

DevOpsCon

**Migrate 3 million websites
without anybody noticing**

Vincent Cassé & Horacio Gonzalez

2020-10-13




@vcasse




@LostInBrittany





Who are we?

Introducing ourselves and
introducing ~~OVH~~ OVHcloud



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Vincent Cassé



@vcasse

Host for millions of websites.
Breakfast and HTTPS included.

Engineering Manager at





DevOpsCon

Horacio Gonzalez


@LostInBrittany

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



DevOpsCon

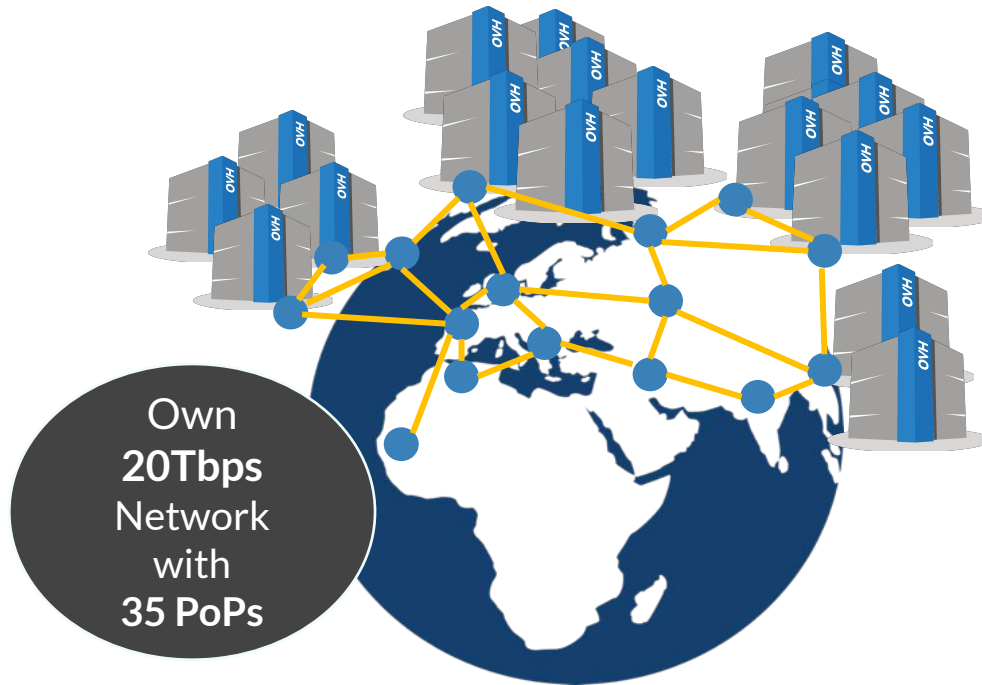

OVHcloud: A Global Leader



200k Private cloud
VMs running



1 Dedicated IaaS
Europe




Own
20Tbps
Network
with
35 PoPs

30 Datacenters

Hosting capacity:
1.3M Physical
Servers


360k
Servers already
deployed

> 1.3M Customers in 138 Countries



DevOpsCon

OVHcloud: 4 Universes of Products



WebCloud

- Domain / Email**
 - Domain names, DNS, SSL, Redirect
 - Email, Open-Xchange, Exchange
 - Collaborative Tools, NextCloud
- PaaS for Web**
 - Mutu, CloudWeb
 - Plesk, CPANEL
 - PaaS with Platform.sh
- Virtual servers**
 - VPS, Dedicated Server
- SaaS**
 - Wordpress, Magento, Prestashop
 - CRM, Billing, Payment, Stats
 - MarketPlace
- Support, Managed**
 - Support Basic
 - Support thought Partners
 - Managed services

Baremetal Cloud

- Standalone, Cluster**
 - General Purpose
 - SuperPlan
 - Game
 - T2 >20e
 - Virtualization
 - T3 >80e
 - Storage
 - T4 >300e
 - Database
 - T5 >600e
 - Bigdata
 - 12KVA /32KVA
 - HCI
 - AI
 - VDI Cloud Game
 - Network
- VPS aaS**
 - pCC DC
 - Virtuozzo Cloud
- Wholesales**
 - IT Integrators, Cloud Storage,
 - CDN, Database, ISV, WebHosting
 - High Intensive CPU/GPU,
- Encrypt**
 - KMS, HSM
 - Encrypt (SGX, Network, Storage)

Public Cloud

- Compute**
 - VM
 - K8S, IA IaaS
 - Baremetal
 - PaaS for DevOps
- Storage**
 - File, Block, Object, Archive
- Databases**
 - SQL, noSQL, Messaging,
 - Dashboard
- Network**
 - IP FO, NAT, LB, VPN, Router,
 - DNS, DHCP, TCP/SSL Offload
- Security**
 - IAM, MFA, Encrypt, KMS
- IA, DL**
 - Standard Tools for AI, AI Studio,
 - IA IaaS, Hosting API AI
- Bigdata, ML, Analytics**
 - Datalake, ML, Dashboard

Hosted Private Cloud


- Hosted Private Cloud**
 - VMware**
 - SDDC, vSAN 1AZ / 2AZ
 - vCD, Tanzu, Horizon, DBaaS,
 - DRaaS
 - Nutanix**
 - HCI 1AZ / 2AZ, Databases,
 - DRaaS, VDI
 - OpenStack**
 - IAM, Compute (VM, K8S)
 - Storage, Network, Databases
 - Storage**
 - Ontap Select, Nutanix File
 - OpenIO, MinIO, CEPH
 - Zerto, Veeam, Atempo
 - AI**
 - ElementAI, HuggingFace,
 - Deepopmatic, Systran,
 - EarthCube
 - Bigdata / Analytics / ML**
 - Cloudera over S3, Dataiku,
 - Saagie, Tableau,
 - Hybrid Cloud**
 - vRack Connect, Edge-DC, Private DC
 - Dell, HP, Cisco, OCP, MultiCloud
 - Secured Cloud**
 - GOV, FinTech, Retail, HealtCare



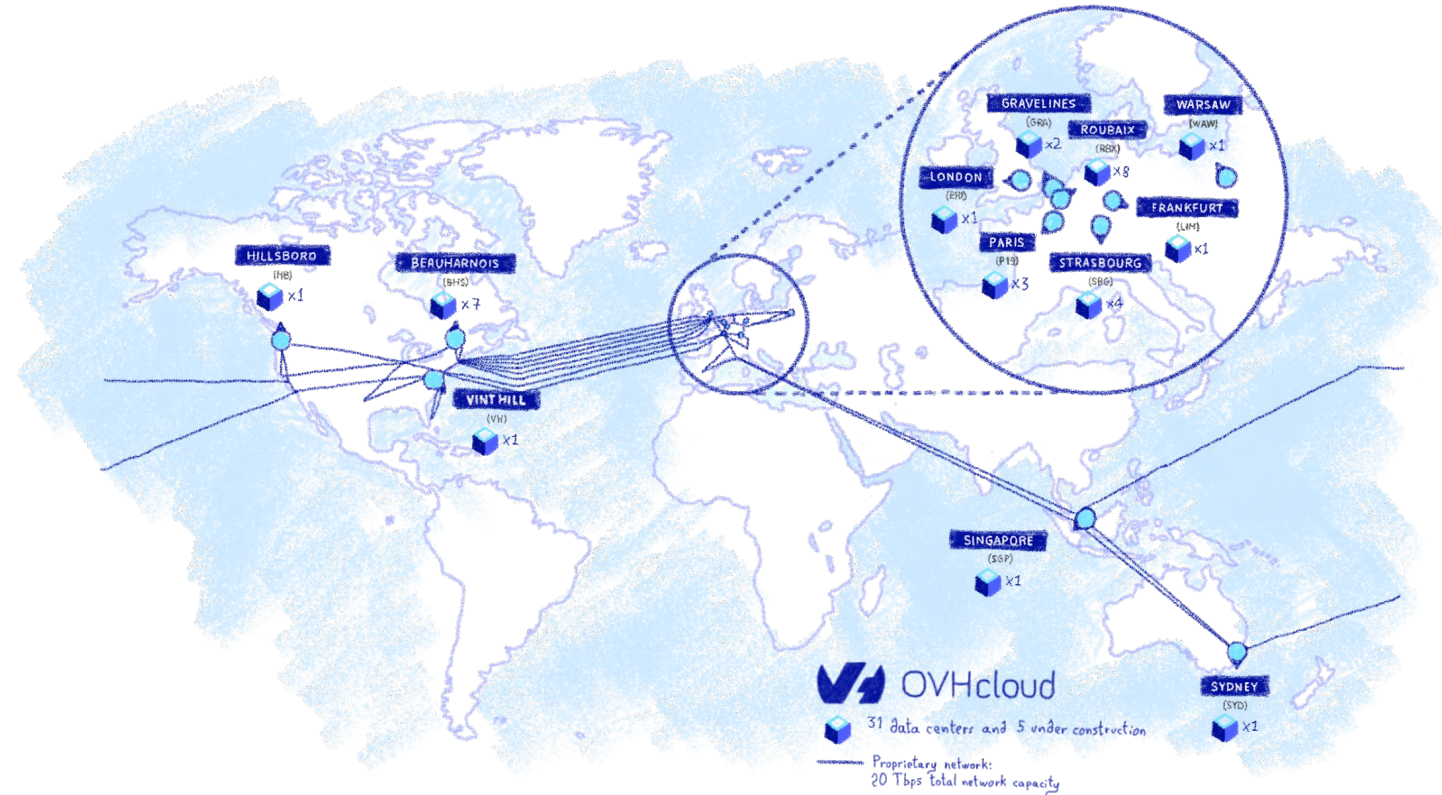
“OVH: We Host You” (On vous héberge)



Webhosting at OVHcloud



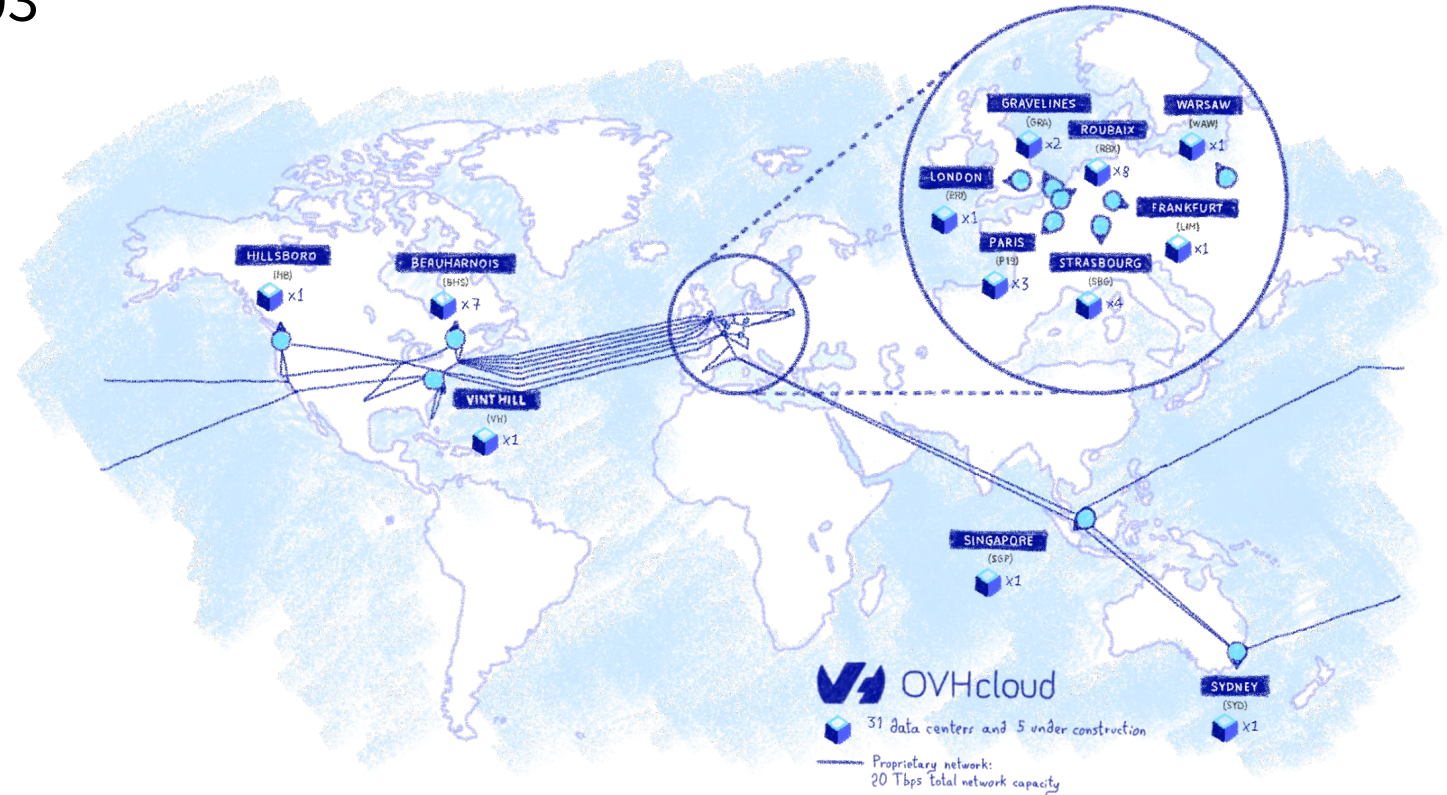
- Biggest webhoster in Europe
- 6 millions websites
- 60 Gb/s
- 6 billions HTTP requests (except CDN caches)
- 15 000 web servers




DevOpsCon

Webhosting at OVHcloud: small history

- Hosting in P19 (Paris) since 2003
- Web have changed from 1999
- New datacenter opening : Gravelines in 2016



What's the hoster's job?




Are you kidding me?



DevOpsCon

apt-get install apache2 php7 mysql-server?



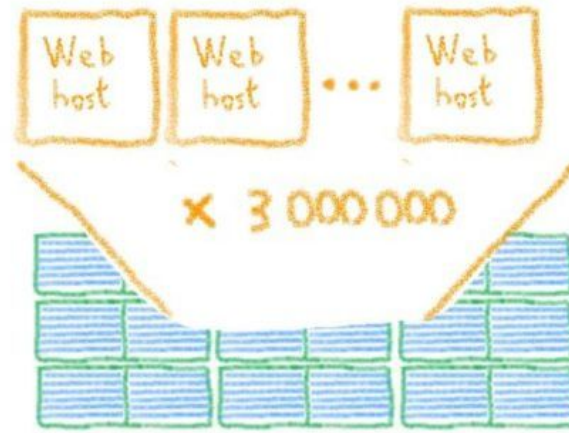


- Store data
- Run code source




DevOpsCon

Why did we want to leave Paris?

- Hardware end of life
- Too slow natural decreasing



P19 Datacenter



Gravelines Datacenter




@LostInBrittany

DevOpsCon

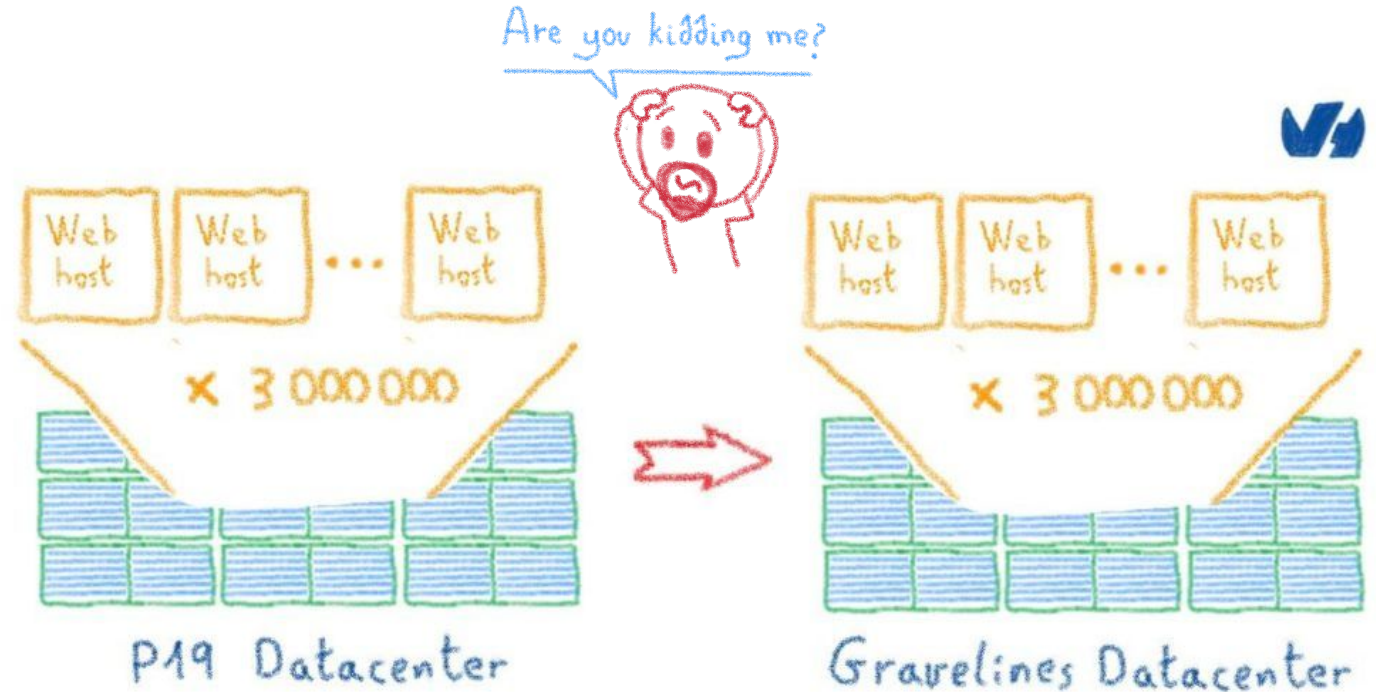
Why it was difficult?

Risk management




Probability by magnitude

- 0,1% for 1 website: 1 in 1000 chance
- 0,1% for 100 websites : 1 in 10 chance
- 0,1% for 3 millions : 3 000 times



DevOpsCon


Risk management



Probability by magnitude

- 0,1% for 1 website: 1 in 1000 chance
- 0,1% for 100 websites : 1 in 10 chance
- 0,1% for 3 millions : 3 000 times


Risk = Impact * Probability



@LostInBrittany

DevOpsCon

Split brain definition



Split-brain is a computer term, based on an analogy with the medical [Split-brain](#) syndrome. It indicates data or availability inconsistencies originating from the maintenance of two separate data sets with overlap in scope, either because of servers in a [network design](#), or a failure condition based on servers not communicating and synchronizing their data to each other.

[https://en.wikipedia.org/wiki/Split-brain_\(computing\)](https://en.wikipedia.org/wiki/Split-brain_(computing))



@vcasse





OVHcloud



@LostInBrittany

DevOpsCon

Hosting architecture. Vue for one website



@vcasse



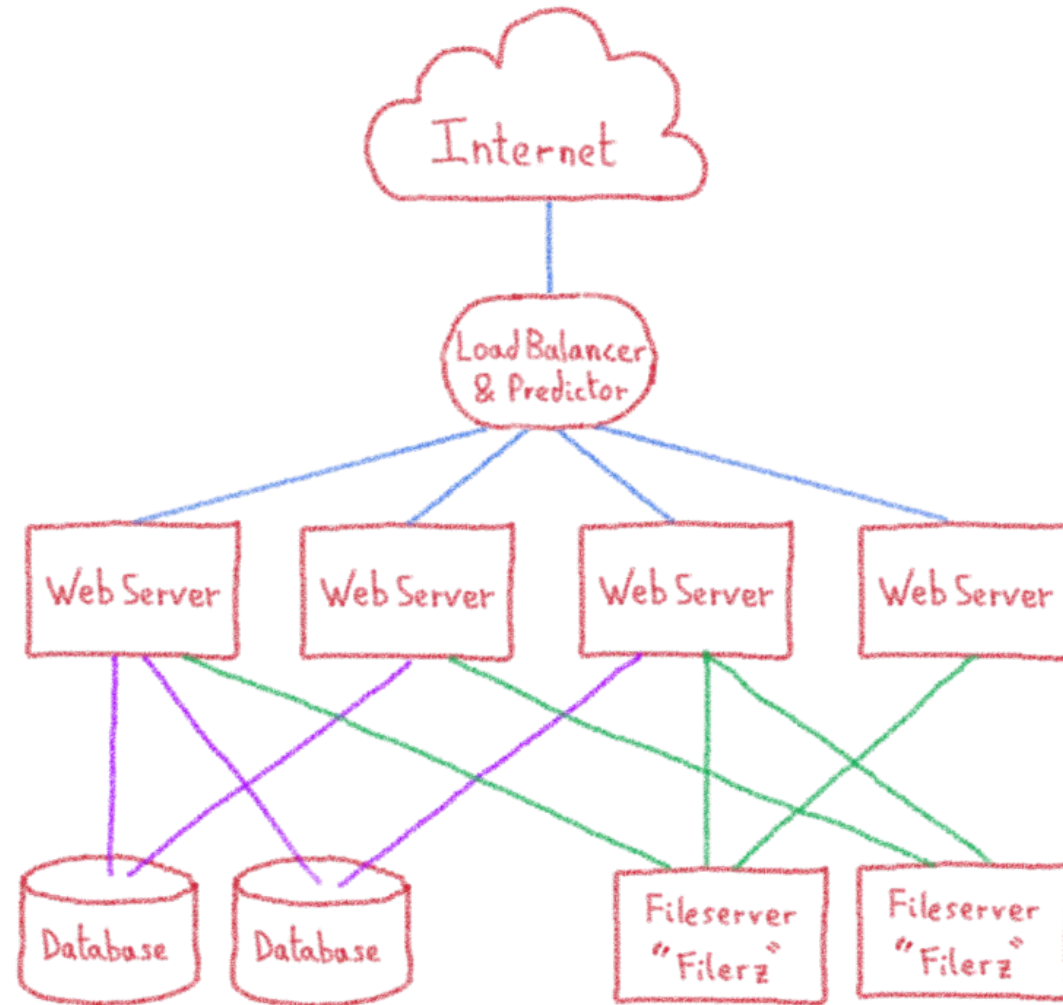

OVHcloud



@LostInBrittany

DevOpsCon

Load balancing and fault tolerance



Date



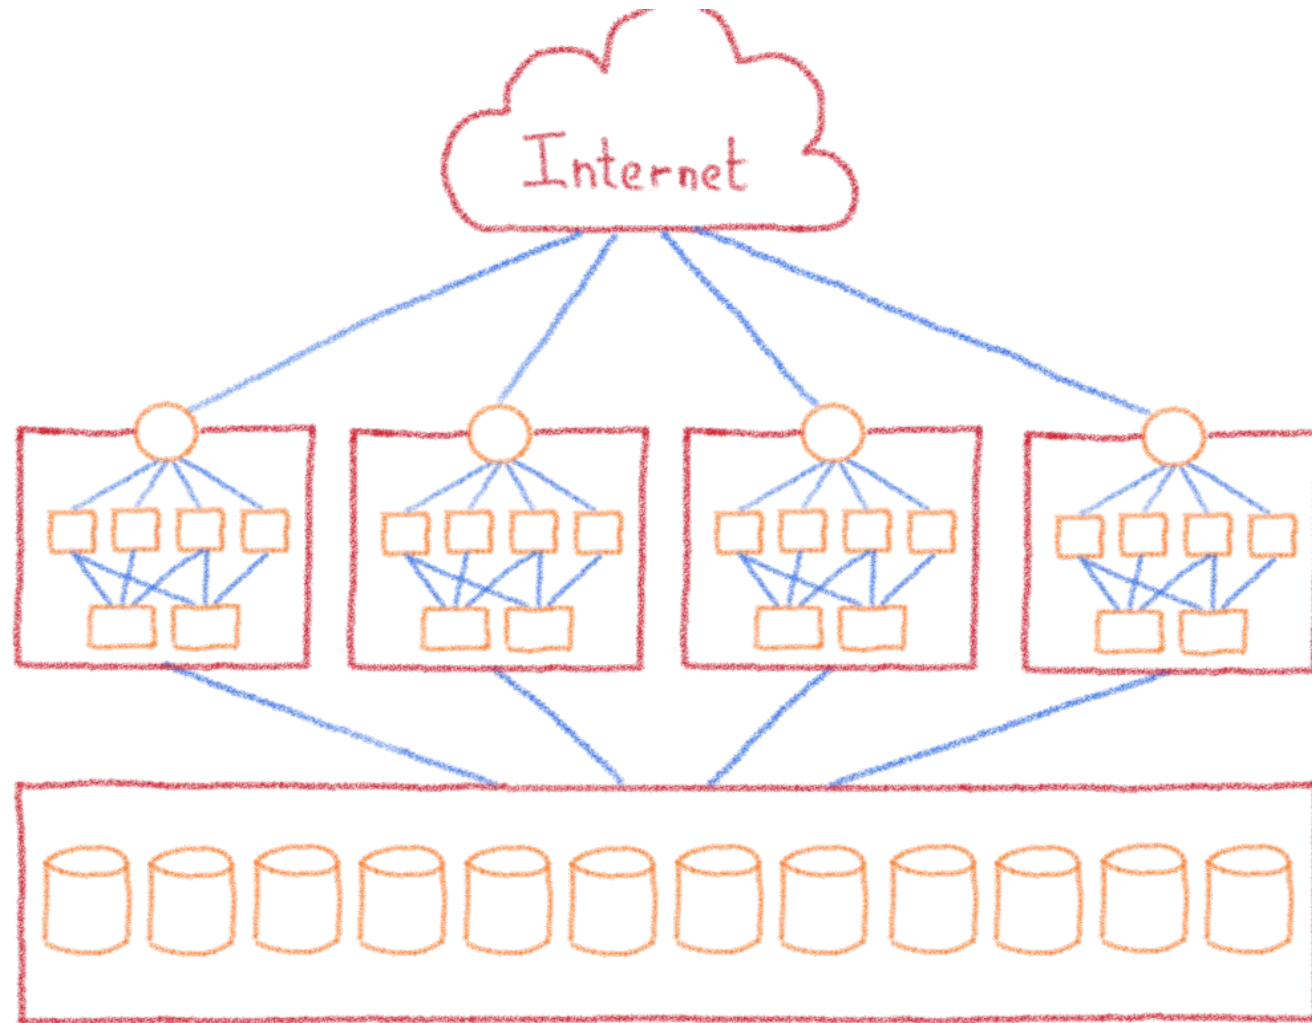

@vcasse



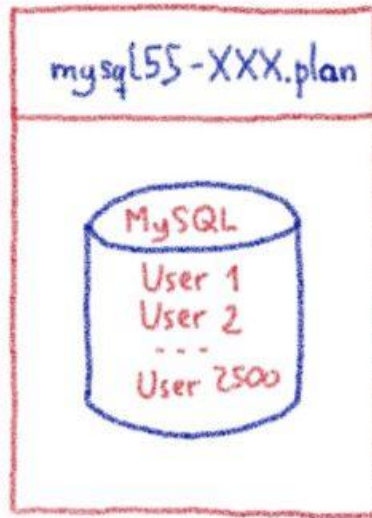

@LostInBrittany

DevOpsCon

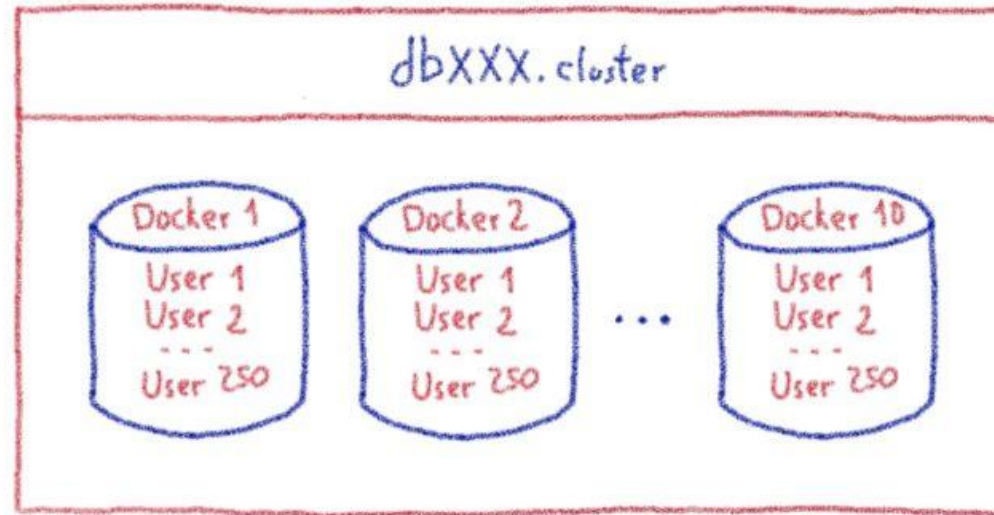
Fault domain



Difference between P19 & Gravelines



PARIS P19



GRAVELINES

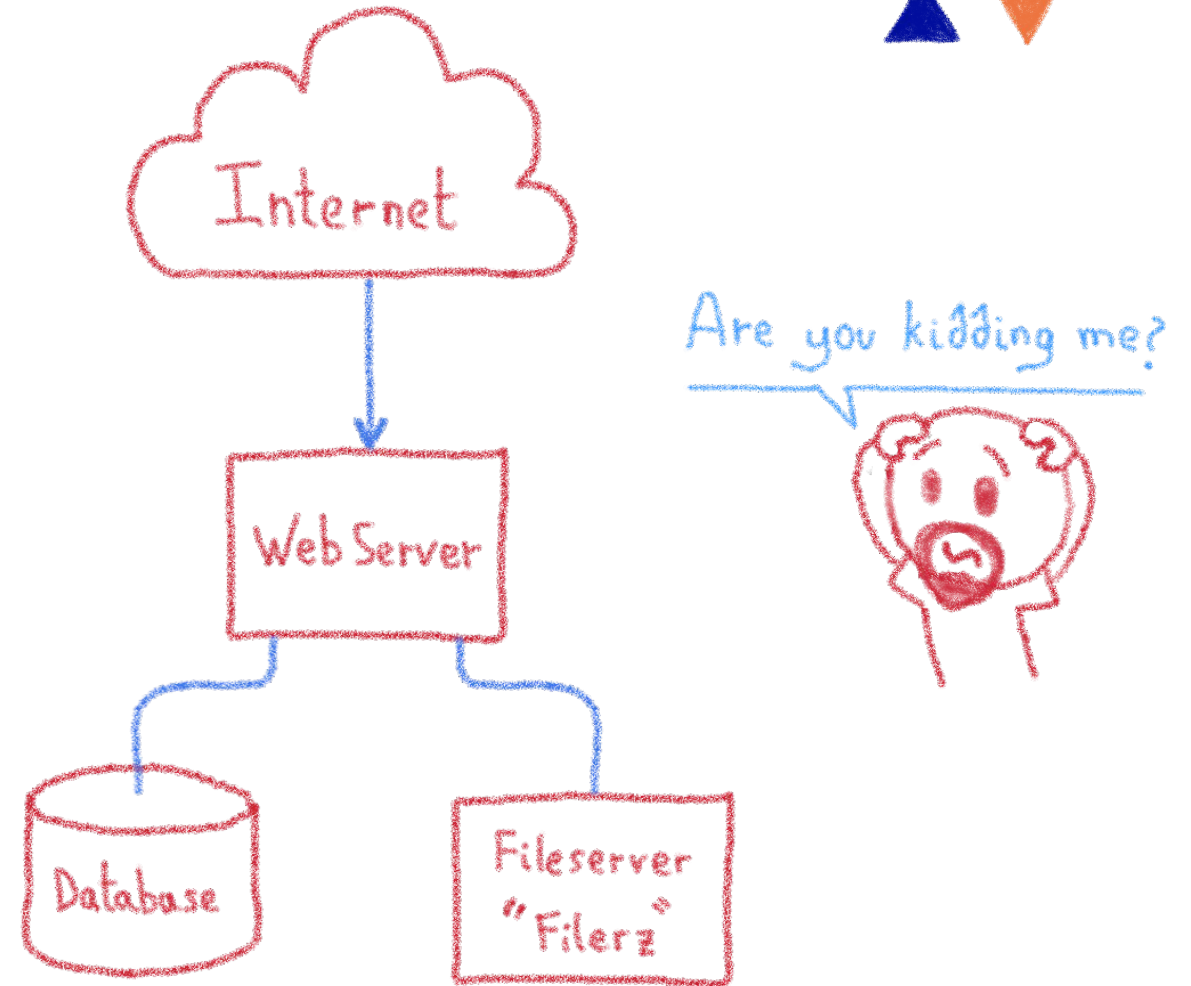


@LostInBrittany

DevOpsCon


Files constraint

- Customer dependencies:
source code / images / javascript...
- Rsync limitations
- Bloc copy implies to migrate
all customer of a filerz




Clusters constraint

- High cost infras are shared by cluster (load balancer, IP...)
- DNS zone relies on customer configurations
- IP migration implies to migrate all cluster customers



Database constraint

- Database linked to one hosting account but...
- Exhaustive knowledge = comprehensive mastery of source code
- Break zero website implies to migrate in same time all websites



Are you kidding me?



@vcasse



OVHcloud

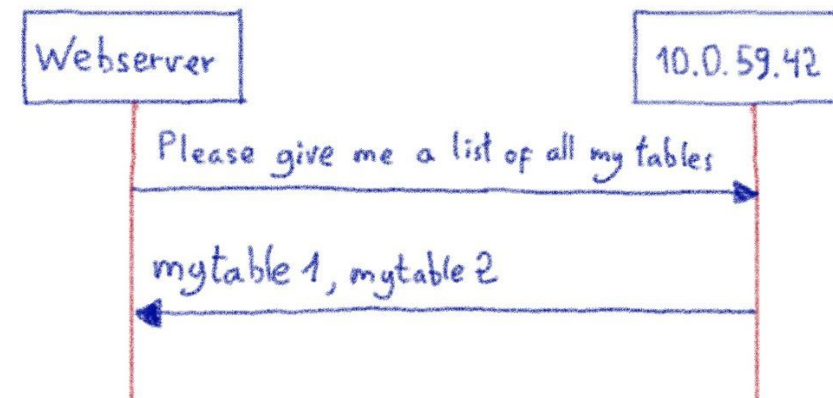
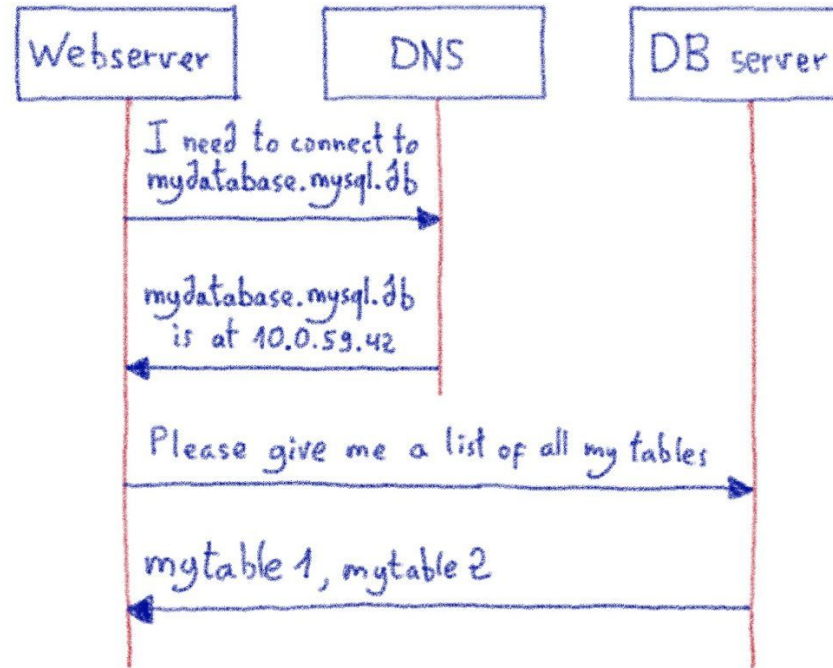



@LostInBrittany

DevOpsCon

Database constraint²

- Database naming use subdomain of mysql.db
- But “recent” feature (5 years)
- Old usages incompatible in Gravelines datacenter





@vcasse




OVHcloud




@LostInBrittany

DevOpsCon

So, how migrate?



Be punk! Break the rules



If we take all constraints:

- Either migrate the sites one by one knowing their website
- Either migrate all at the same time (TCP over Trucks ?)

Are you kidding me?



@vcasse




OVHcloud



@LostInBrittany

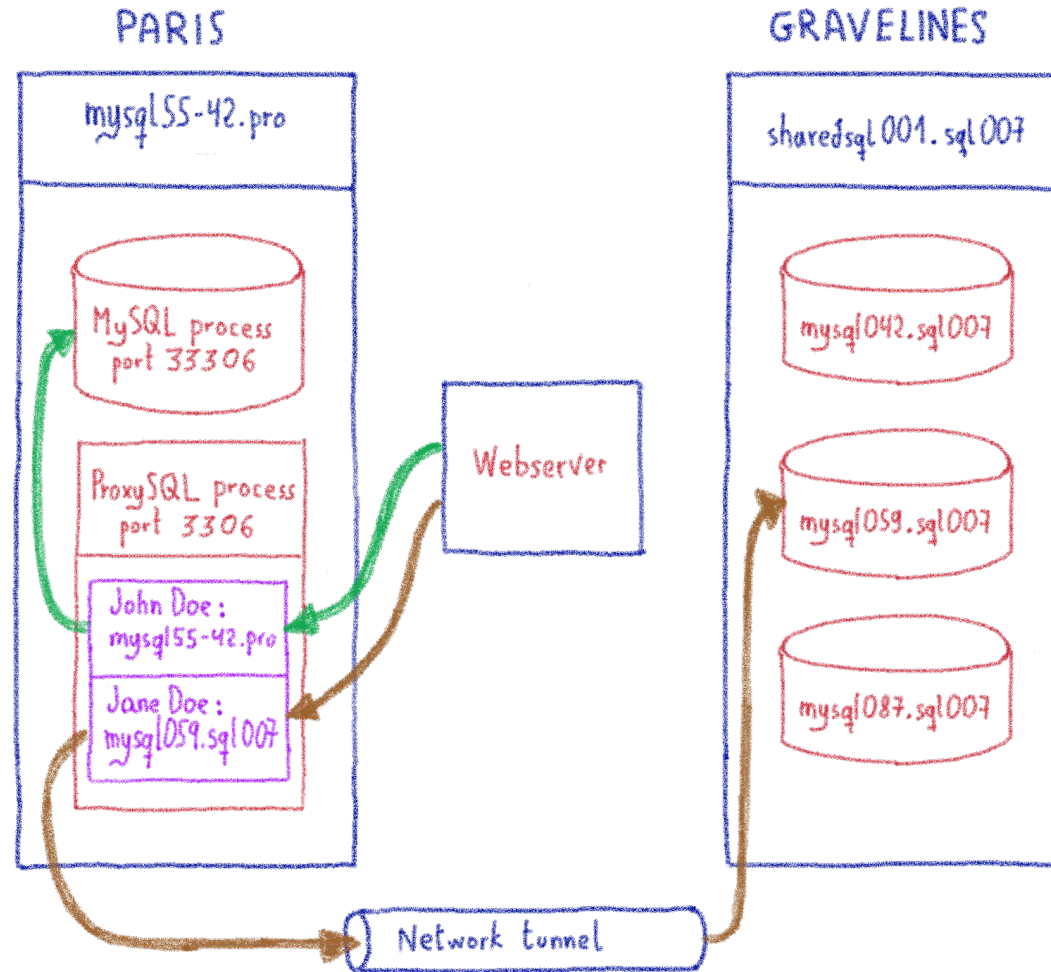

DevOpsCon

Database naming



DevOpsCon

Database naming : ProxySQL

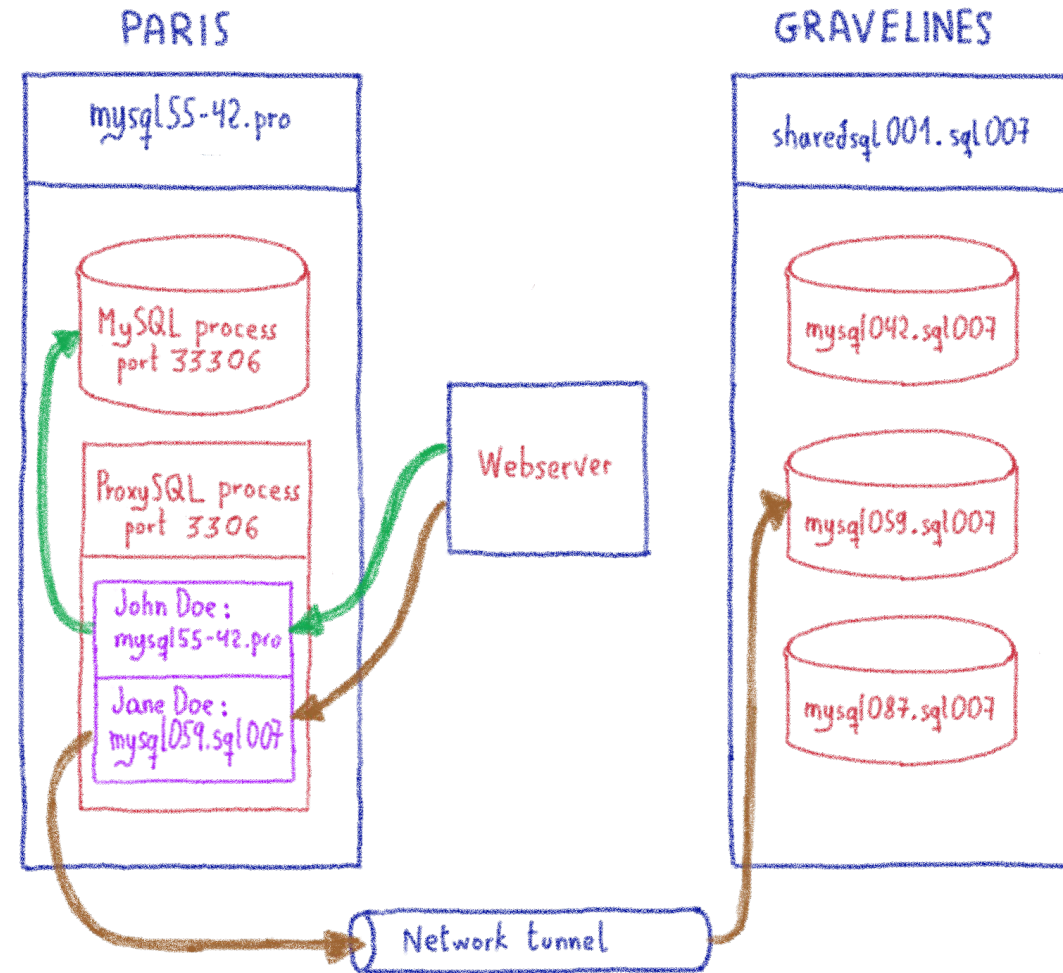


@LostInBrittany

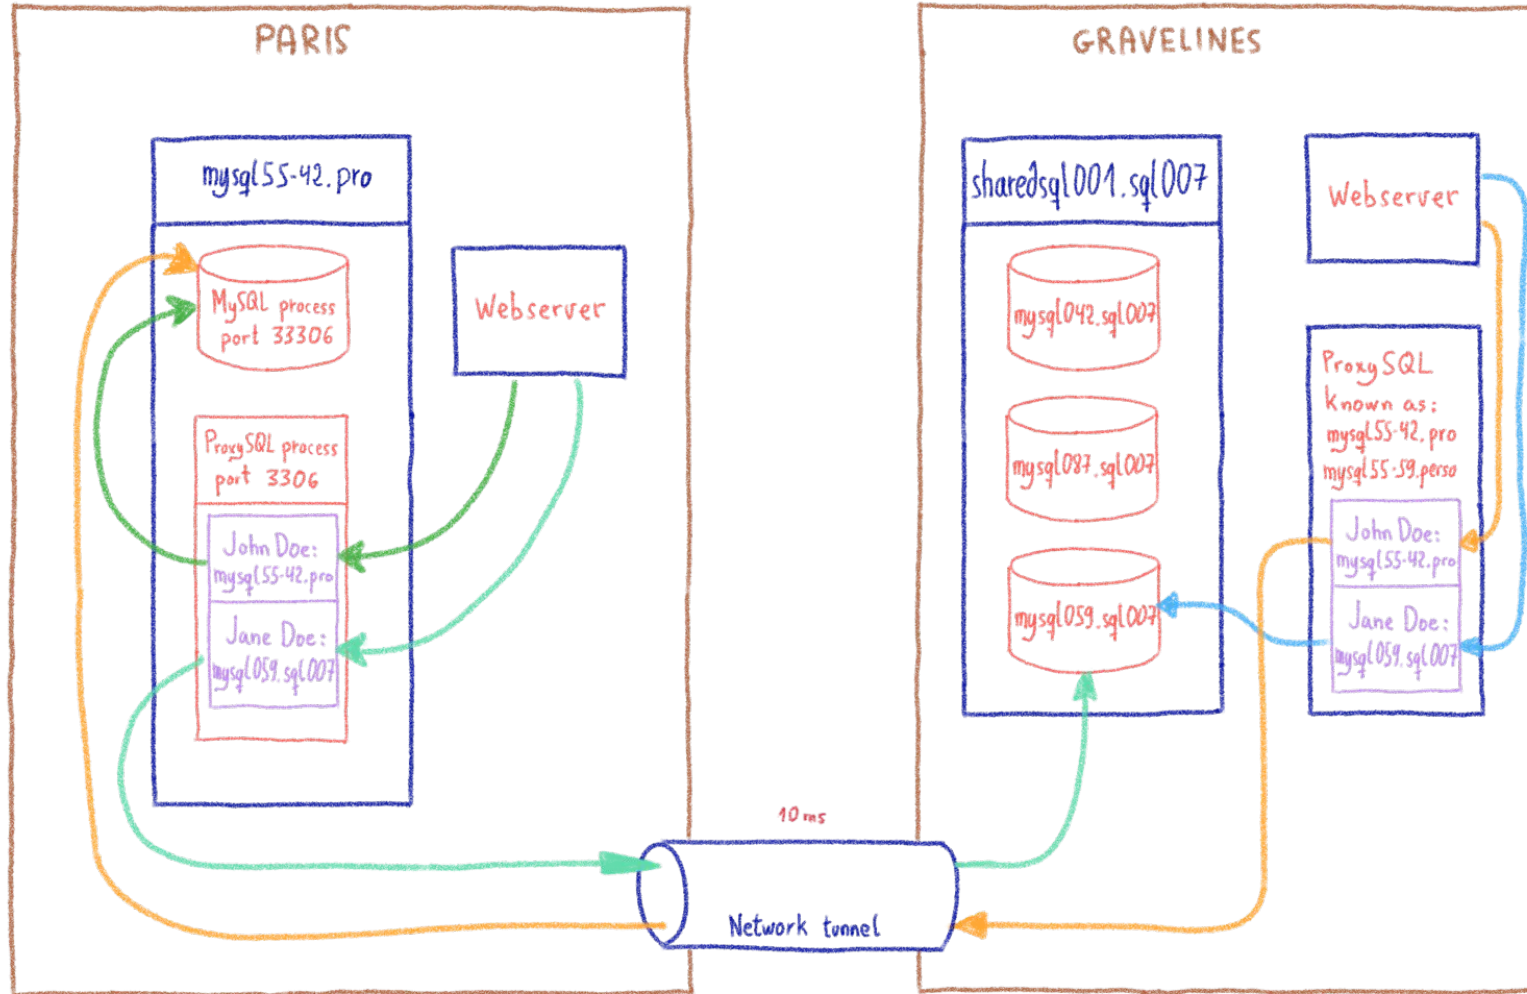

DevOpsCon

Database naming not know

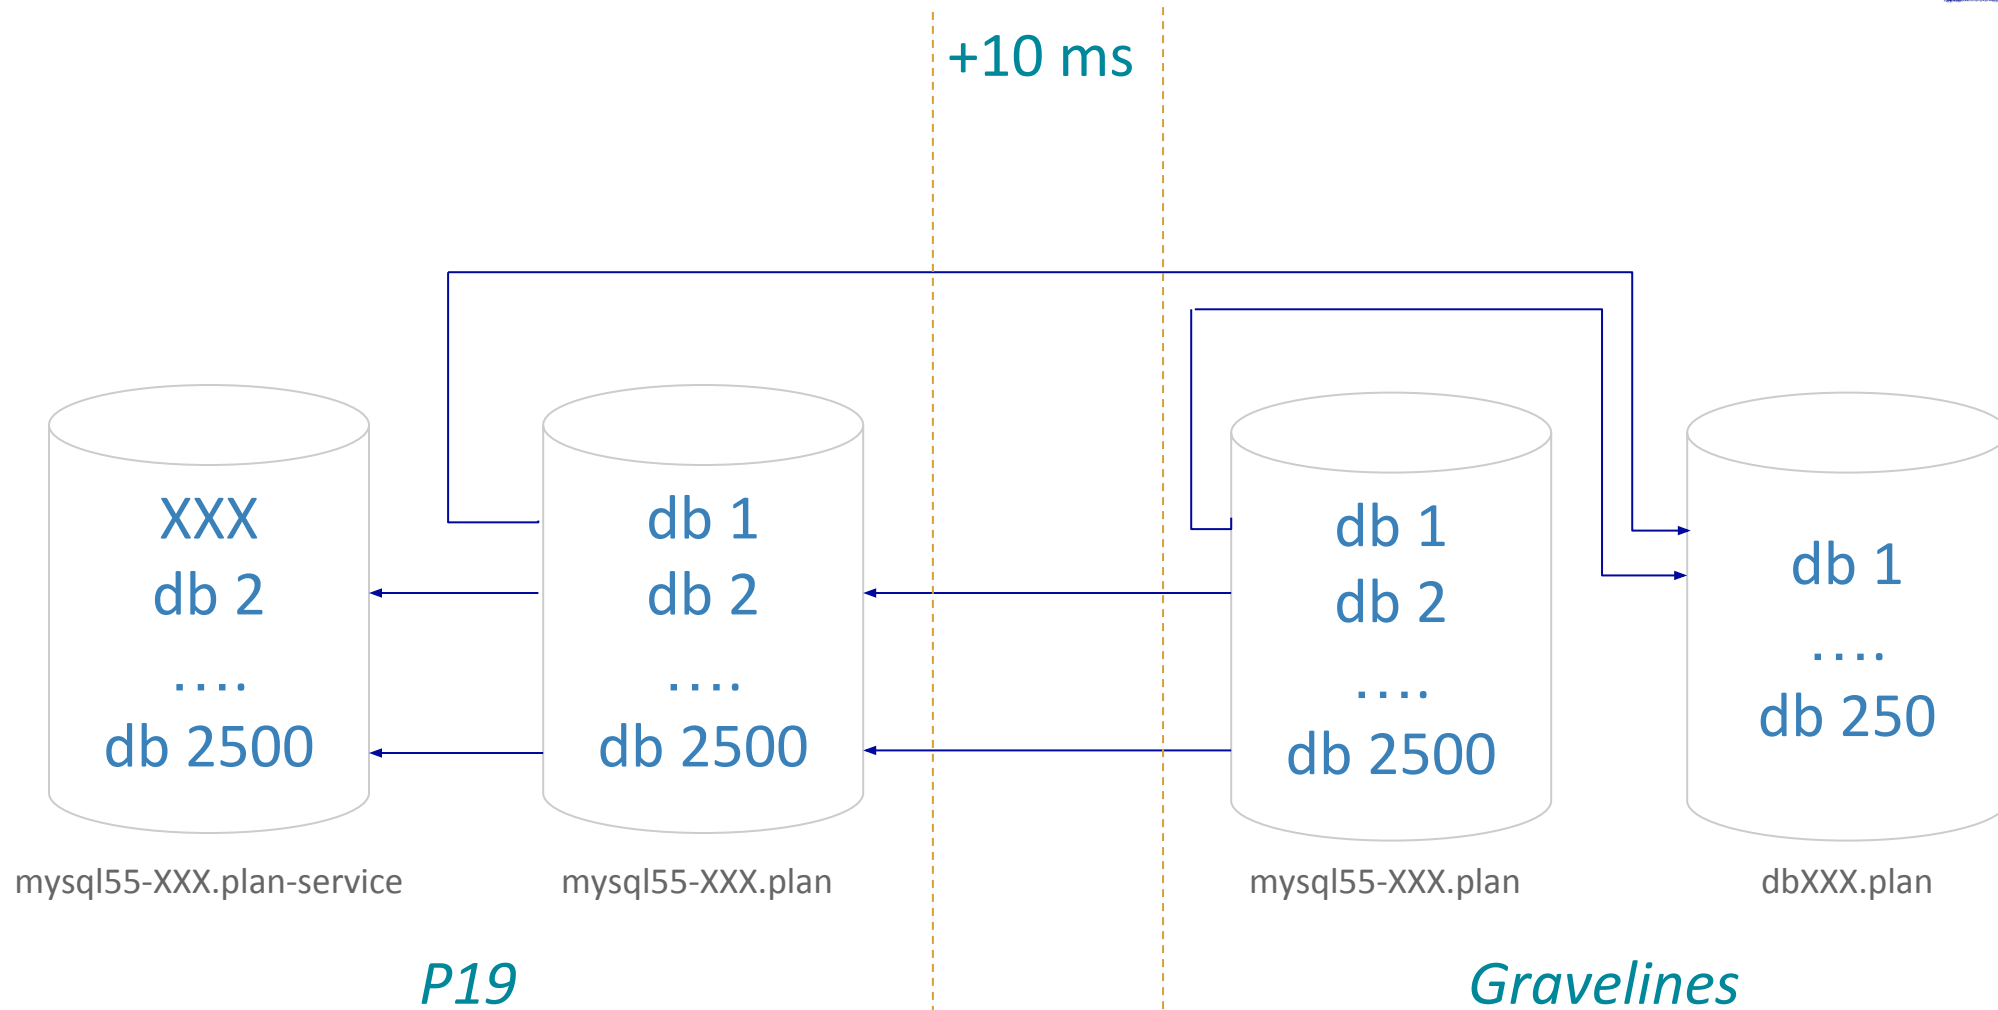

- Network tunnel between the two datacenters
- Impact : + 10ms latency for each request
- « Best effort »



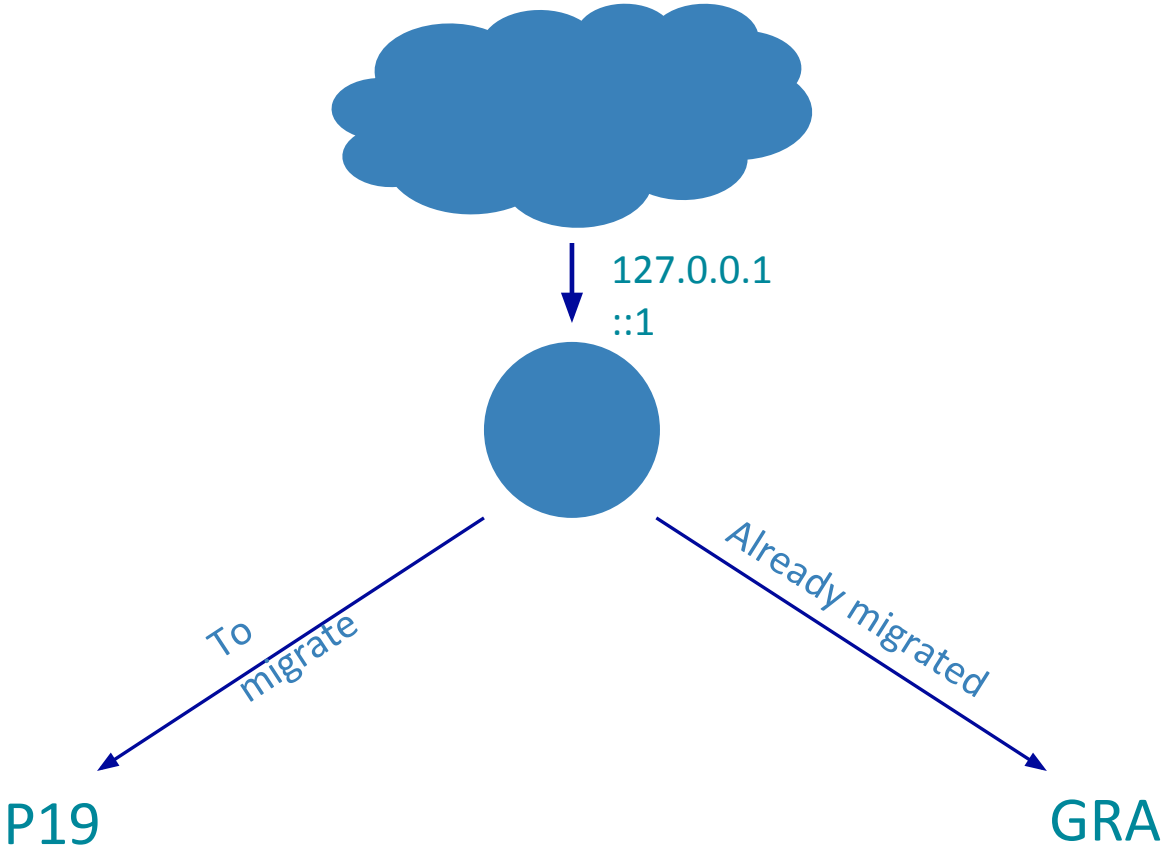

ProxySQL and latency





SQL proxy and latency



Shared IP constraint




Shared IP constraint



P19


To migrate



GRA



File constraints



Are we able to migrate filerz customer all at the same time?



@vcasse



OVHcloud




@LostInBrittany

DevOpsCon

Party time! Let's migrate!



IP switch



- Information system adaptation
- Load balancer patch
- Network tunnel
- Tools & monitoring



@vcasse




OVHcloud




@LostInBrittany

DevOpsCon

ProxySQL



- Configuration automatisation
- Risk management deployment :
1 / 10 / 100 / 1000
- SQL proxying at scale: some surprises
 - MySQL and password storage format...
 - ARP Table
 - Old database management



ProxySQL



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Migration plan



- Migration filerz by filerz.
- Database related to hosting migrated, migrate at the same time
- 1 IP switch d'IP at a time. So 1 cluster at a time
- Cluster migration order by risk level. Less risky to more risky



@vcasse



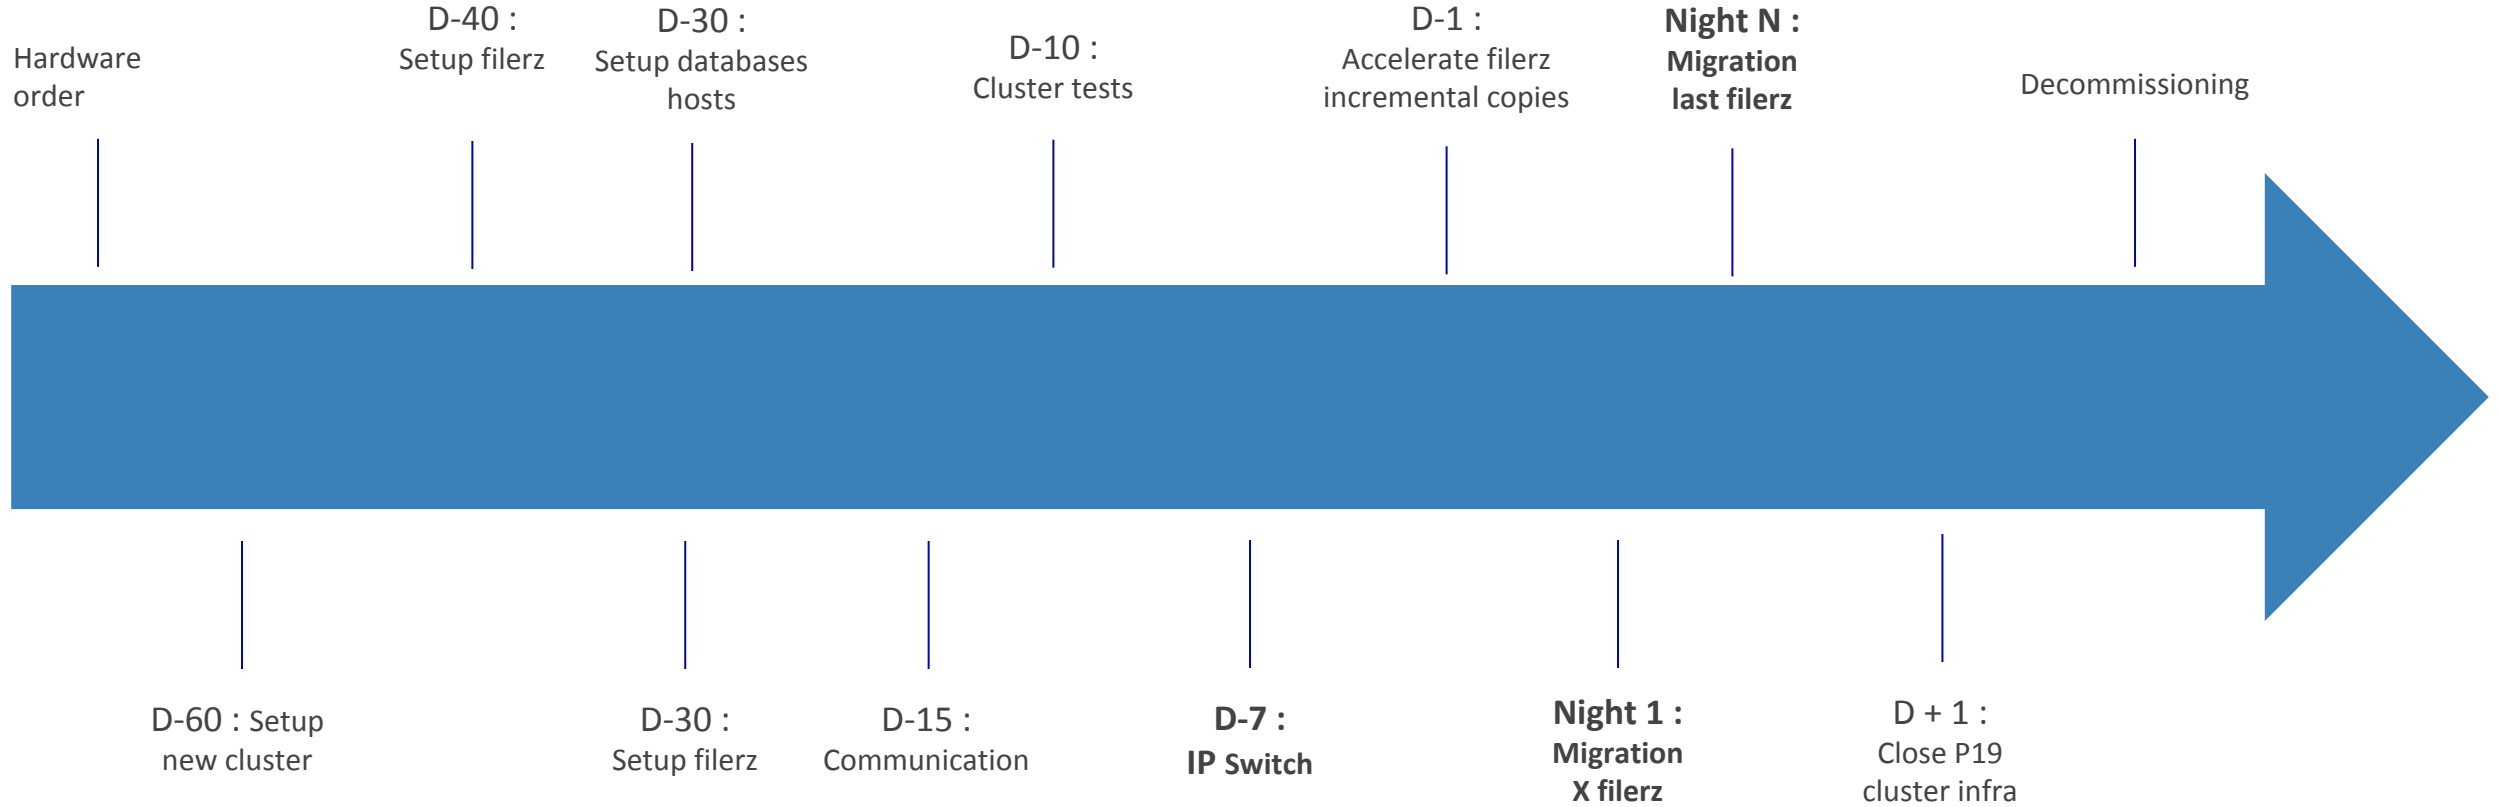

OVHcloud



@LostInBrittany


DevOpsCon

Chronological timeline



DevOpsCon

IP Switch (D-7)



1. Destination cluster a network tunnel tests
2. Send communication to support and customers
3. SSL jobs redirections
4. Setup all SSL on destination load balancer
5. For each IPv4 / v6 addresses!
 - Route IP to new load balancer
 - Tester websites at Paris and Gravelines
6. Route CDN to new infra



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Filerz migration: during the night



1. Cluster websites tests
2. Cut monitoring of the cluster
3. Launch incremental
4. Close website (maintenance mode)
5. Wait PHP timeout
6. Close file access from the filerz
7. Launch last incremental
8. When data are in Gravelines: launch database migration
9. Update configurations of migrated hosting (IS, infrastructure...)
10. Reopen hostings accounts
11. Wait end of database migration
12. Test website again and check all is ok
13. Enable cluster monitoring
14. Prevent customer about the end of operations
15. Go to bed!



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Migration: and databases?



For each databases:

1. Put database in read-only mode
2. Dump database
3. Import database on destination cluster and put the new in read-write mode
4. Redirect DNS name to the new server
5. Setup SQL proxy to new server
6. Close old database in Paris



@vcasse




OVHcloud



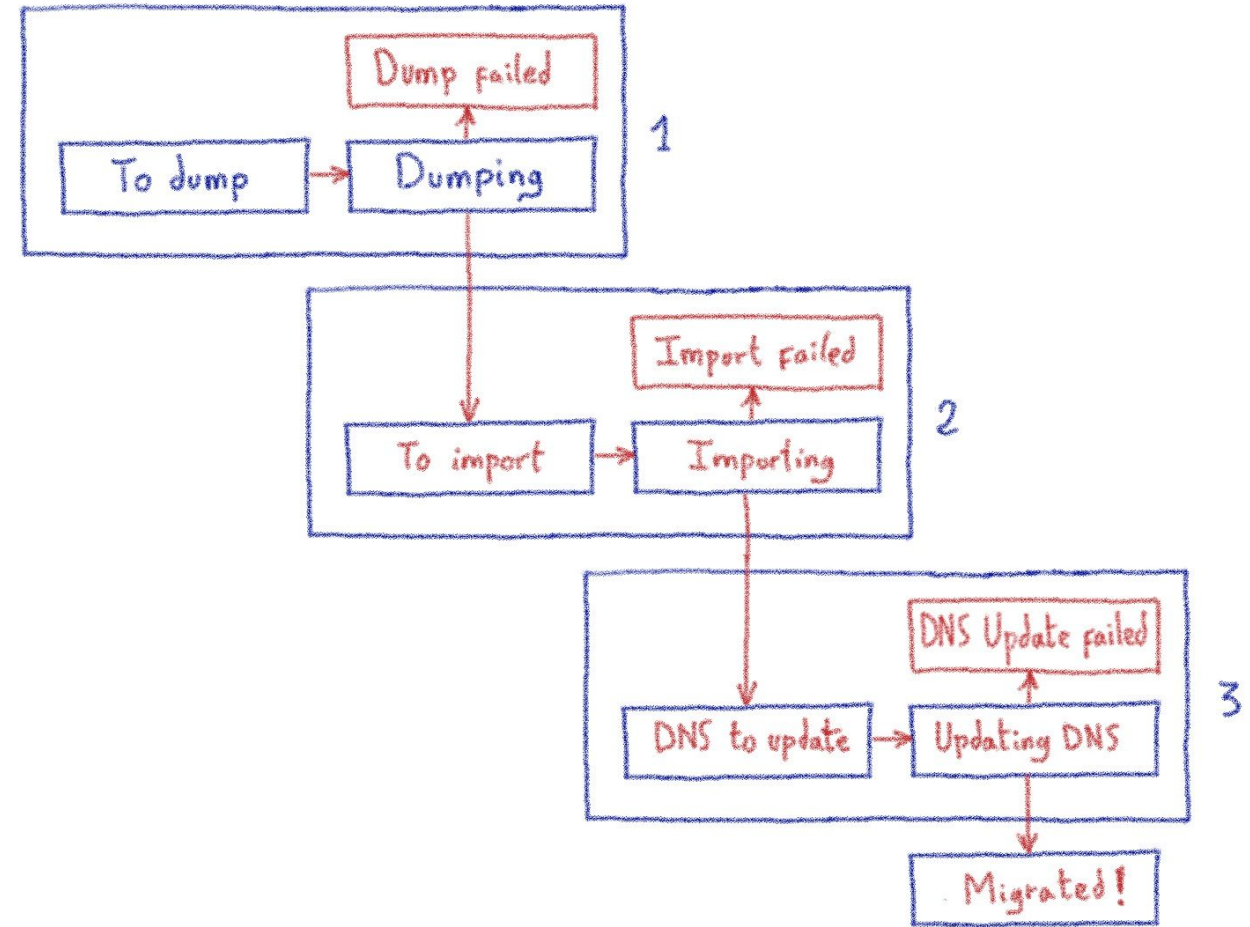
@LostInBrittany

DevOpsCon


Migration: and databases?



- Distribution of operations on all servers
- Orchestration
 - storing information inside... a database !

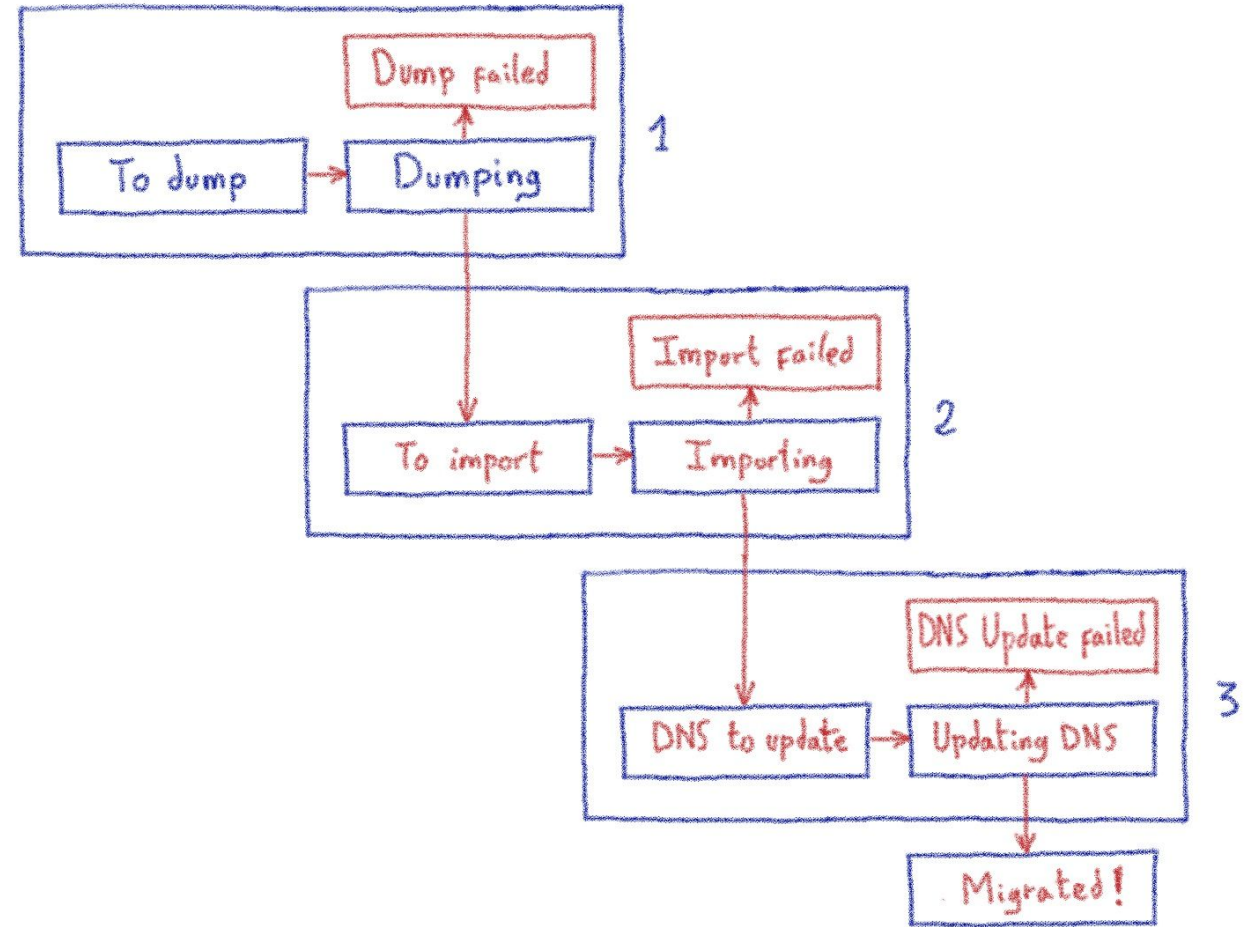


Migration: and databases?



- Distribution of operations on all servers
- Orchestration
 - storing information inside... a database !

Record
13 502 databases migrated
in 1 hour 13 minutes




@LostInBrittany

DevOpsCon

Organisation



Challenges



- Technical. But it was this presentation up to this slide
- Infrastructure work splitted in specialized teams (database, web servers, storage servers, datacenters, server factory, support, load balancers, cdn, network...)
- Legacy
- Loooong migration



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Continue improvement organisation



1. Build migration plan
2. Implement and test the plan
3. Migrate then improve migration after each week



@vcasse




OVHcloud



@LostInBrittany

DevOpsCon

Change management




<input type="checkbox"/> 33 <input type="checkbox"/>	done ▾	Jun 26, 2019 12:01 AM	Jun 26, 2019 12:07 AM	6 minutes	Lock every hosting on the filerz : chmod 000	webhosting	<pre>perl manual/migration_p19/migrate_filerz/lock_filerz.pl --cluster \$SRCCLUSTER --filerz=\$SRCFILERZ --storageid=\$FULLTODO</pre> <p>&</p> <p>Wait 5 min (PHP timeout)</p> <p>export LOCKTODO=<the outputted todo id></p>
<input type="checkbox"/> 34 <input type="checkbox"/>	done ▾	Jun 26, 2019 12:07 AM	Jun 26, 2019 12:19 AM	12 minutes	Wait end of storage lock todo	Mutu	<pre>perl manual/migration_p19/migrate_filerz/storage/check_todo.pl --id "\$LOCKTODO"</pre>
<input type="checkbox"/> 35 <input type="checkbox"/>	done ▾	Jun 26, 2019 12:19 AM	Jun 26, 2019 12:20 AM	a minute	Request final zMotion to storageTeam	Mutu	<pre>perl manual/migration_p19/migrate_filerz/storage/post_check.pl --cluster="\$SRCCLUSTER" --src="\$SRCFILERZ" --dst="\$DSTFILERZ" --lasttodo="\$LOCKTODO"</pre> <p>export LASTTODO="<the outputted todo id>"</p>
<input type="checkbox"/> 36 <input type="checkbox"/>	done ▾	Jun 26, 2019 12:20 AM	Jun 26, 2019 12:24 AM	4 minutes	Wait that Storage has finished the zMotion	Mutu	<pre>perl manual/migration_p19/migrate_filerz/storage/check_todo.pl --id="\$LASTTODO"</pre>
<input type="checkbox"/> 37 <input type="checkbox"/>	done ▾	Jun 26, 2019 12:52 AM	Jun 26, 2019 12:53 AM	a few seconds	Give Go to migration to database team	Mutu	Speak with humans



DevOpsCon

--verbose?



- Why we decided to migrate three million websites?
<https://www.ovh.com/blog/web-hosting-why-we-decided-to-migrate-three-million-websites/>
- How to host 3 million websites?
<https://www.ovh.com/blog/web-hosting-how-to-host-3-million-websites/>
- How to migrate 3 Million web sites?
<https://www.ovh.com/blog/web-hosting-how-to-migrate-3-million-web-sites/>
- How do our databases work?
<https://www.ovh.com/blog/web-hosting-how-do-our-databases-work/>
- How to win at the massive database migration game
<https://www.ovh.com/blog/how-to-win-at-the-massive-database-migration-game/>
- migrate-datacentre –quiet: How do we seamlessly migrate a datacentre?
<https://www.ovh.com/blog/migrate-datacentre-quiet-how-do-we-seamlessly-migrate-a-datacentre/>
- A day in the life of a ProxySQL at OVHcloud
<https://www.ovh.com/blog/a-day-in-the-life-of-a-proxysql-at-ovhcloud/>
- Another day in ProxySQL life: sharing is caring
<https://www.ovh.com/blog/another-day-in-proxysql-life-sharing-is-caring/>

More soon on <https://ovh.com/blog>



@vcasse



OVHcloud



@LostInBrittany

DevOpsCon