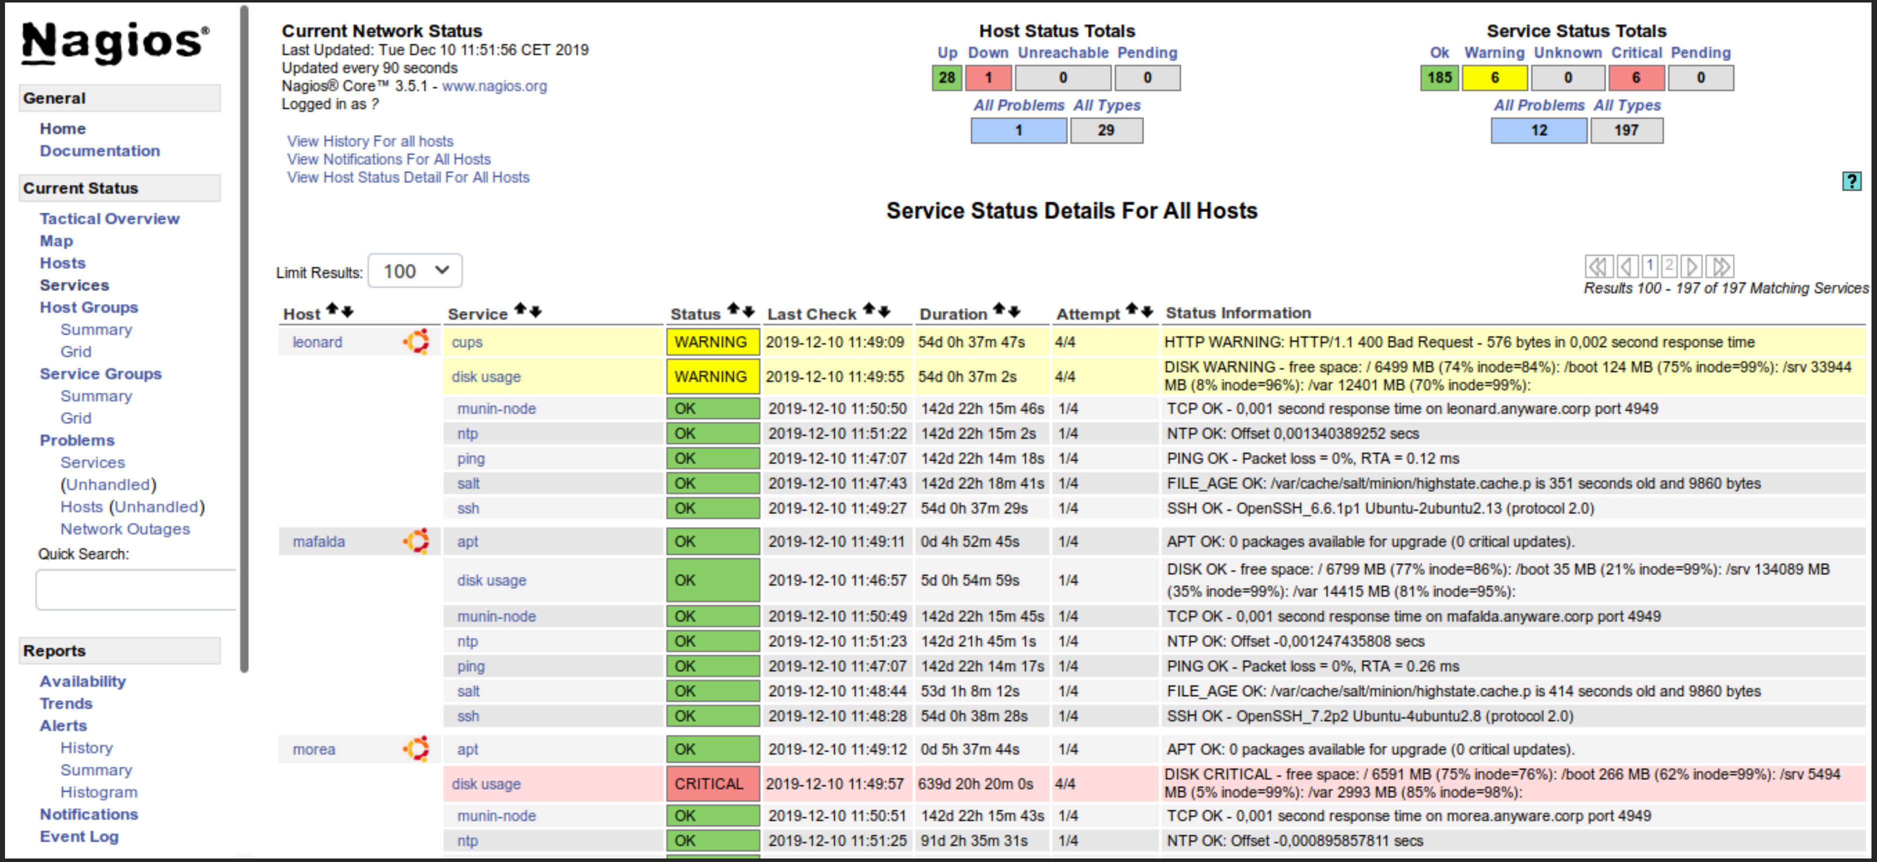

Current Network Status
Last Updated: Tue Dec 10 11:51:56 CET 2019
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as ?

View History For all hosts
View Notifications For All Hosts
View Host Status Detail For All Hosts

Host Status Totals

Up	Down	Unreachable	Pending
28	1	0	0

All Problems	All Types
1	29

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
185	6	0	6	0

All Problems	All Types
12	197

Service Status Details For All Hosts

Limit Results: 100

Results 100 - 197 of 197 Matching Services

Host	Service	Status	Last Check	Duration	Attempt	Status Information
leonard	cups	WARNING	2019-12-10 11:49:09	54d 0h 37m 47s	4/4	HTTP WARNING: HTTP/1.1 400 Bad Request - 576 bytes in 0,002 second response time
	disk usage	WARNING	2019-12-10 11:49:55	54d 0h 37m 2s	4/4	DISK WARNING - free space: / 6499 MB (74% inode=84%): /boot 124 MB (75% inode=99%): /srv 33944 MB (8% inode=96%): /var 12401 MB (70% inode=99%):
	munin-node	OK	2019-12-10 11:50:50	142d 22h 15m 46s	1/4	TCP OK - 0,001 second response time on leonard.anyware.corp port 4949
	ntp	OK	2019-12-10 11:51:22	142d 22h 15m 2s	1/4	NTP OK: Offset 0,001340389252 secs
	ping	OK	2019-12-10 11:47:07	142d 22h 14m 18s	1/4	PING OK - Packet loss = 0%, RTA = 0.12 ms
	salt	OK	2019-12-10 11:47:43	142d 22h 18m 41s	1/4	FILE_AGE OK: /var/cache/salt/minion/highstate.cache.p is 351 seconds old and 9860 bytes
	ssh	OK	2019-12-10 11:49:27	54d 0h 37m 29s	1/4	SSH OK - OpenSSH_6.6.1p1 Ubuntu-2ubuntu2.13 (protocol 2.0)
mafalda	apt	OK	2019-12-10 11:49:11	0d 4h 52m 45s	1/4	APT OK: 0 packages available for upgrade (0 critical updates).
	disk usage	OK	2019-12-10 11:46:57	5d 0h 54m 59s	1/4	DISK OK - free space: / 6799 MB (77% inode=86%): /boot 35 MB (21% inode=99%): /srv 134089 MB (35% inode=99%): /var 14415 MB (81% inode=95%):
	munin-node	OK	2019-12-10 11:50:49	142d 22h 15m 45s	1/4	TCP OK - 0,001 second response time on mafalda.anyware.corp port 4949
	ntp	OK	2019-12-10 11:51:23	142d 21h 45m 1s	1/4	NTP OK: Offset -0,001247435808 secs
	ping	OK	2019-12-10 11:47:07	142d 22h 14m 17s	1/4	PING OK - Packet loss = 0%, RTA = 0.26 ms
	salt	OK	2019-12-10 11:48:44	53d 1h 8m 12s	1/4	FILE_AGE OK: /var/cache/salt/minion/highstate.cache.p is 414 seconds old and 9860 bytes
	ssh	OK	2019-12-10 11:48:28	54d 0h 38m 28s	1/4	SSH OK - OpenSSH_7.2p2 Ubuntu-4ubuntu2.8 (protocol 2.0)
morea	apt	OK	2019-12-10 11:49:12	0d 5h 37m 44s	1/4	APT OK: 0 packages available for upgrade (0 critical updates).
	disk usage	CRITICAL	2019-12-10 11:49:57	639d 20h 20m 0s	4/4	DISK CRITICAL - free space: / 6591 MB (75% inode=76%): /boot 266 MB (62% inode=99%): /srv 5494 MB (5% inode=99%): /var 2993 MB (85% inode=98%):
	munin-node	OK	2019-12-10 11:50:51	142d 22h 15m 43s	1/4	TCP OK - 0,001 second response time on morea.anyware.corp port 4949
	ntp	OK	2019-12-10 11:51:25	91d 2h 35m 31s	1/4	NTP OK: Offset -0,000895857811 secs

Network Operation Center

- ~~Big~~ huge displays
- Network traffic
- Where the operations team "lives"
- Apply Standard Operating Procedures (SOP)
- A bit old fashioned...

Microservices and modern era

- Increase architecture complexity
- Increase number of technical components to monitor
- Moderns infrastructure base on containers are dynamic
- Some components may come from third parties

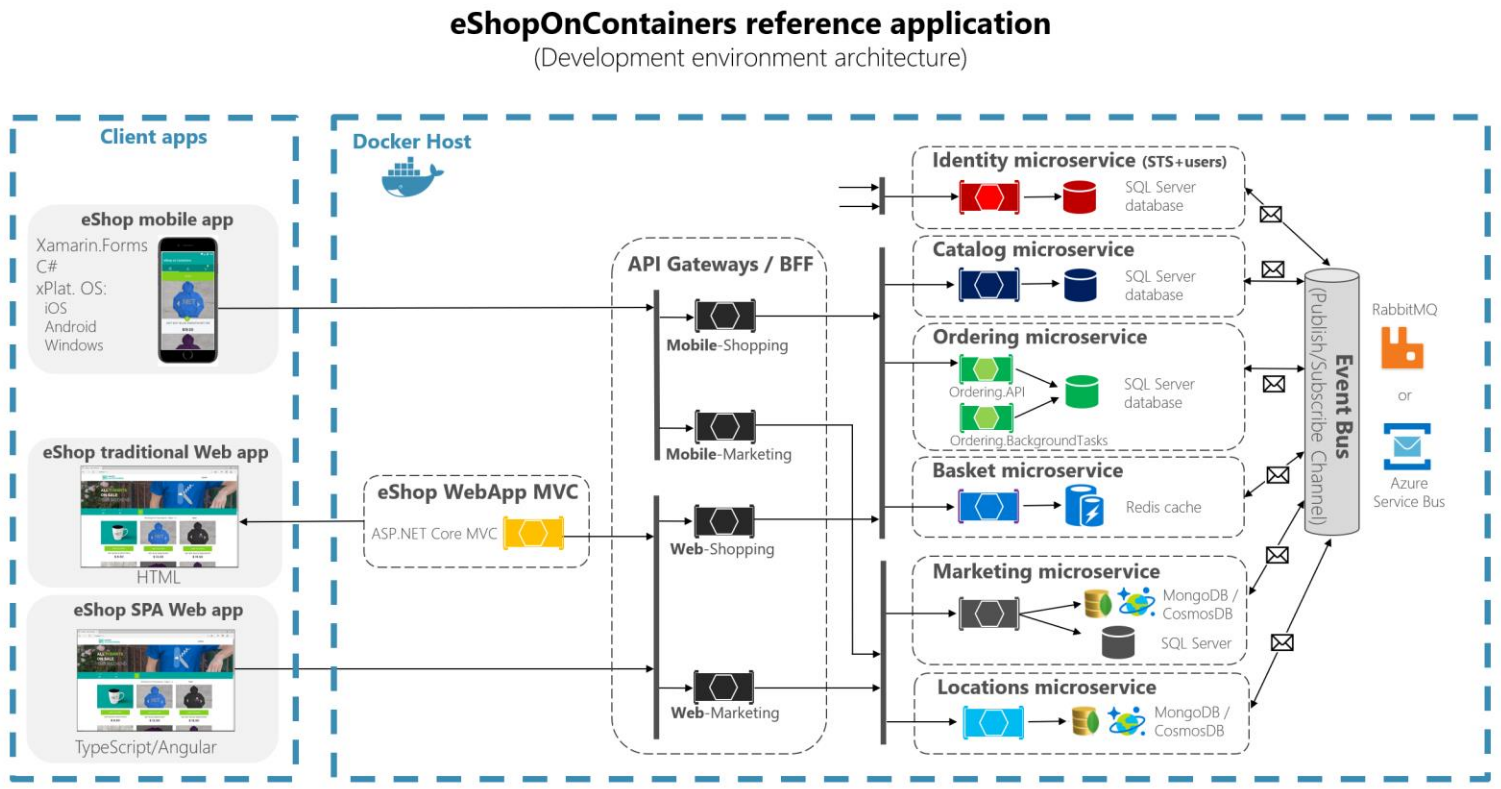

Graphite… and Prometheus

- Graphite change the way we were doing monitoring: metrics became central
- Graphite appeared in 2008
- Prometheus became de-facto standard for monitoring
- Prometheus was "invented" in 2012 at Soundcloud and is now a (graduated) CNCF project
- Ecosystem based on Prometheus: exporters, Grafana, software themselves (Kubernetes, Traefik & many others)

What type of monitoring?

- Monitoring your system infrastructure (servers, etc.)
- Monitoring your network (switch, routers, firewalls)
- Monitoring your application logs (ELK from Elastic)
- Monitoring your application performances (need an agent in the application)

s/monitoring/observability/

- No more monitoring as blackbox: we now know what is inside
- Exports tons of metrics for future usage
- Code need to be instrumented to provide business metrics
- New Buzzword 😎

Three pillars of observability

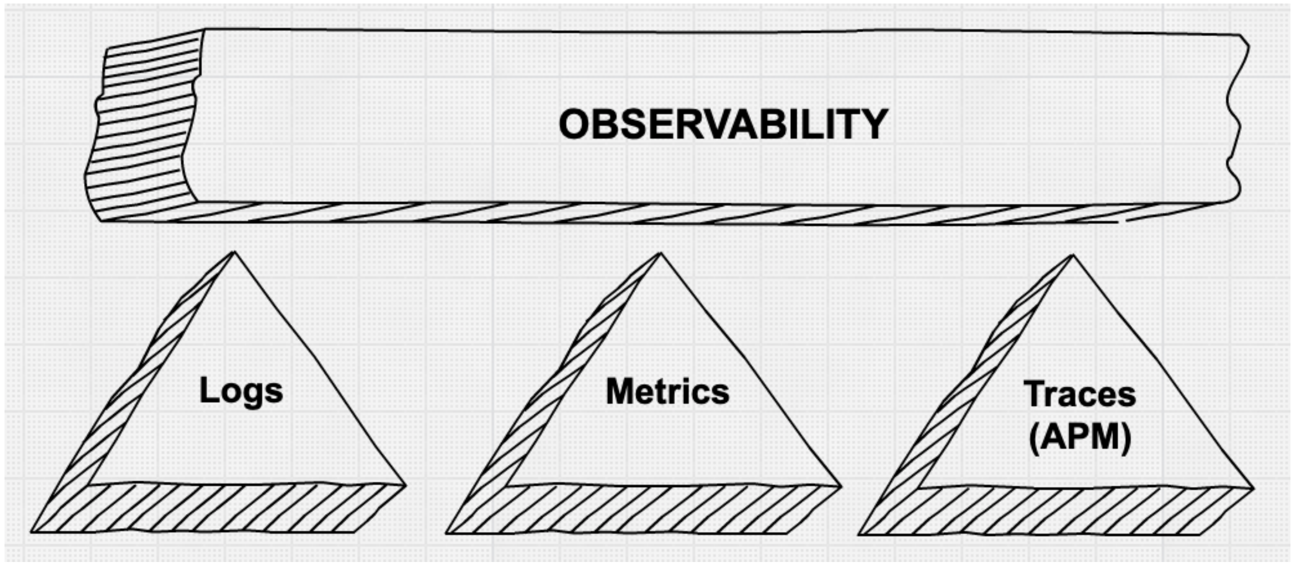

Observability Golden Signals

The RED Method

- (Request) **R**ate - the number of requests, per second, you services are serving.
- (Request) **E**rrors - the number of failed requests per second.
- (Request) **D**uration - distributions of the amount of time each request takes.

The USE Method

- (Ressource) **U**tilization: as a percent over a time interval. eg, "one disk is running at 90% utilization".
- (Ressource) **S**aturation: as a queue length. eg, "the CPUs have an average run queue length of four".
- (Ressource) **E**rrors: scalar counts. eg, "this network interface has had fifty late collisions".

Prometheus exporters

- Prometheus exporters export on web page metrics (basic plain text page with a metric per line)
- Prometheus poll regularly those endpoints
- Prometheus exporters exist for almost everything:

  https://prometheus.io/docs/instrumenting/exporters/

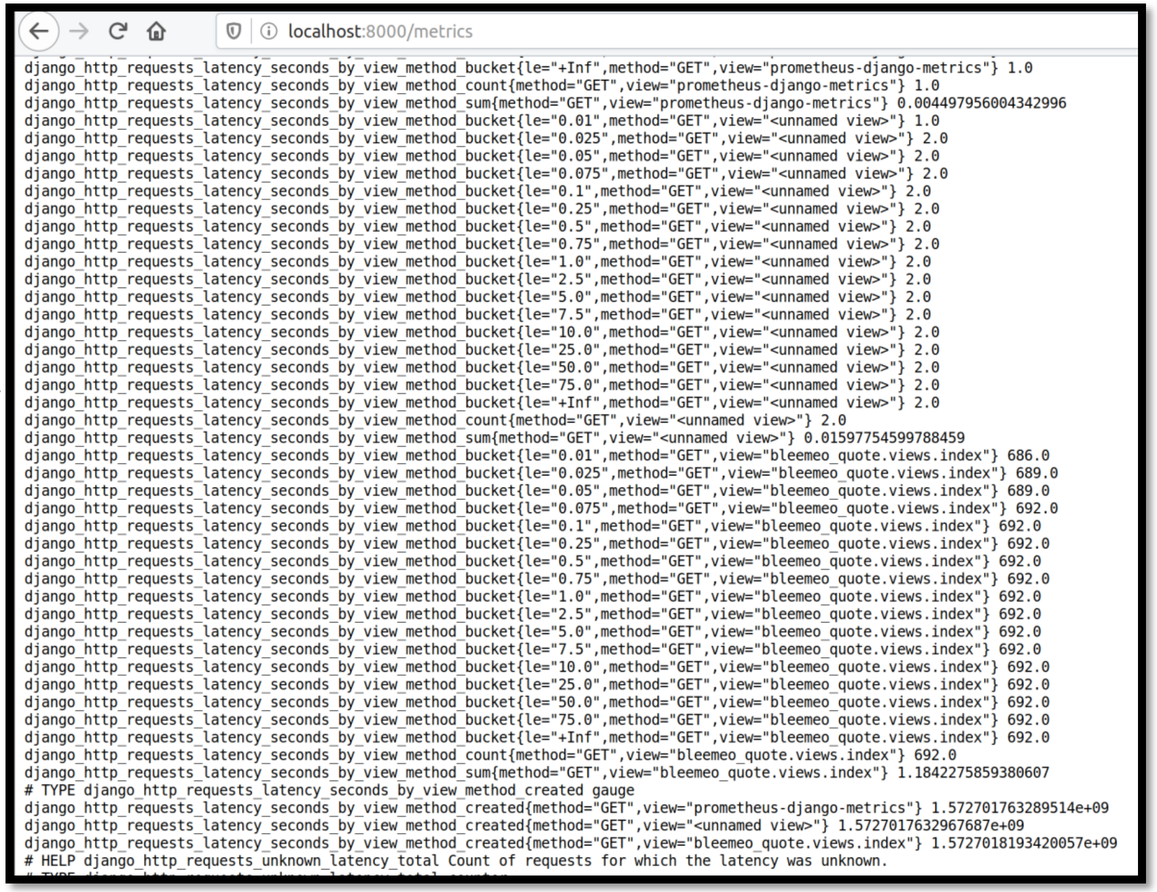

```
django_http_requests_latency_seconds_by_view_method_bucket{le="+Inf",method="GET",view="prometheus-django-metrics"} 1.0
django_http_requests_latency_seconds_by_view_method_count{method="GET",view="prometheus-django-metrics"} 1.0
django_http_requests_latency_seconds_by_view_method_sum{method="GET",view="prometheus-django-metrics"} 0.004497956004342996
django_http_requests_latency_seconds_by_view_method_bucket{le="0.01",method="GET",view="<unnamed view>"} 1.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.025",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.05",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.075",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.1",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.25",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.5",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.75",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="1.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="2.5",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="5.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="7.5",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="10.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="25.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="50.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="75.0",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_bucket{le="+Inf",method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_count{method="GET",view="<unnamed view>"} 2.0
django_http_requests_latency_seconds_by_view_method_sum{method="GET",view="<unnamed view>"} 0.01597754599788459
django_http_requests_latency_seconds_by_view_method_bucket{le="0.01",method="GET",view="bleemeo_quote.views.index"} 686.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.025",method="GET",view="bleemeo_quote.views.index"} 689.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.05",method="GET",view="bleemeo_quote.views.index"} 689.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.075",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.1",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.25",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.5",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="0.75",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="1.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="2.5",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="5.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="7.5",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="10.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="25.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="50.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="75.0",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_bucket{le="+Inf",method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_count{method="GET",view="bleemeo_quote.views.index"} 692.0
django_http_requests_latency_seconds_by_view_method_sum{method="GET",view="bleemeo_quote.views.index"} 1.1842275859380607
# TYPE django_http_requests_latency_seconds_by_view_method_created gauge
django_http_requests_latency_seconds_by_view_method_created{method="GET",view="prometheus-django-metrics"} 1.572701763289514e+09
django_http_requests_latency_seconds_by_view_method_created{method="GET",view="<unnamed view>"} 1.5727017632967687e+09
django_http_requests_latency_seconds_by_view_method_created{method="GET",view="bleemeo_quote.views.index"} 1.5727018193420057e+09
# HELP django_http_requests_unknown_latency_total Count of requests for which the latency was unknown.
```

Instrument your code

- For Django application: django-prometheus
- Django Middleware for metrics
- Health checks and custom checks with django-health-checks

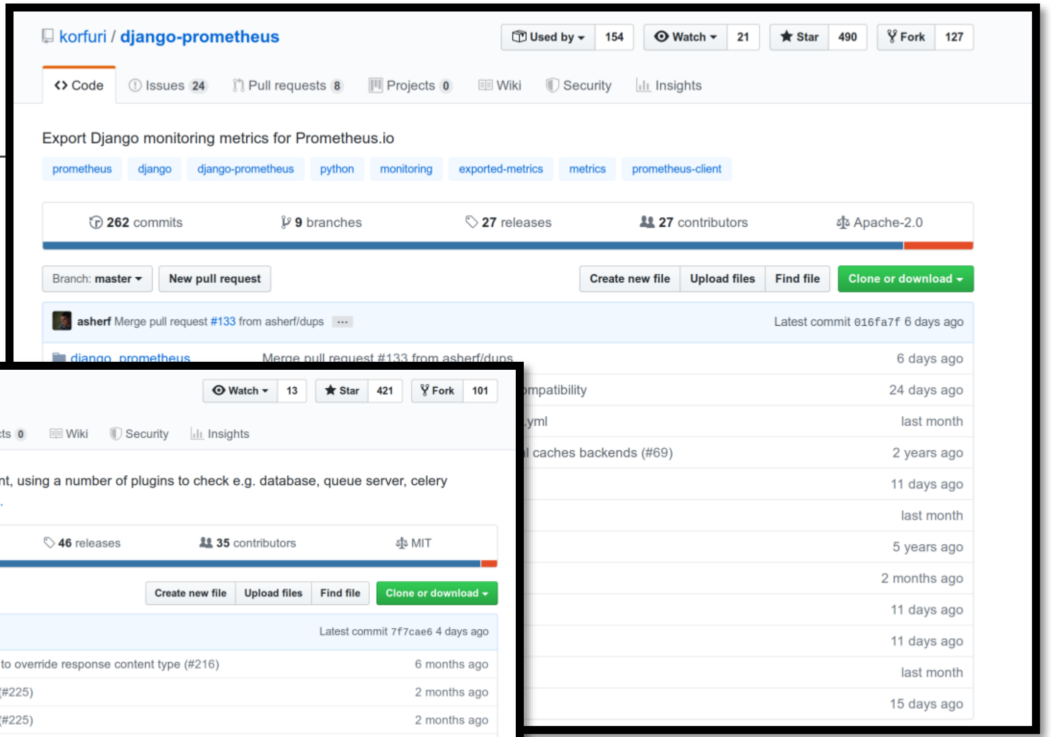

korfuri / django-prometheus

Export Django monitoring metrics for Prometheus.io

prometheus django django-prometheus python monitoring exported-metrics metrics prometheus-client

262 commits | 9 branches | 27 releases | 27 contributors | Apache-2.0

asherf Merge pull request #133 from asherf/dups … Latest commit 016fa7f 6 days ago

django_prometheus	Merge pull request #133 from asherf/dups	6 days ago
	ompatibility	24 days ago
	.yml	last month
	l caches backends (#69)	2 years ago
		11 days ago
		last month
		5 years ago
		2 months ago
		11 days ago
		11 days ago
		last month
		15 days ago

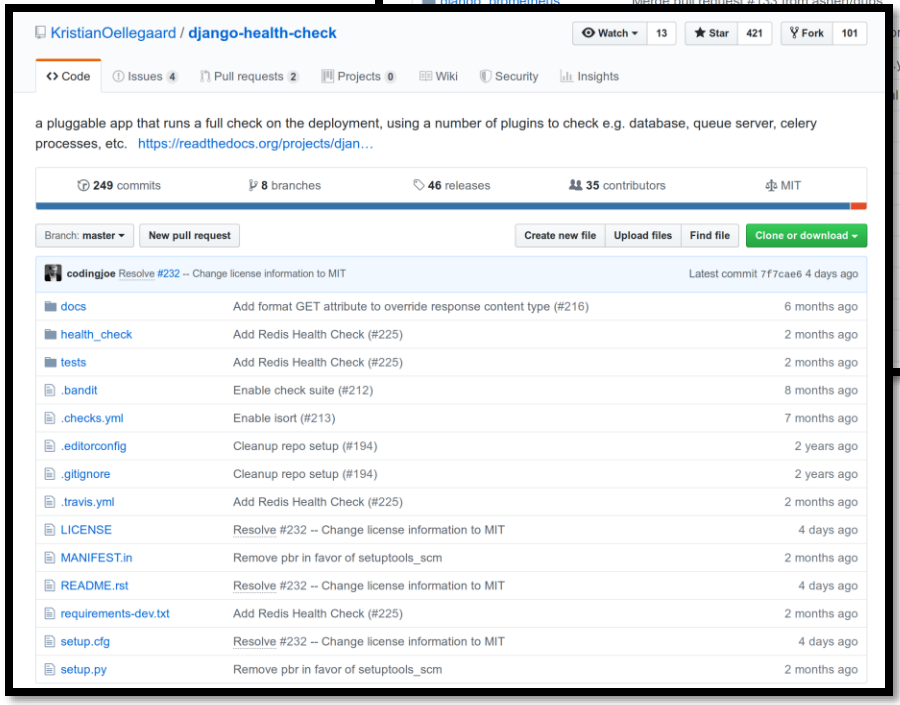

KristianOellegaard / django-health-check

a pluggable app that runs a full check on the deployment, using a number of plugins to check e.g. database, queue server, celery processes, etc. https://readthedocs.org/projects/djan…

249 commits | 8 branches | 46 releases | 35 contributors | MIT

codingjoe Resolve #232 -- Change license information to MIT Latest commit 7f7cae6 4 days ago

docs	Add format GET attribute to override response content type (#216)	6 months ago
health_check	Add Redis Health Check (#225)	2 months ago
tests	Add Redis Health Check (#225)	2 months ago
.bandit	Enable check suite (#212)	8 months ago
.checks.yml	Enable isort (#213)	7 months ago
.editorconfig	Cleanup repo setup (#194)	2 years ago
.gitignore	Cleanup repo setup (#194)	2 years ago
.travis.yml	Add Redis Health Check (#225)	2 months ago
LICENSE	Resolve #232 -- Change license information to MIT	4 days ago
MANIFEST.in	Remove pbr in favor of setuptools_scm	2 months ago
README.rst	Resolve #232 -- Change license information to MIT	4 days ago
requirements-dev.txt	Add Redis Health Check (#225)	2 months ago
setup.cfg	Resolve #232 -- Change license information to MIT	4 days ago
setup.py	Remove pbr in favor of setuptools_scm	2 months ago

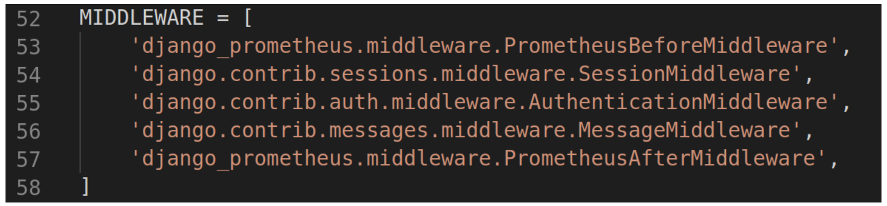

```
MIDDLEWARE = [
    'django_prometheus.middleware.PrometheusBeforeMiddleware',
    'django.contrib.sessions.middleware.SessionMiddleware',
    'django.contrib.auth.middleware.AuthenticationMiddleware',
    'django.contrib.messages.middleware.MessageMiddleware',
    'django_prometheus.middleware.PrometheusAfterMiddleware',
]
```

Prometheus Ecosystem

- Exporters we just talked about
- Grafana for graphing your metrics
- Alert Manager for generating alerts
- Libraries to instrument your code (replacing "proprietary" APM)

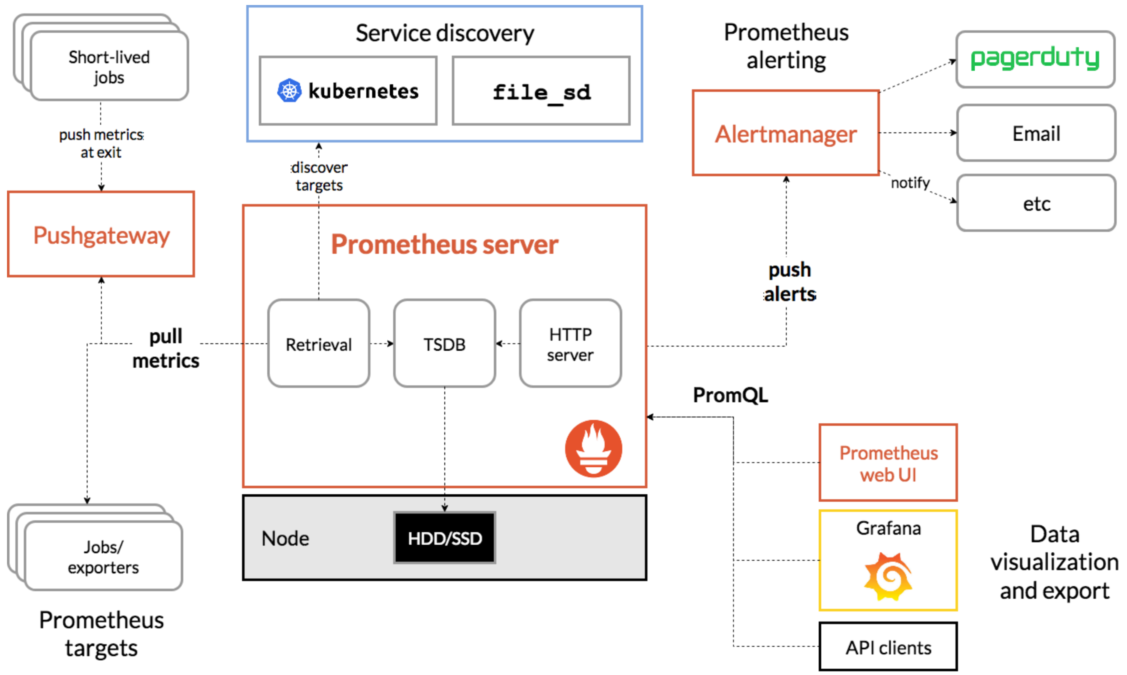

Logs management

- ELK from Elastic or Loki from Grafana Labs

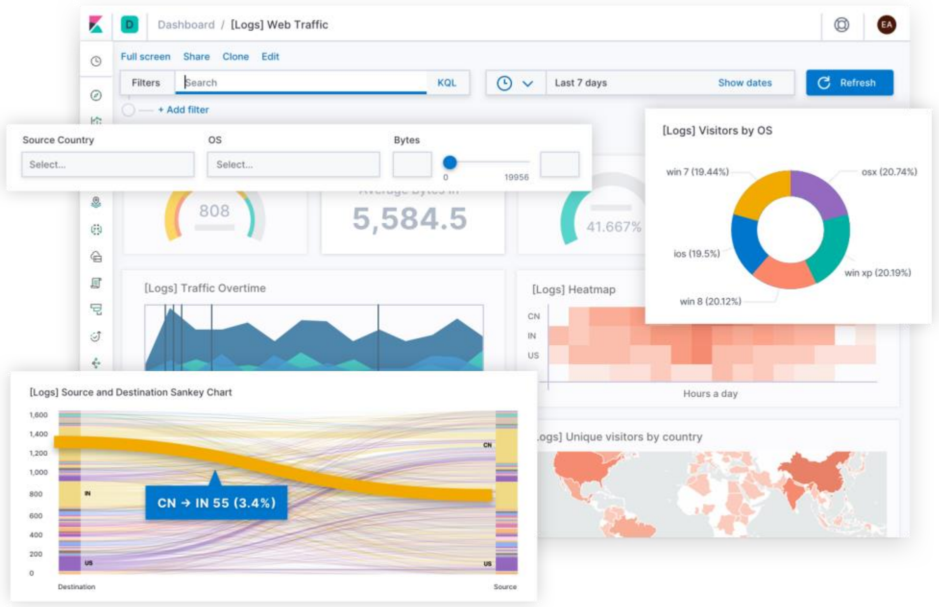

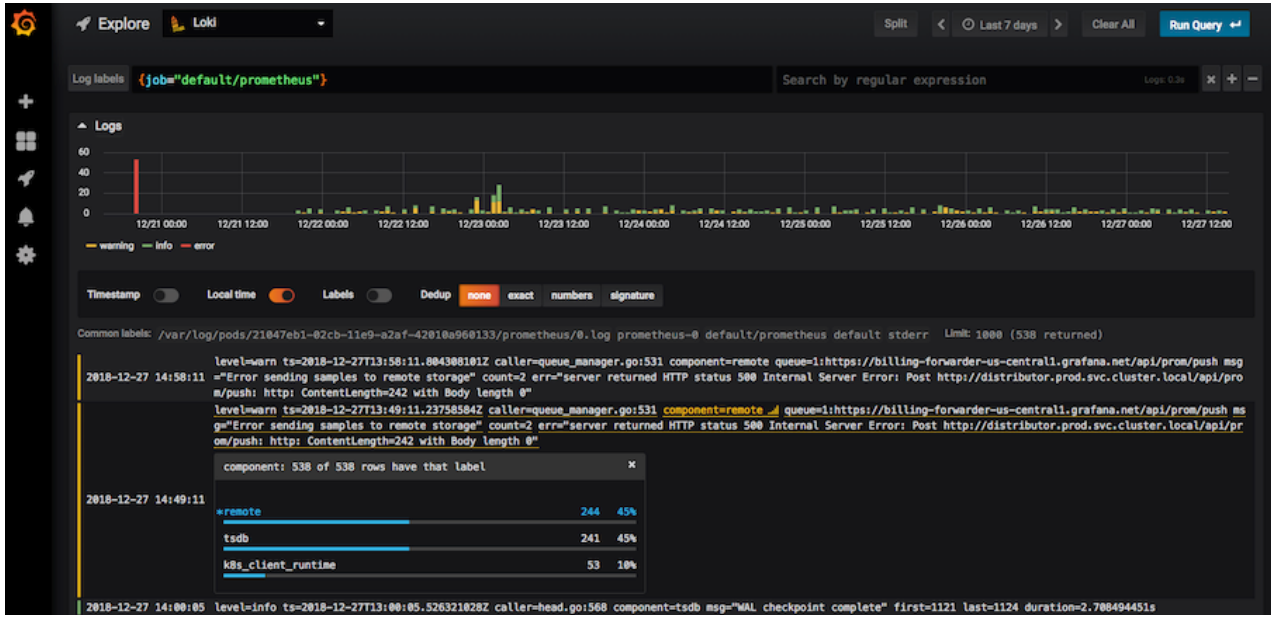

```
{job="default/prometheus"}

Common labels: /var/log/pods/21047eb1-02cb-11e9-a2af-42010a960133/prometheus/0.log prometheus-0 default/prometheus default stderr   Limit: 1000 (538 returned)

2018-12-27 14:58:11  level=warn ts=2018-12-27T13:58:11.804308101Z caller=queue_manager.go:531 component=remote queue=1:https://billing-forwarder-us-central1.grafana.net/api/prom/push msg
                     ="Error sending samples to remote storage" count=2 err="server returned HTTP status 500 Internal Server Error: Post http://distributor.prod.svc.cluster.local/api/pro
                     m/push: http: ContentLength=242 with Body length 0"
                     level=warn ts=2018-12-27T13:49:11.237585842 caller=queue_manager.go:531 component=remote queue=1:https://billing-forwarder-us-central1.grafana.net/api/prom/push ms
                     g="Error sending samples to remote storage" count=2 err="server returned HTTP status 500 Internal Server Error: Post http://distributor.prod.svc.cluster.local/api/pr
                     om/push: http: ContentLength=242 with Body length 0"
                     component: 538 of 538 rows have that label
2018-12-27 14:49:11  *remote                 244  45%
                     tsdb                    241  45%
                     k8s_client_runtime       53  10%
2018-12-27 14:00:05  level=info ts=2018-12-27T13:00:05.526321028Z caller=head.go:568 component=tsdb msg="WAL checkpoint complete" first=1121 last=1124 duration=2.708494451s
```

Tracing

- Jaeger or Zipkin
- Useful with microservices

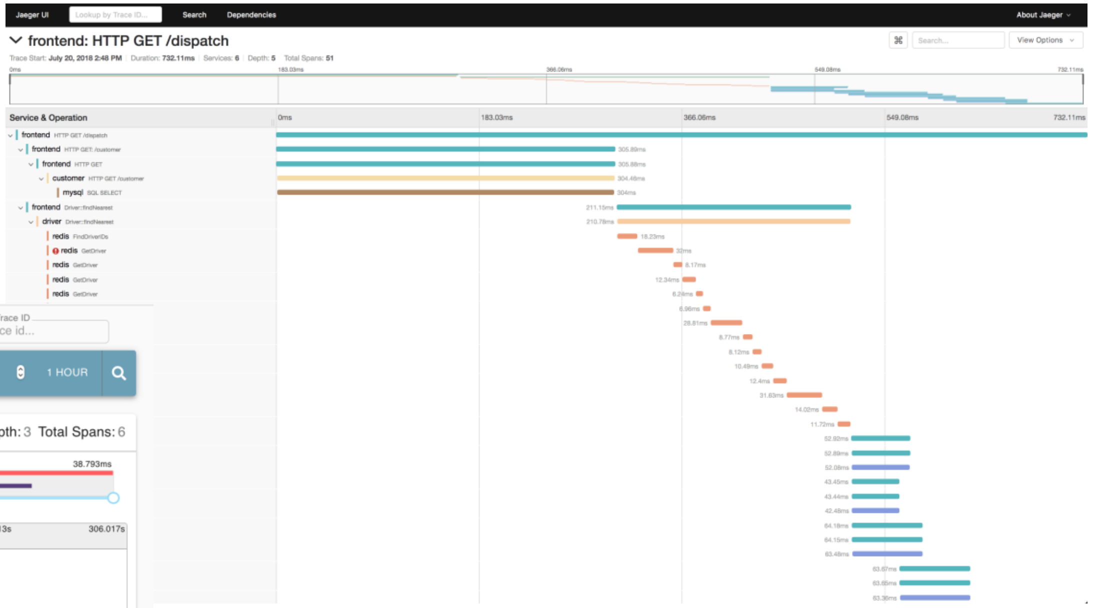

frontend: HTTP GET /dispatch

Trace Start: July 20, 2018 2:48 PM | Duration: 732.11ms | Services: 6 | Depth: 5 | Total Spans: 51

Service & Operation
frontend HTTP GET /dispatch
frontend HTTP GET: /customer
frontend HTTP GET
customer HTTP GET /customer
mysql SQL SELECT
frontend Driver::findNearest
driver Driver::findNearest
redis FindDriverIDs
redis GetDriver
redis GetDriver
redis GetDriver
redis GetDriver

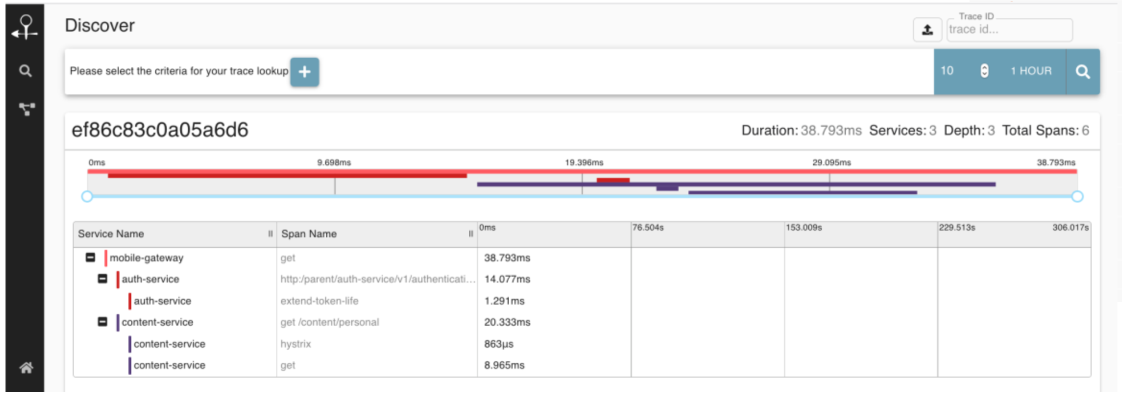

ef86c83c0a05a6d6 Duration: 38.793ms Services: 3 Depth: 3 Total Spans: 6

Service Name	Span Name	0ms
mobile-gateway	get	38.793ms
auth-service	http:/parent/auth-service/v1/authenticati...	14.077ms
auth-service	extend-token-life	1.291ms
content-service	get /content/personal	20.333ms
content-service	hystrix	863µs
content-service	get	8.965ms

More Readings

Google SRE books (2 books available online for Free):

https://landing.google.com/sre/books/

Prometheus Overview:

https://prometheus.io/docs/introduction/overview/

Next Week Hackathon

Code & Deploy an application following observability best practices.

We will provide:

- Application objectives :)
- Links to libraries and examples to help you to build your application
- Access to an AWS server (one EC2 instance per team)

https://events.bleemeo.com/bleemeo/hackathon-epitech-2020/

Questions?

☞ https://bleemeo.com/jobs