SNOWCAMP 2019

Monitoring OVH

300k servers, 27 DCs...
and one Metrics platform

Horacio Gonzalez
@LostInBrittany

OVH
Innovation for Freedom

Sommaire temporaire

- Intro we and OVH (5 minutes)
- Intro our talk (2 minutes)
- Make Better Decisions By using Numbers (5 minutes)
- Building OVH Metrics (10 minutes)
- Conclusion (2 minutes)
- Bye bye (1 minute)

OVH
Innovation for Freedom

Who are we?

Introducing myself and introducing OVH

Horacio Gonzalez

OVH
Innovation for Freedom

@LostInBrittany

Spaniard lost in Brittany,
developer, dreamer and
all-around geek

DevFest du
Bout du Monde

Finist
Devs

WARP 10
MEETUP

OVH
Team DevRel

Google Developers
Experts
2015
Web Technologies
GDE

OVH : Key Figures

OVH
Innovation for Freedom

1.3M Customers worldwide in 138 Countries
1.5 Billions euros investment over five years
30 Datacenters (growing)
350k Dedicated Servers
200k Private cloud VMs running
650k Public cloud Instances created in a month
15TB bandwidth capacity
35 Points of presence
4TB Anti DDoS capacity
Hosting capacity : 1.3M Physical Servers

+ 2 500 Employees in 19 countries
18 Years of Innovation

OVH: A Global Leader on Cloud

OVH
Innovation for Freedom

200k Private cloud VMs running

1 **Dedicated IaaS Europe**

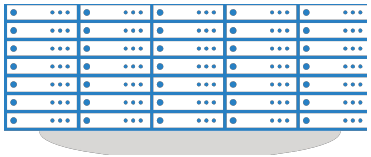

Hosting capacity : **1.3M** Physical Servers

360k Servers already deployed

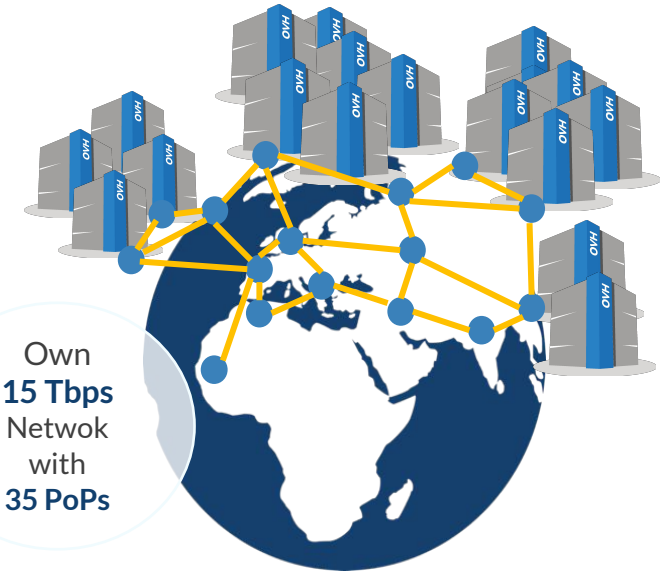

Own **15 Tbps** Netwok with **35 PoPs**

2018
27 Datacenters

2020
50 Datacenters

> **1.3M** Customers in **138** Countries

Ranking & Recognition

OVH
Innovation for Freedom

1st **European** Cloud Provider*

1st **Hosting** provider in Europe

1st **Provider** Microsoft Exchange

Certified vCloud Datacenter

Certified Kubernetes platform (CNCF)

Vmware **Global Service Provider** 2013-2016

Veeam Best Cloud Partner of the year (2018)

OVH

OVH: Our solutions

OVH
Innovation for Freedom

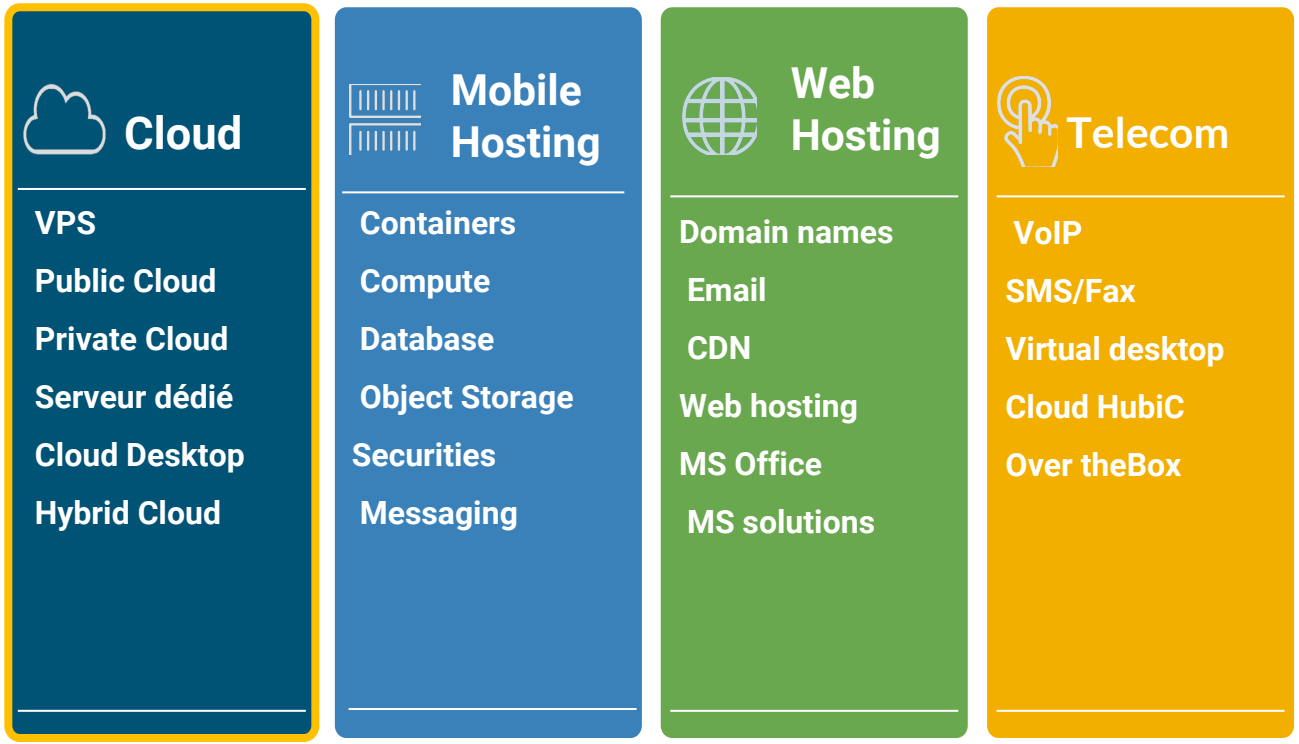

Once upon a time...

Because I love telling tales

This talk is about a tale...

OVH
Innovation for Freedom

A true one nevertheless

And as in most tales

It begins with a mission

And a band of heroes

OVH
Innovation for Freedom

Engulfed into the adventure

They fight against mishaps

OVH
Innovation for Freedom

And all kind of foes

They build mighty fortresses

OVH
Innovation for Freedom

Pushing the limits of possible

And defend them day after day

OVH
Innovation for Freedom

Against all odds

But we don't know yet the end

OVH
Innovation for Freedom

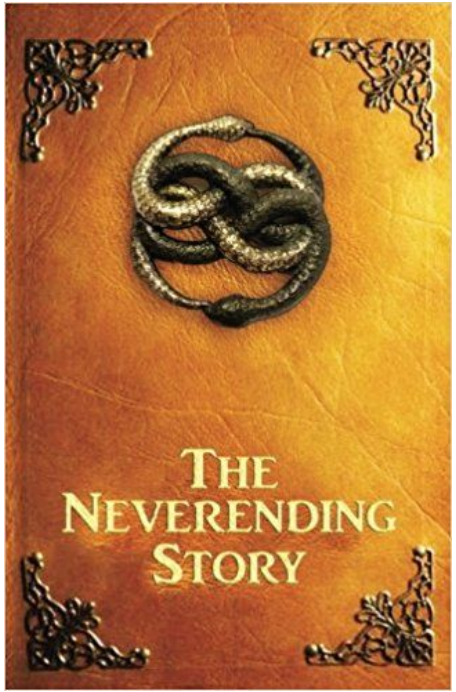

Because this tale isn't finished yet

It begins with a mission

Build a metrics platform for OVH

Why do we need metrics?

OVH
Innovation for Freedom

To make better **decisions**
by using **numbers**

Why do we need metrics?

OVH
Innovation for Freedom

We want our **code** to add **value**

Why do we need metrics?

OVH
Innovation for Freedom

We need to make better
decisions about our **code**

Why do we need metrics?

OVH
Innovation for Freedom

Code adds **value** when it **runs**
not when we write it

Why do we need metrics?

OVH
Innovation for Freedom

We need to know what our code **does** when it **runs**

Why do we need metrics?

OVH
Innovation for Freedom

We can’t do this
unless we **measure** it

Why do we need metrics?

OVH
Innovation for Freedom

We have a **mental model**
of what our code **does**

Why do we need metrics?

OVH
Innovation for Freedom

This **representation**
can be **wrong**

Why do we need metrics?

OVH
Innovation for Freedom

We can't **know** until
we **measure** it

Find the bottleneck

“The app is slow.” - User

Find the bottleneck

“The app is slow.” - User

“The page takes 500ms!” - Ops

Find the bottleneck

SQL Query?
Template Rendering?
Session Storage?

Find the bottleneck

We don't know

Find the bottleneck

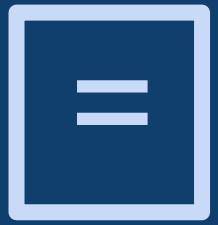

With observability:

SQL Query.....53ms

Template Rendering.....1ms

Session Storage.....315ms

Find the bottleneck

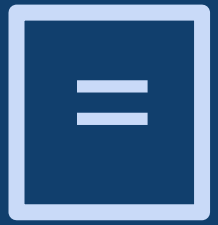

With observability:

SQL Query.....53ms

Template Rendering.....1ms

Session Storage.....315ms

Why do we need metrics?

OVH
Innovation for Freedom

We improve our mental model by **measuring** what our code **does**

Why do we need metrics?

OVH
Innovation for Freedom

We use our **mental model**
to **decide** what to do

Why do we need metrics?

OVH
Innovation for Freedom

A better **mental model** makes us better at **deciding** what to do

Why do we need metrics?

OVH
Innovation for Freedom

Better **decisions** makes us
better at generating **value**

Why do we need metrics?

OVH
Innovation for Freedom

Measuring make your App **better**

It began with a mission

OVH
Innovation for Freedom

Build a **metrics** platform for **OVH**

A metrics platform for OVH

For **all** OVH

OVH
Innovation for Freedom

Building OVH Metrics

One Platform to unify them all,
One Platform to find them,
One Platform to bring them all
and in the Metrics monitor them

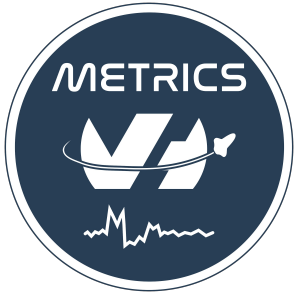

What is OVH Metrics?

OVH
Innovation for Freedom

Managed Cloud Platform
for Time Series

OVH monitoring story

OVH
Innovation for Freedom

We had lots of partial solutions...

OpenTSDB

mongoDB.

influxdb

OVH monitoring story

OVH
Innovation for Freedom

One Platform to unify them all

What should we build it on?

OVH monitoring story

OVH
Innovation for Freedom

Including a really big

OPENTSDB

OpenTSDB drawbacks

OVH
Innovation for Freedom

OpenTSDB RowKey Design

metrics timestamp tagk1 tagv1 tagk2 tagv2

OpenTSDB Rowkey design flaws

OVH Innovation for Freedom

- .*regex.* => full table scans
- High cardinality issues (Query latencies)

We needed something able to manage

hundreds of millions time series

OpenTSBD didn't **scale** for us

OpenTSDB other flaws

OVH
Innovation for Freedom

- Compaction (or append writes)
- /api/query : 1 endpoint per function?
- Asynchronous
- Unauthenticated
- ...

Scaling OpenTSDB

OVH
Innovation for Freedom

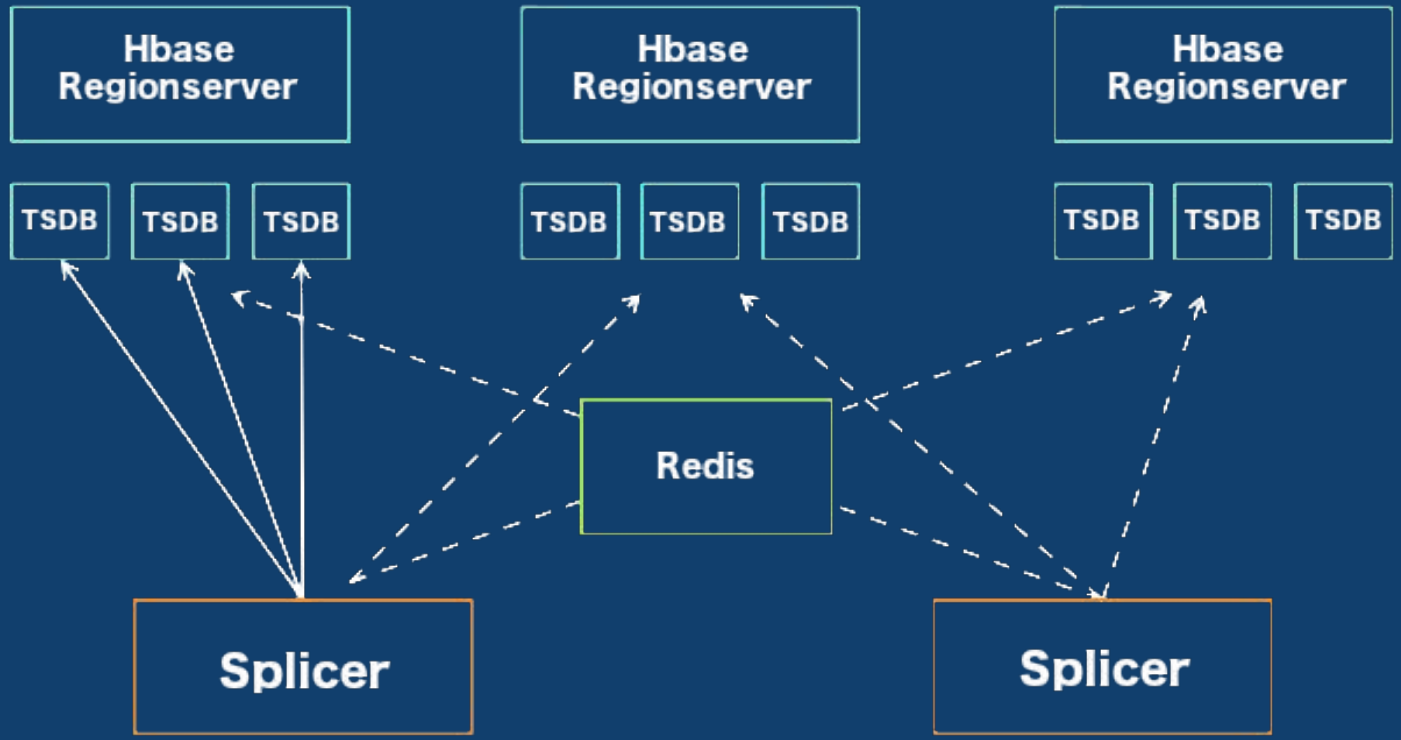

Metrics needs

OVH
Innovation for Freedom

First **need**:

To be **massively** scalable

Analytics is the key to success

OVH
Innovation for Freedom

Fetching data is only the tip of the iceberg

Analysing metrics data

OVH
Innovation for Freedom

To be scalable, analysis must be done in the database, not in user's computer

Metrics needs

OVH
Innovation for Freedom

Second **need**:

To have **rich query** capabilities

OVH
Innovation for Freedom

Enter Warp 10...

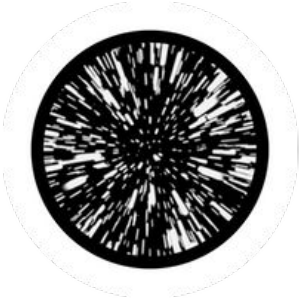

WARP10

Open-source
Time series
Database

Cityzen Data

PRODUIT EN BRETAGNE

More than a Time Series DB

OVH
Innovation for Freedom

Warp 10 is a software platform that

- Ingests and stores time series
- Manipulates and analyzes time series

WARP 10

A true Time Series analysis toolbox

- Hundreds of functions
- Manipulation frameworks
- Analysis workflow

OVH Innovation for Freedom

A Time Series manipulation language

Did you say scalability?

OVH
Innovation for Freedom

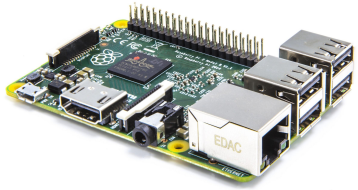

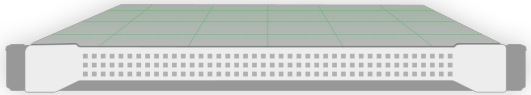

From the smallest to the largest...

More Warp 10 goodness

OVH
Innovation for Freedom

- Secured & multi tenant
- In memory Index
- No cardinality issues
- Lockfree ingestion
- WarpScript Query Language
- Support more data types
- Synchronous (transactions)
- Better Performance
- Better Scalability
- Versatile

  (standalone, distributed)

Metrics Data Platform

OVH
Innovation for Freedom

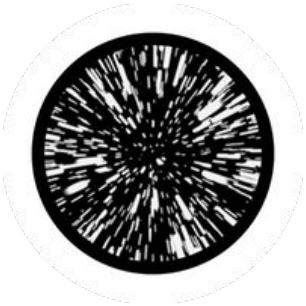

+

+

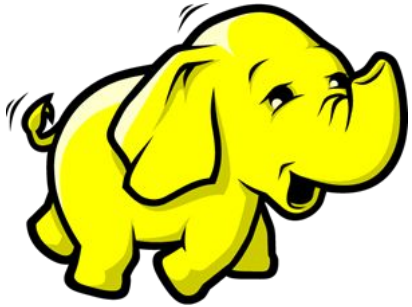

Metrics Data Platform

OVH
Innovation for Freedom

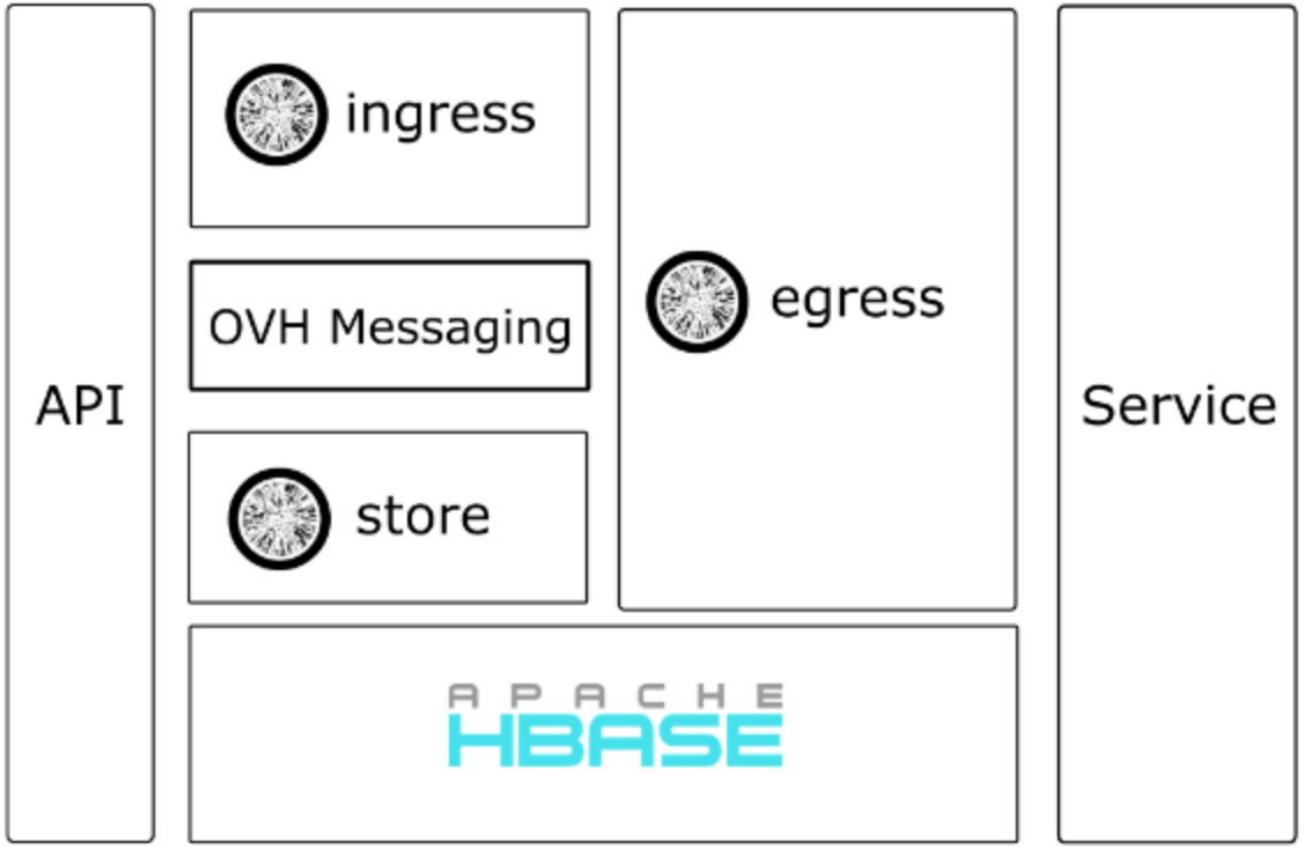

OVH
Innovation for Freedom

Building an ecosystem

From Warp 10 to OVH Metrics

OVH
Innovation for Freedom

Multi-protocol

Why to choose? We need them all!

Open source monitoring tools

OVH
Innovation for Freedom

Open source monitoring tools

OVH
Innovation for Freedom

Open source monitoring tools

OVH
Innovation for Freedom

OpenTSDB

Open source monitoring tools

OVH
Innovation for Freedom

Open source monitoring tools

OVH
Innovation for Freedom

Open source monitoring tools

OVH
Innovation for Freedom

Open source monitoring tools

OVH
Innovation for Freedom

Why choose?

Let’s support all of them!

Metrics Platform

OVH
Innovation for Freedom

Metrics Platform

OVH
Innovation for Freedom

graphite

influx

https:// **opentsdb** .<region>.metrics.ovh.net

prometheus

warp10

...

OVH
Innovation for Freedom

Metrics Live

In-memory, high-performance Metrics instances

In-memory: Metrics live

OVH
Innovation for Freedom

+120 million of writes/s

In-memory: Metrics live

OVH
Innovation for Freedom

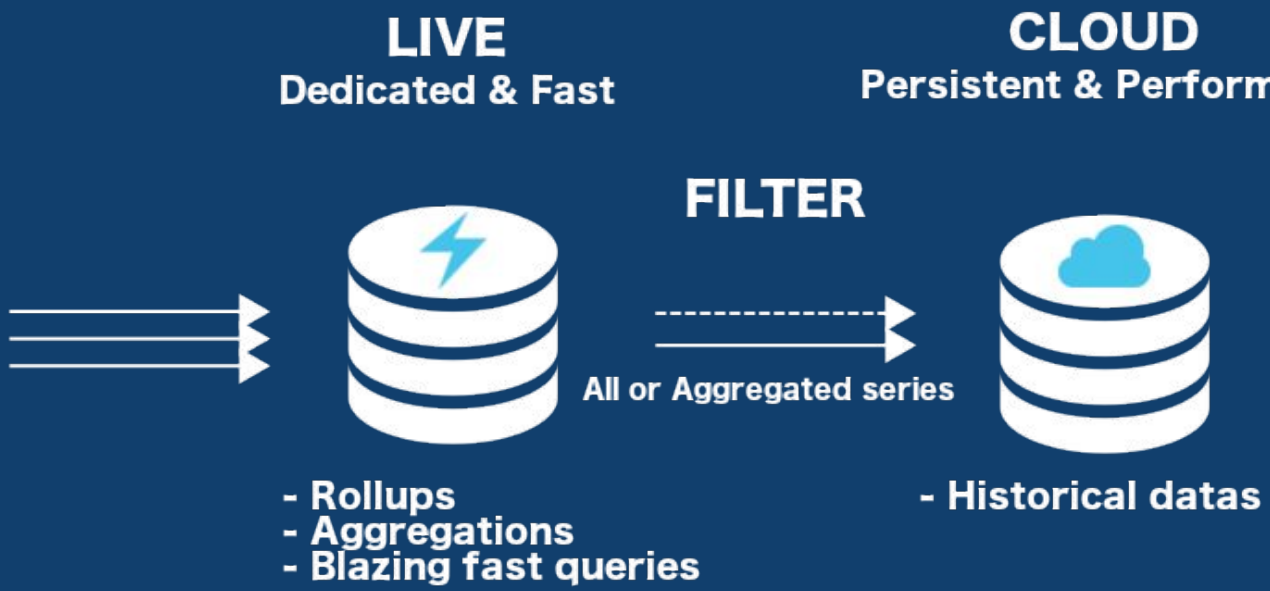

In-memory: Metrics live

OVH
Innovation for Freedom

STAGE 1
- Short retention - hours
- Fine grained monitoring
- Raw data

STAGE 2
- Short retention - days
- Consolidated aggregations
- Global infra monitoring

STAGE 3
- Customer metrics
- Historical datas

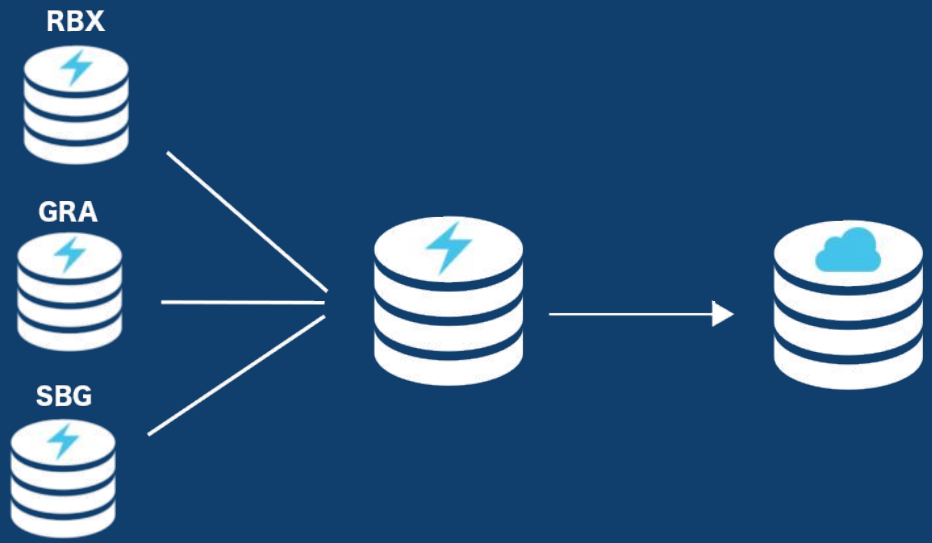

OVH
Innovation for Freedom

Monitoring is only the beginning

OVH Metrics answer to many other use cases

Use cases families

- Billing ..... (e.g. bill on monthly max consumption)
- Monitoring ..... (APM, infrastructure,appliances,...)
- IoT ..... (Manage devices, operator integration, ...)
- Geo Location ..... (Manage localized fleets)

Use cases

- DC Temperature/Elec/Cooling map
- Pay as you go billing (PCI/IPLB)
- GSCAN
- Monitoring
- ML Model scoring (Anti-Fraude)
- Pattern Detection for medical applications

OVH
Innovation for Freedom

SREing Metrics

With a great power comes a great responsibility

Metrics' own metrics

432 000 000 000
datapoints / day

Metrics' own metrics

10 Tb / day

Metrics' own metrics

5 000 000 dp/s

Metrics' own metrics

500 000 000 series

Our clusters size

OVH
Innovation for Freedom

GRA:

- 150 nodes
- 2 PB
- 1.1 Gbps

BHS:

- 30 nodes
- 400 TB
- 120 Mbps

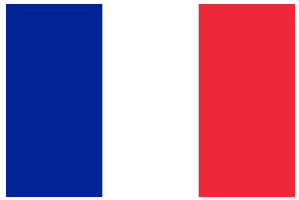

Our cluster architecture

OVH
Innovation for Freedom

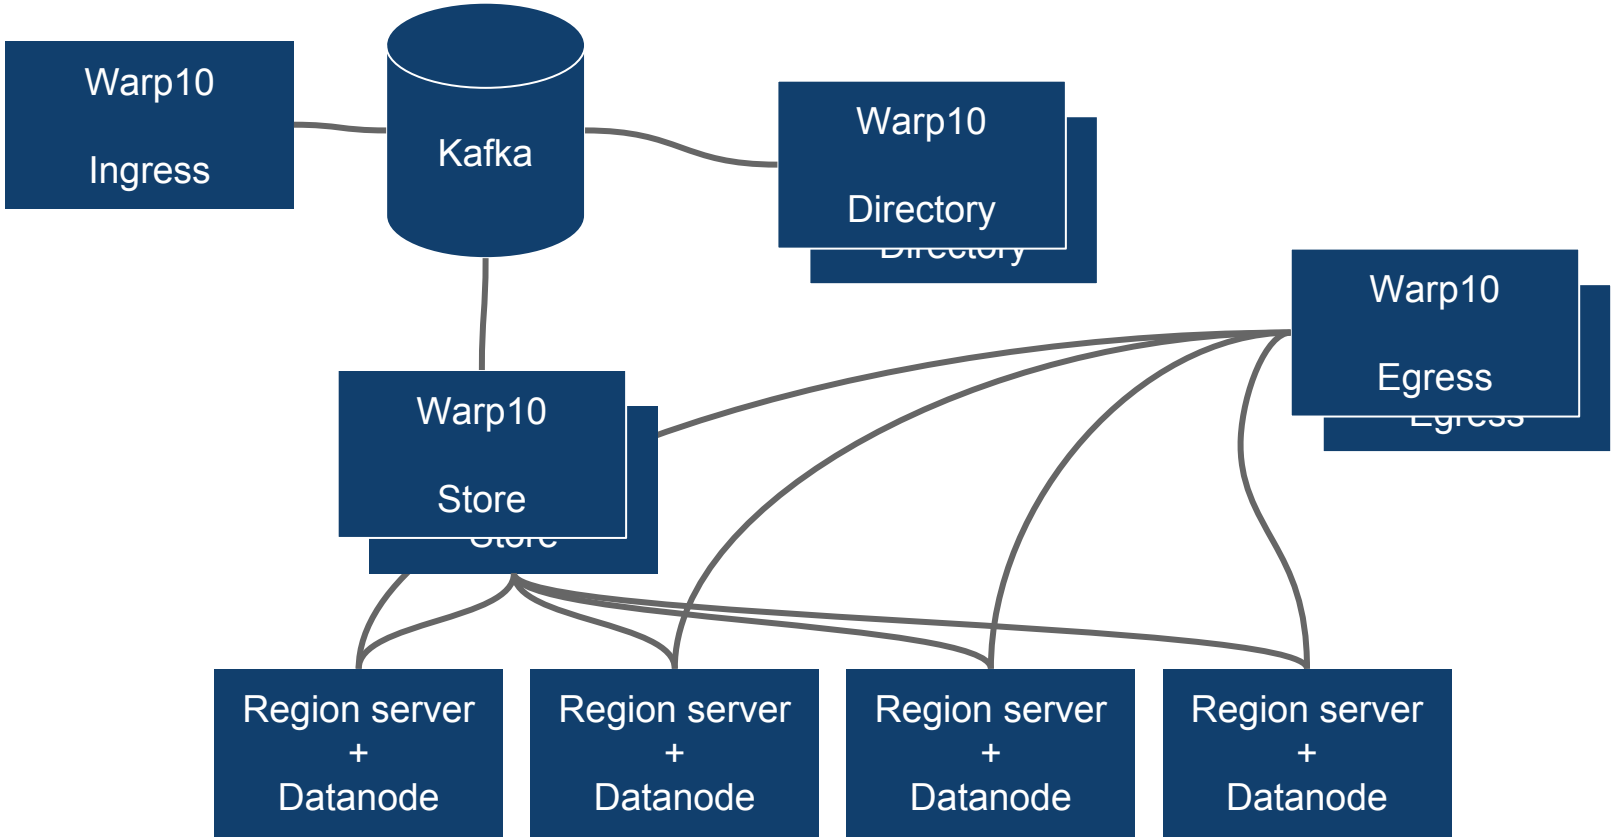

Detecting errors

Before it's too late

Extract errors from logs

OVH
Innovation for Freedom

Tailor

Forward logs and extract metrics!

Monitoring the JVM

OVH
Innovation for Freedom

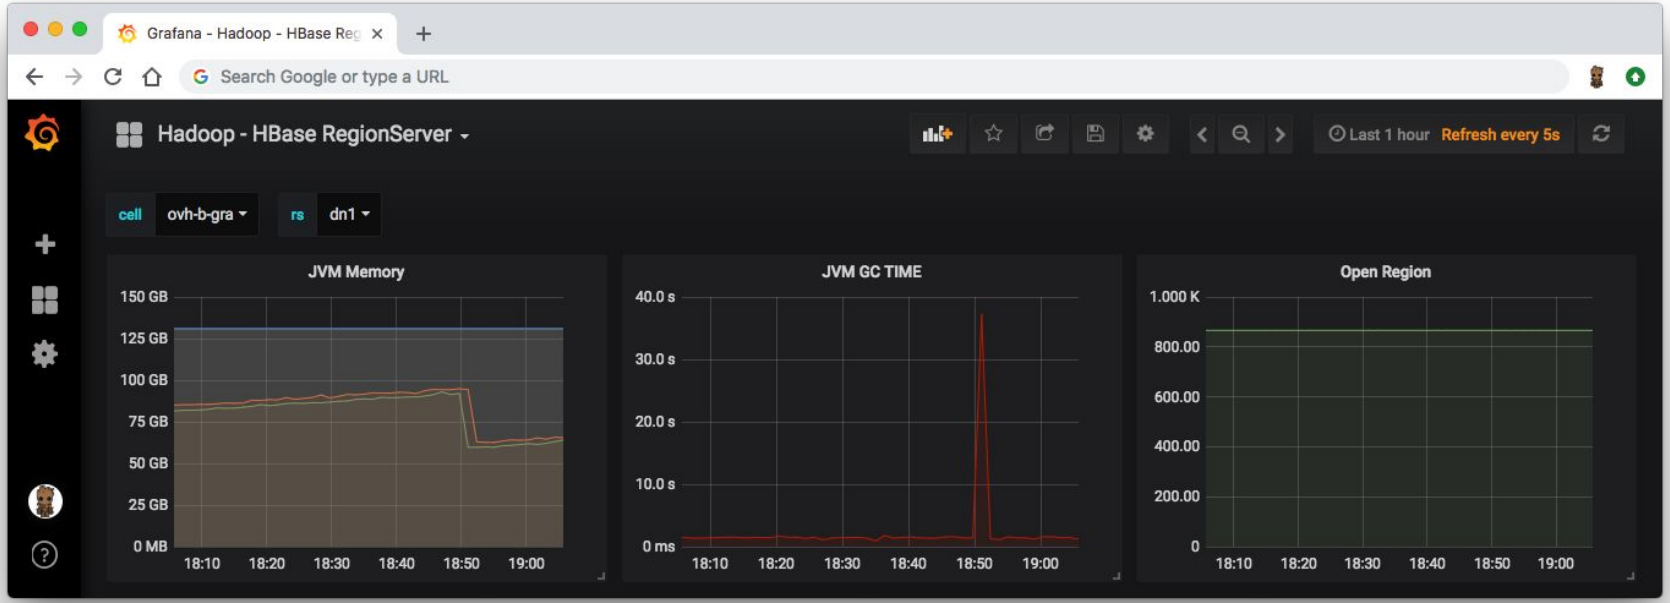

Documentation

OVH
Innovation for Freedom

OVH
Innovation for Freedom

JVM GC

The good, the bad and the ugly

The good

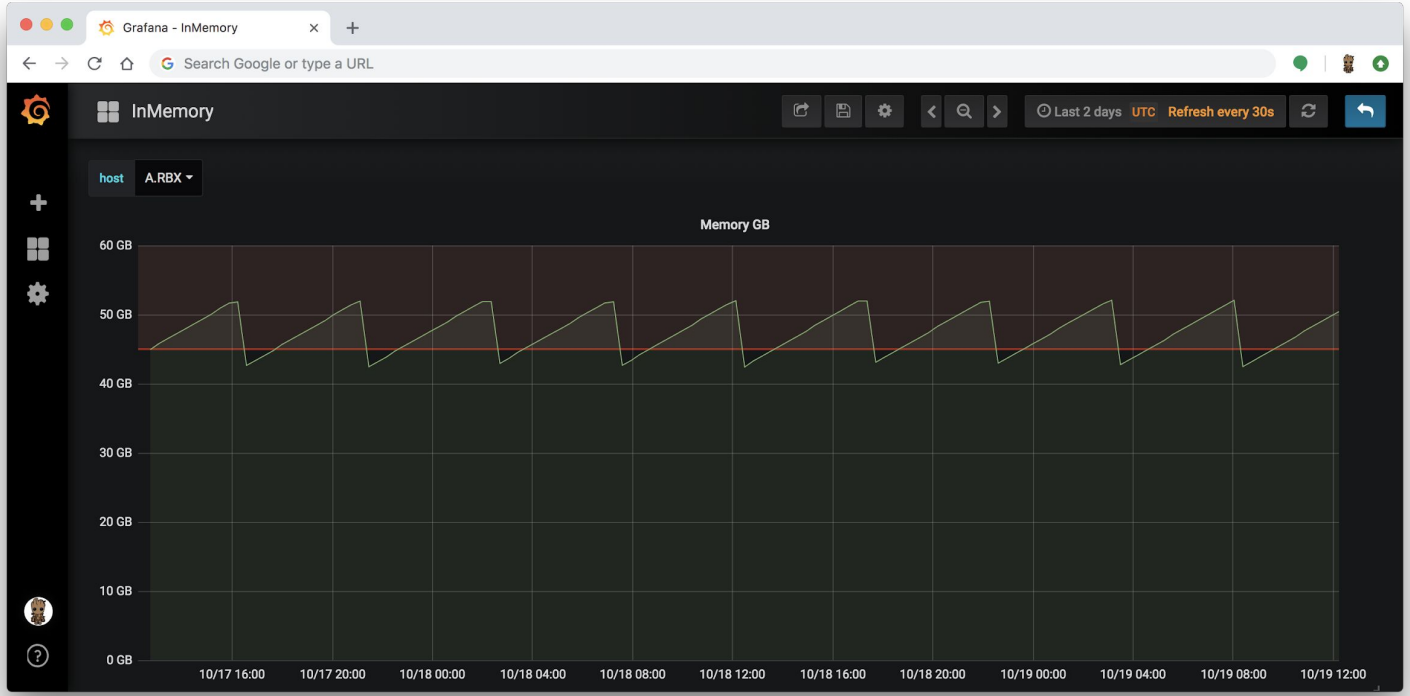
Grafana - InMemory
Search Google or type a URL
InMemory
Last 2 days
UTC
Refresh every 30s
host
A.RBX
Memory GB
60 GB
50 GB
40 GB
30 GB
20 GB
10 GB
0 GB
10/17 16:00
10/17 20:00
10/18 00:00
10/18 04:00
10/18 08:00
10/18 12:00
10/18 16:00
10/18 20:00
10/19 00:00
10/19 04:00
10/19 08:00
10/19 12:00

The bad

OVH
Innovation for Freedom

… and the ugly

#java #jdk11 #zgc

Monitoring HBase

OVH
Innovation for Freedom

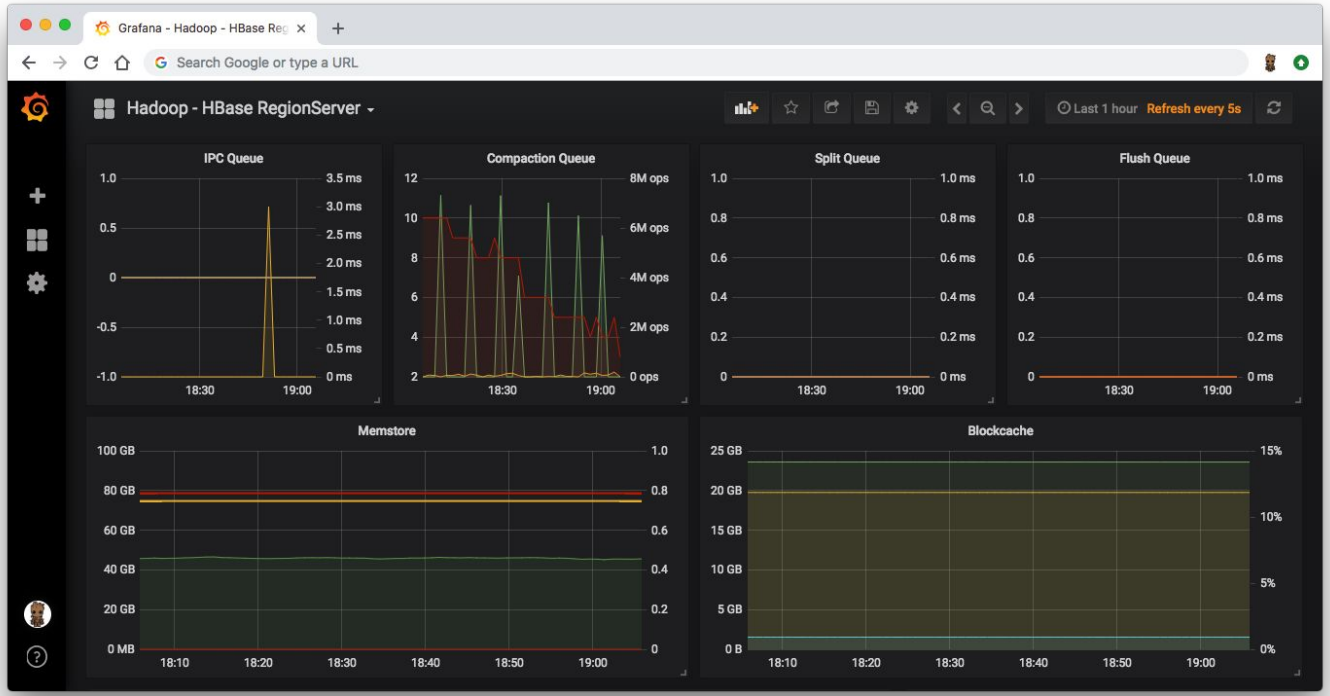

Number of open regions

OVH
Innovation for Freedom

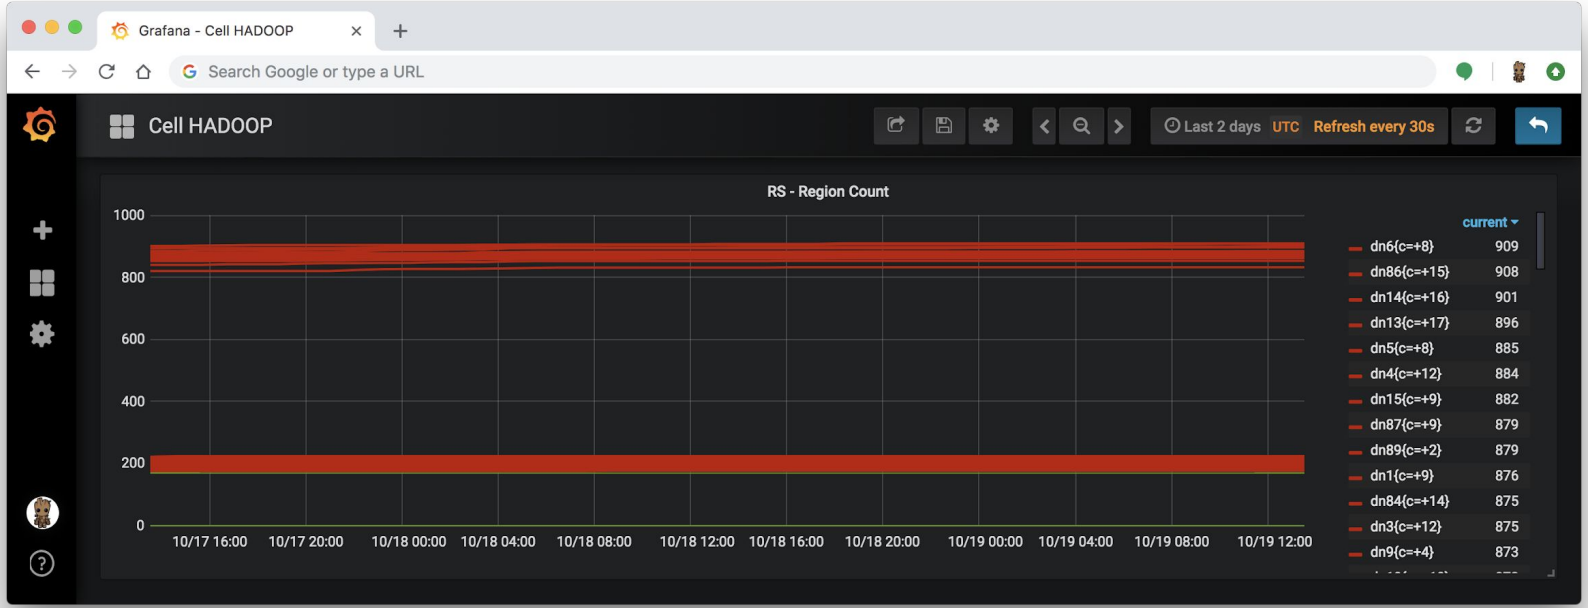

Queues length

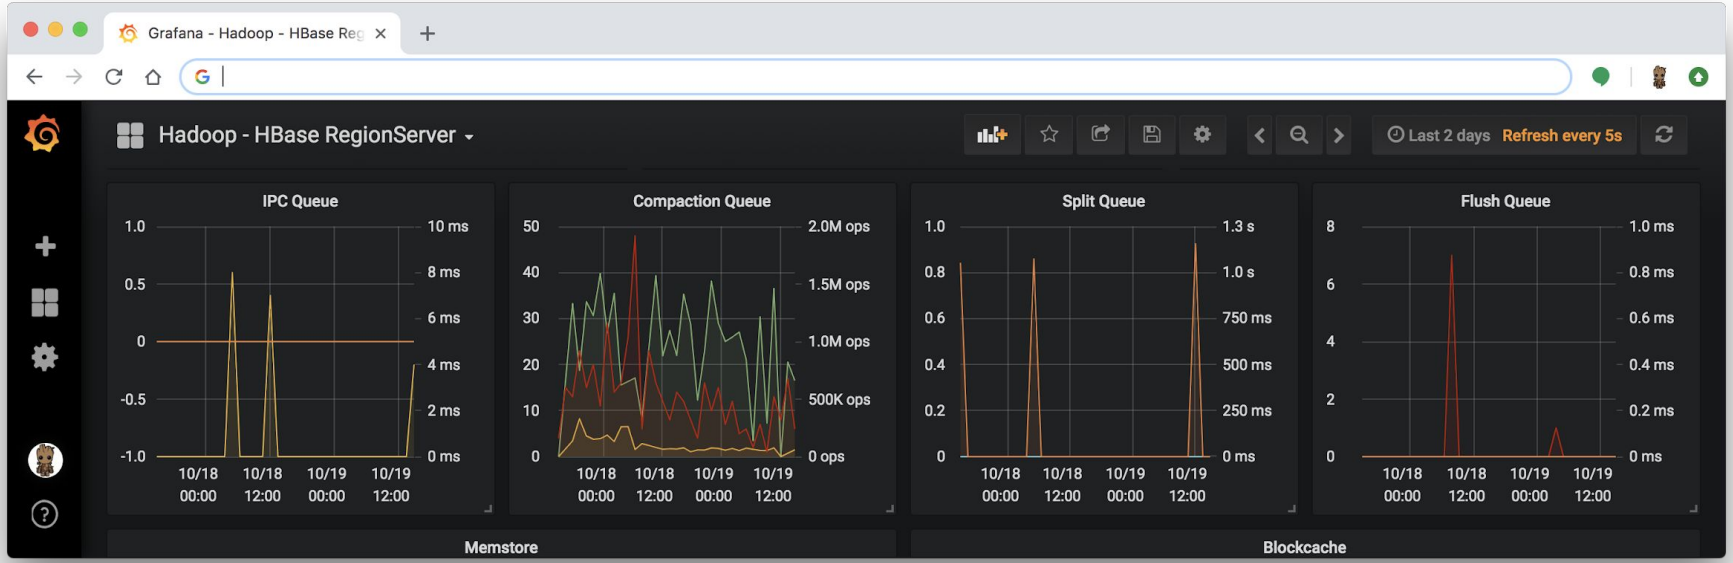

Number of read and write requests

OVH
Innovation for Freedom

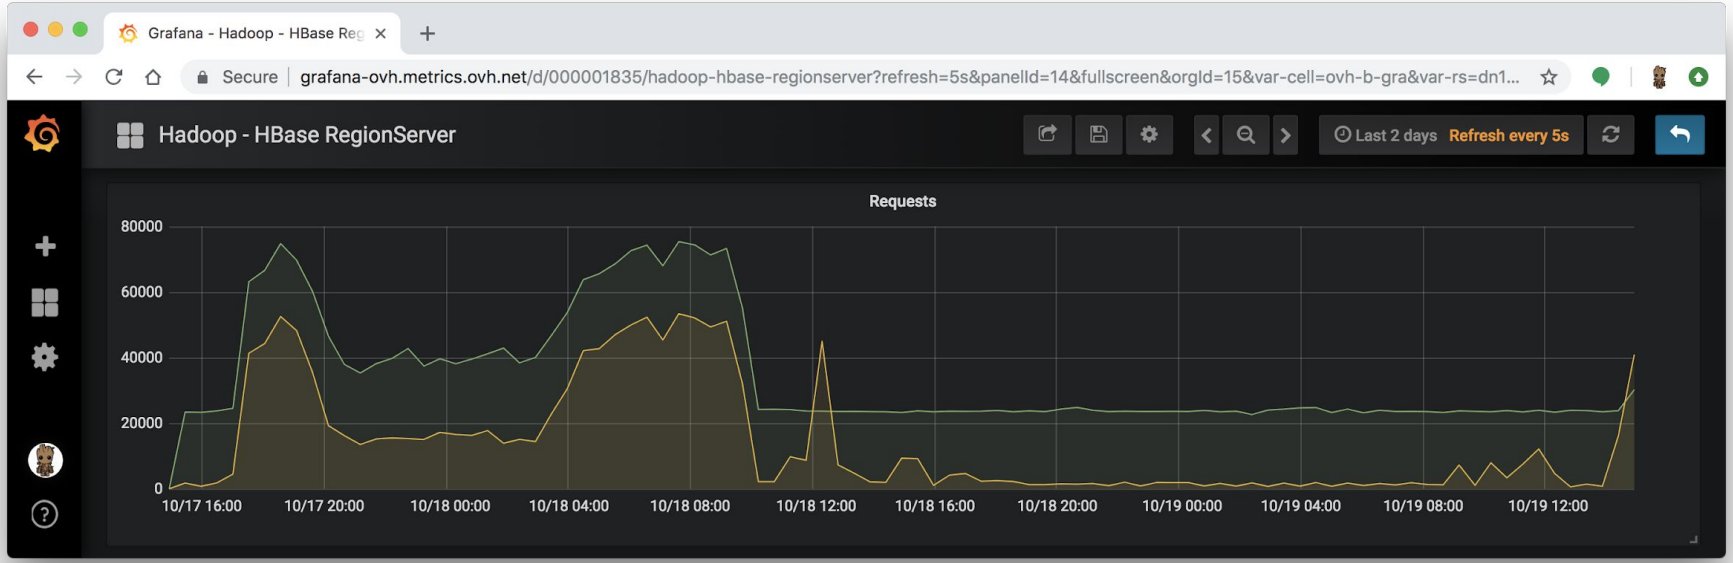

Preserve data locality

OVH
Innovation for Freedom

Host health

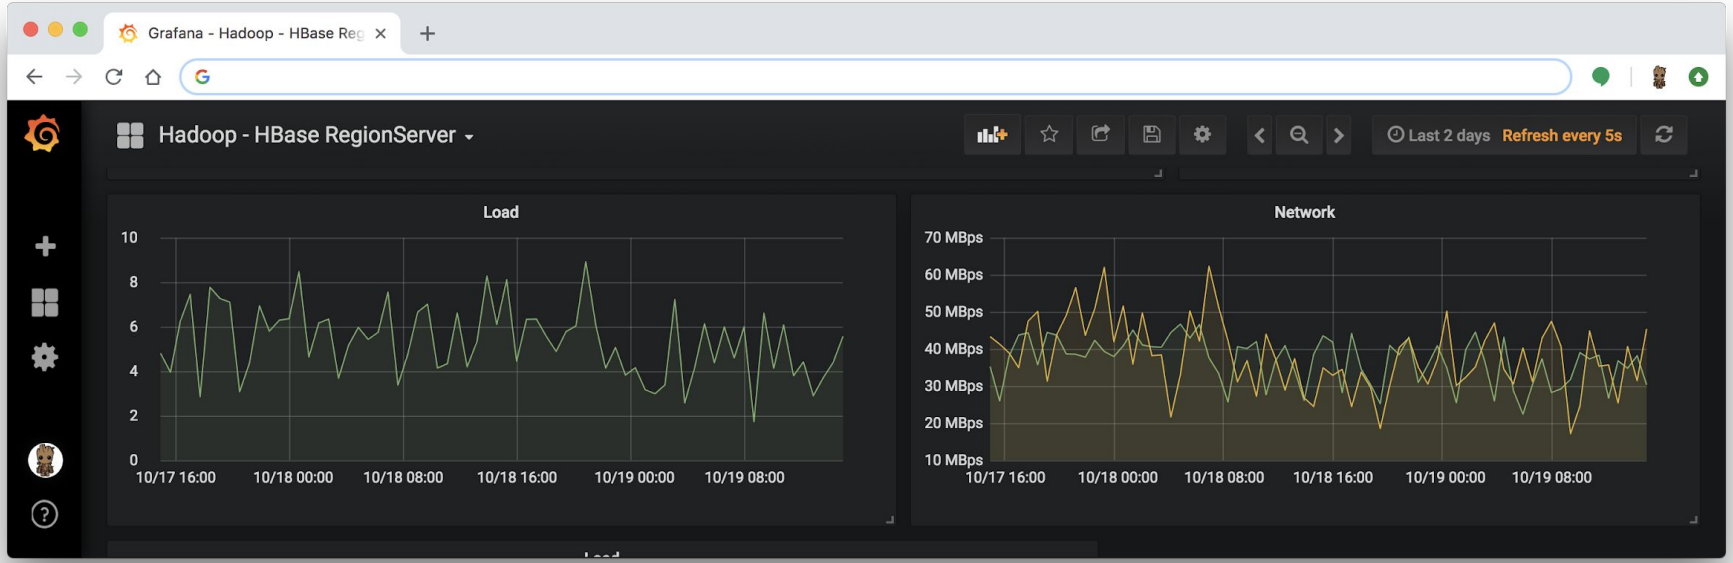
Grafana - Hadoop - HBase Reg
Hadoop - HBase RegionServer
Last 2 days
Refresh every 5s
Load
10
8
6
4
2
0
10/17 16:00
10/18 00:00
10/18 08:00
10/18 16:00
10/19 00:00
10/19 08:00
Network
70 MBps
60 MBps
50 MBps
40 MBps
30 MBps
20 MBps
10 MBps
10/17 16:00
10/18 00:00
10/18 08:00
10/18 16:00
10/19 00:00
10/19 08:00

OVH
Innovation for Freedom

Pokédex

Inventory all animals.

Merging all data sources

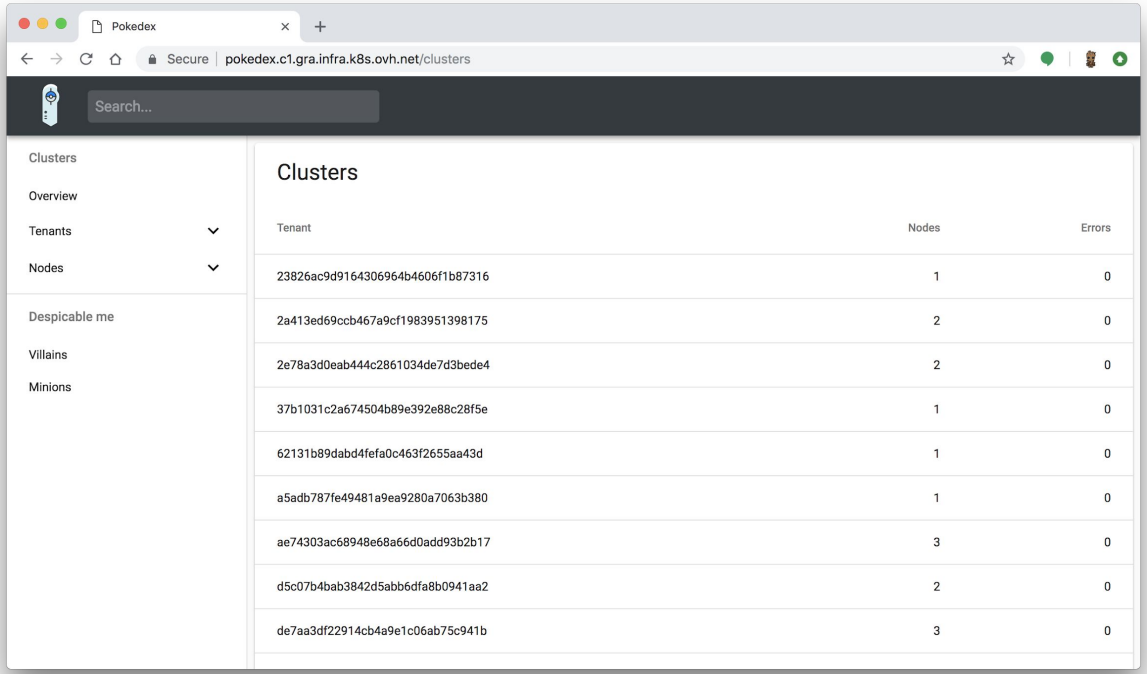

Global visualization

OVH
Innovation for Freedom

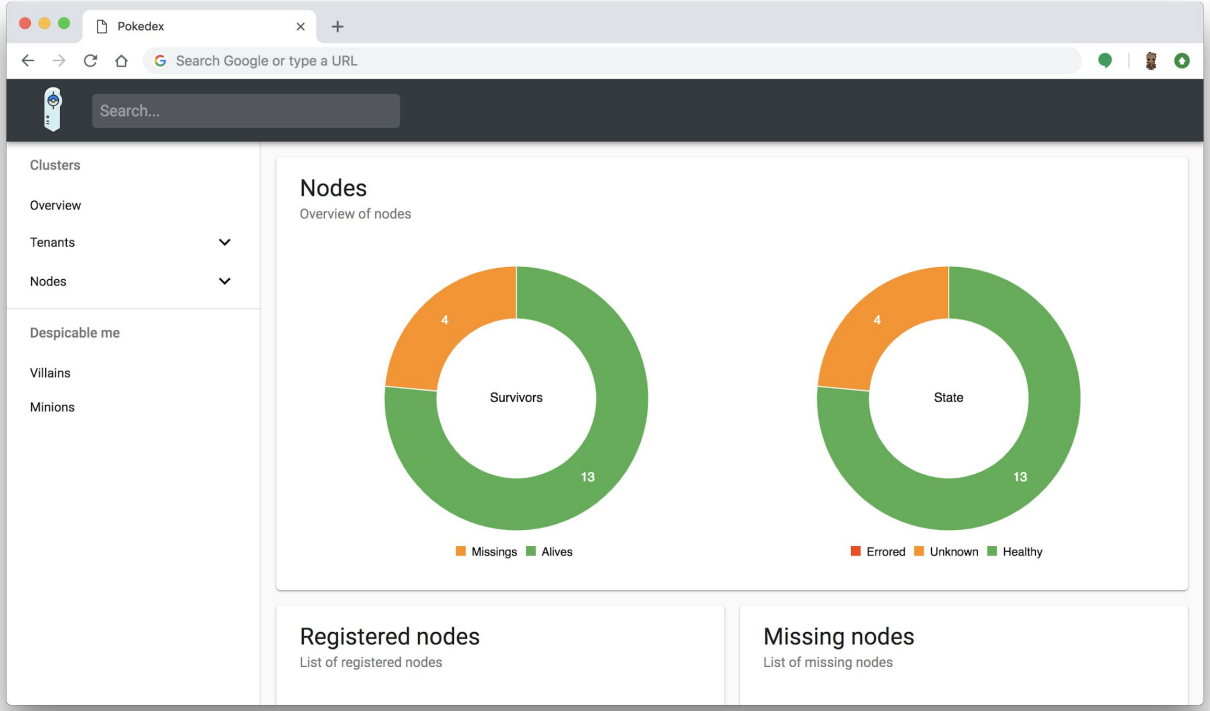

Correlate information

OVH
Innovation for Freedom

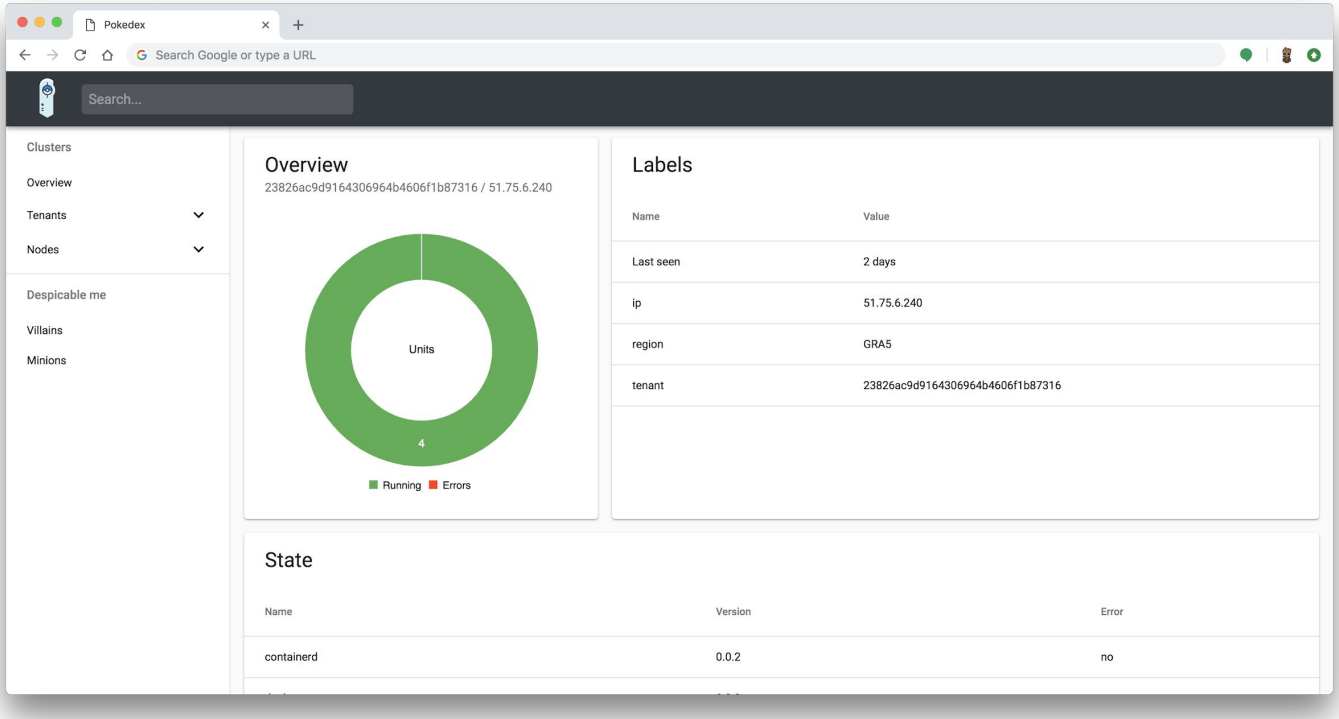

OVH
Innovation for Freedom

Sacha

The best tamer

An awesome CLI

OVH
Innovation for Freedom

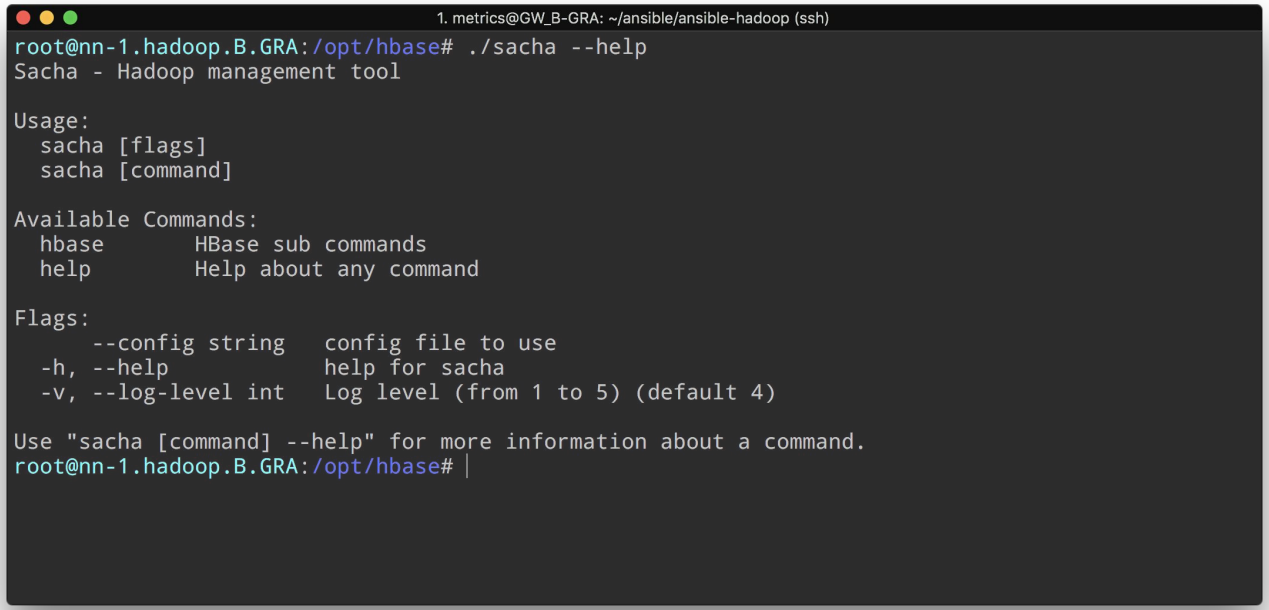

Retrieving bare informations

OVH
Innovation for Freedom

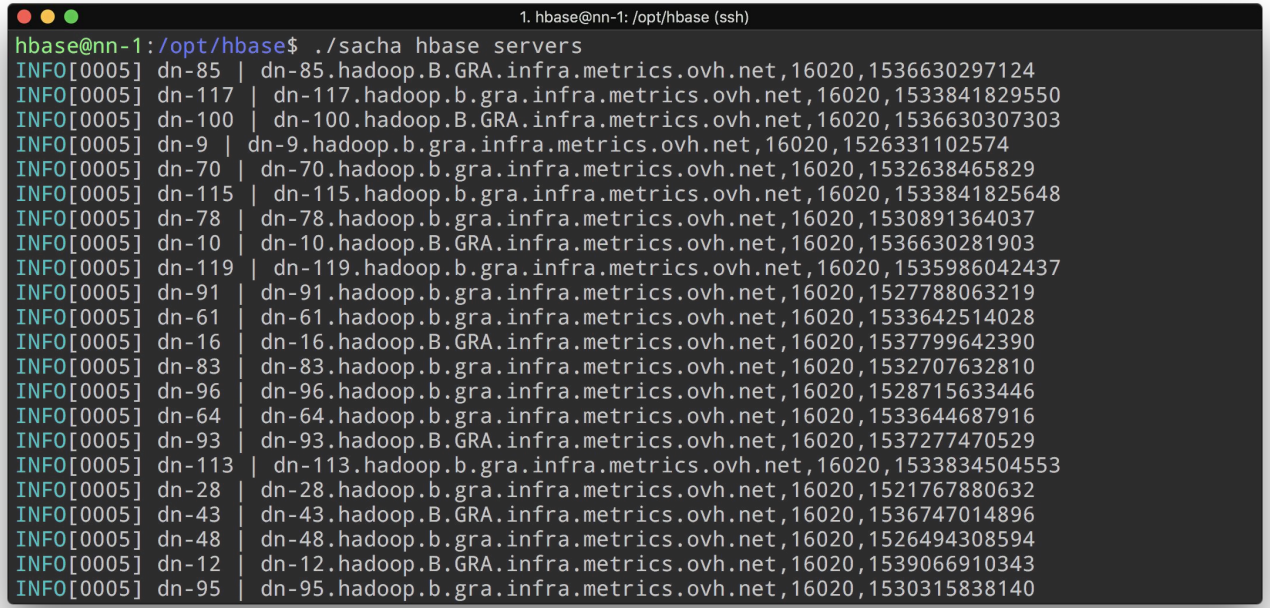

Create region map

OVH
Innovation for Freedom

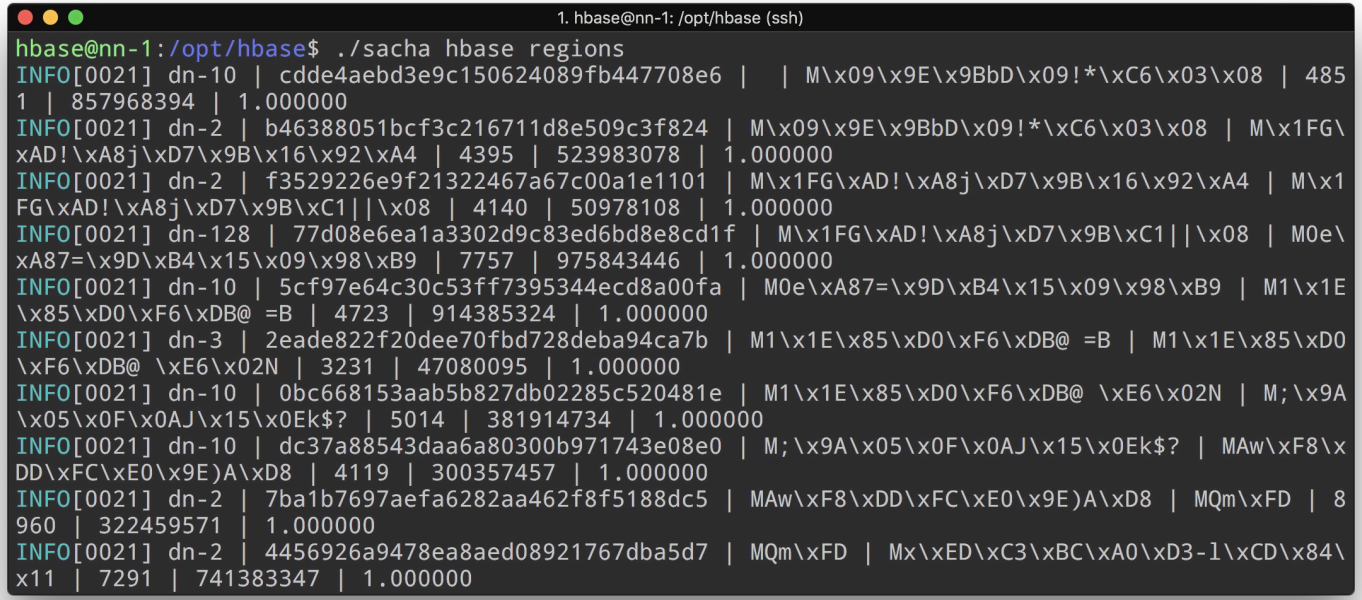

Move region to another region server

OVH
Innovation for Freedom

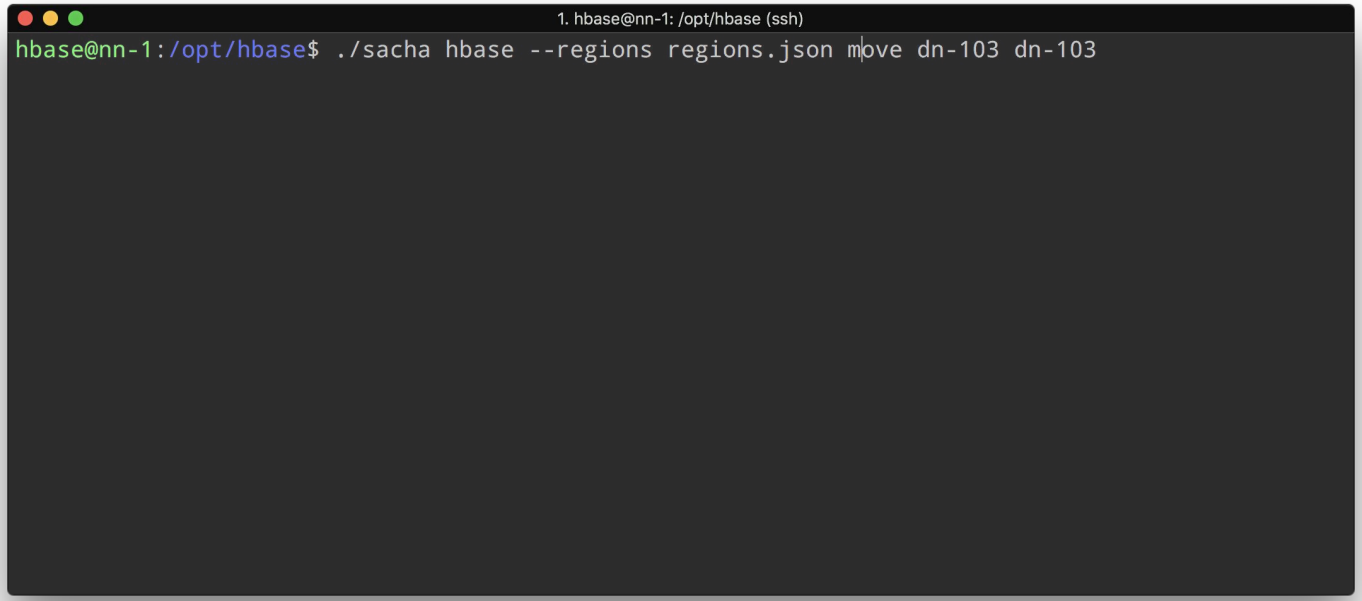

Drain regions of the region server

OVH
Innovation for Freedom

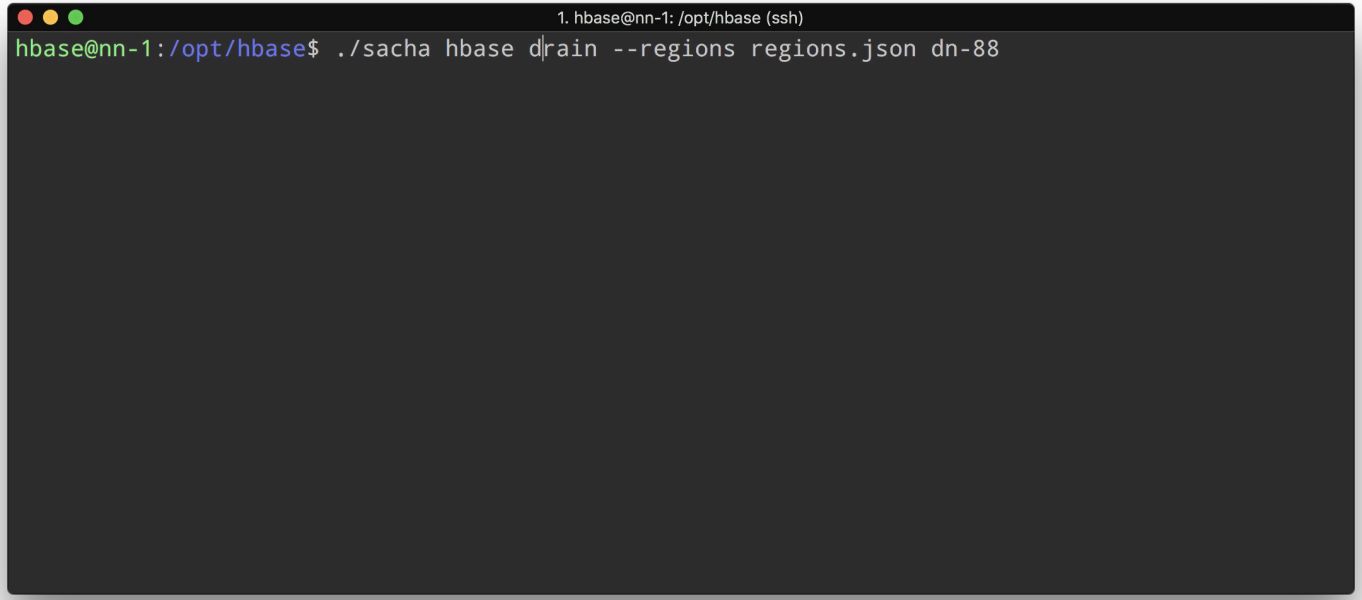

Managing multiple hardware profiles

OVH
Innovation for Freedom

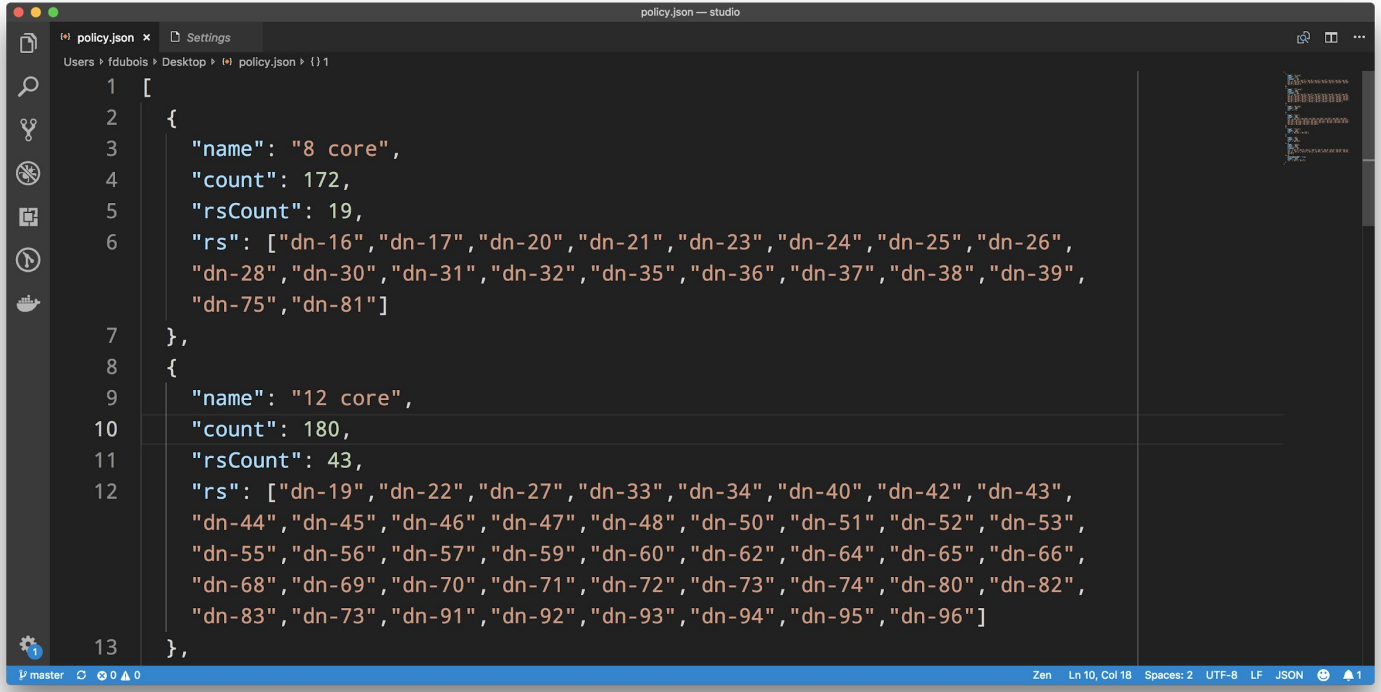

Balance the cluster

OVH
Innovation for Freedom

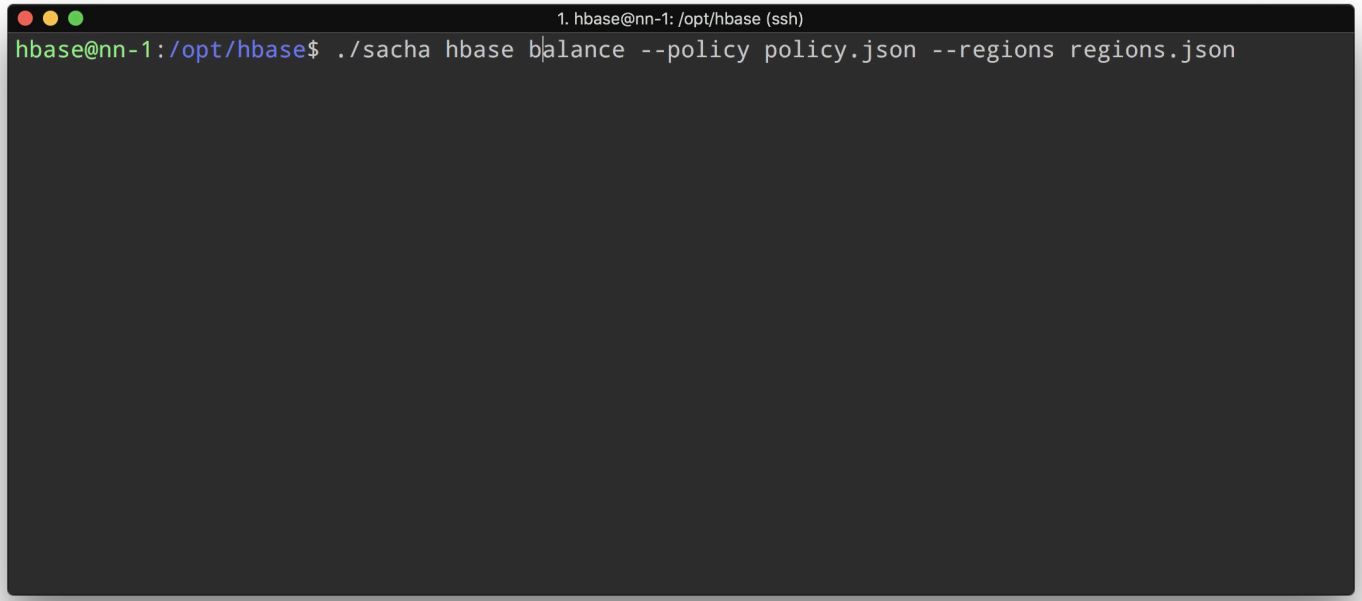

OVH
Innovation for Freedom

Conclusion

That's all folks!