



CASE STUDIES BOOK



DISCOVER ORIGINAL INFRASTRUCTURES DEPLOYED AT OVH.COM

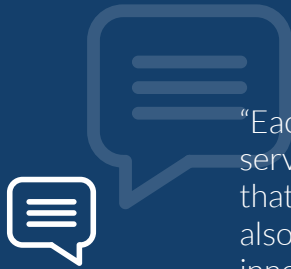


 OVH.com



Online media, applications and business software, collaborative tools, smart music playlists at points of sale, total outsourcing of an urban community's IT...





“Each year, OVH is able to offer several hundred new services. After listening to you, we have come to realize that launching new services is not enough. We should also assist and guide you in the adoption of these innovations. Often you just need to see some practical use cases. That is the goal of the “Case Studies Book”, to provide you with such examples. And of course, if you need advice, our customer advocates and our Professional Services team are always available. Enjoy your reading!”

Octave Klaba, Founder OVH

CASE STUDIES BOOK

CASE STUDIES BOOK

This publication has been edited by the OVH Group, 2 rue Kellermann 59100 Roubaix (FRANCE) - RCS Lille Métropole 424 761 419 00045.

Written by: Hugo Bonnaffé
in collaboration with:

Rémy Vandepoel, Vincent Cassé, Jean-Daniel Bonnetot, Félicien Lainé, Jonathan Guinez, Sylvain Wallez et Alexandre Morel.

Translation: Michael Kapus

Layout: Marie Delattre

Graphic Design: Carl Lambert, Lucille Folens and Élodie Lenin

Photographers: Élycia Husse, Frédéric Anne, Alban Gernigon, Stéphane Bureau du Colombier, iStockphoto

Printer: Nord'imprim

Special thanks is given to each customer that made this project possible by revealing the details of their infrastructures and their willingness to share their expertise with the OVH Community. Congratulations to everyone who took to the stage during the OVH World Tour to present their projects and explain their technical choices.

Summary

#01

Photography and Video p. 6 to 11

Giroptic Hosts the Content of Its Users on OVH.com Public Cloud Object Storage

#02

Expertise Accounting p. 14 to 21

Implementation of an HA Infrastructure and DRP for Hosting a Critical Application (ERP)

#03

Industry p. 24 to 27

High Availability and Collaborative Applications on Two Dedicated Servers

#04

Online Media p. 30 to 38

NextRadioTV:
A Behind the Scenes Look Into the Infrastructure Hosting the Digital Activities of the Media Group

#05

Local Community p. 42 to 49

VDI, IS and Video Protection: The Urban Community of Bourget Airport Uses Dedicated Connect and Dedicated Cloud to Outsource Its IT

#06

Sound Design p. 52 to 56

RadioShop:
A Highly Available Infrastructure, Resistant to Increases in Traffic that Provides Background Music for 7,000 Points of Sale Worldwide

#07

Social Network p. 60 to 65

Simplify Implementation of a Backup Strategy with Veeam Managed Backup

#08

Online Media p. 68 to 75

A Scalable Infrastructure Capable of Handling Large Peaks in Load

#09

Over the Top Video Service p. 76 to 81

Vodkaster: A Flexible Infrastructure for a Remote DVD Player Service Application

#10

Collaborative Work p. 84 to 93

Wimi: An Infrastructure with High Availability Combining Dedicated Cloud and Dedicated Servers Through vRack

#11

Connected Farm p. 96 to 101

Medria: OVH.com Servers to Monitor Connected Livestock

#12

Automobile p. 104 to 108

Keyzee: A Highly Available, Secure and Load Resistant Infrastructure for Hosting a Virtual Key Application

A black Giroptic camera device is shown in the foreground, slightly out of focus. It has two large lenses on top and a small display on the front that says 'Photos' in white dots. Below the display is a row of five small, colored LEDs (red, orange, green, blue, purple). The background is a blurred sunset or sunrise over water.

#01

Giroptic Hosts the Content of Its Users on OVH.com Public Cloud Object Storage

With more than 1.4 million dollars raised on Kickstarter for its camera project that shoots high definition video and photos at 360°, the start-up from France, Giroptic, holds the French record for crowdfunding. To host the photos of its camera users, Giroptic has built a scalable platform based on VPS and Public Cloud Object Storage solutions. Learn more: 360.tv



Specifications

- Ease of platform deployment and administration
(system administration is managed by the company which is made up of mostly developers and engineers specializing in embedded electronics).
- Unknown number of future users
- The volume of storage employed by users is difficult to estimate
(users have the possibility to store photos on their own computer, NAS, or on any of the many cloud storage services).

Photo Storage | International Users

Web Hosting | VPS

Public Cloud Object Storage | IP Load Balancing



Key Points

July 2014: start of crowdfunding on Kickstarter

March 2015: shipped the first “360cams” to 80 developers that participated in the crowdfunding by purchasing the pre-sales “developer pack”. The developers kit included access to unlimited storage on the platform.

Sept. 2015: shipped the 3,500 preordered “360cams” and launched the storage platform.

Average size of a 360° photo: **3 MB**

Average size of a 360° video: **5 to 6 GB/hour**

Pre-registered buyers from **all over the world**
(mainly Europe, Asia, North America)



Performance Hosting Plan for Promotion of Its Project

To launch its Kickstarter crowdfunding campaign, Giroptic opted for two “Performance” hosting plans, one for its web presence (360.tv) and the other for a testing and pre-production environment. Powered by WordPress, the site supported all traffic, even with significant increases in traffic caused by frequent mentions of the project in the media. The possibility to “boost”



hosting resources by upgrading the hosting offer (four levels to choose from with different amount of guaranteed resources depending on level option) for a period of a month, allowing the site to handle increased traffic. This was especially beneficial during the launch phase of the project which received considerable buzz on social media. Thanks to the GeoCache accelerator, which is included with the offer (a CDN system which takes advantage of 17 points of presence on the OVH worldwide network, locally storing in cache the static content of the site), international visitors (nearly 80% of traffic) enjoy optimal access times.

Cloud VPS the Infrastructure's Core

To address the delivery of the first orders, Giroptic moved its web application onto two Cloud range virtual private servers, one in France and the other in Canada. On one hand, it was necessary to obtain optimal performance in anticipation of the first users of the 360° photo storage platform. While on the other hand, it was essential to have root access to the machine, in order to make specific configurations for such tasks as image processing, guaranteeing rapid execution of scripts. It was also a wise idea to place the application in close proximity to North American users, thus providing them

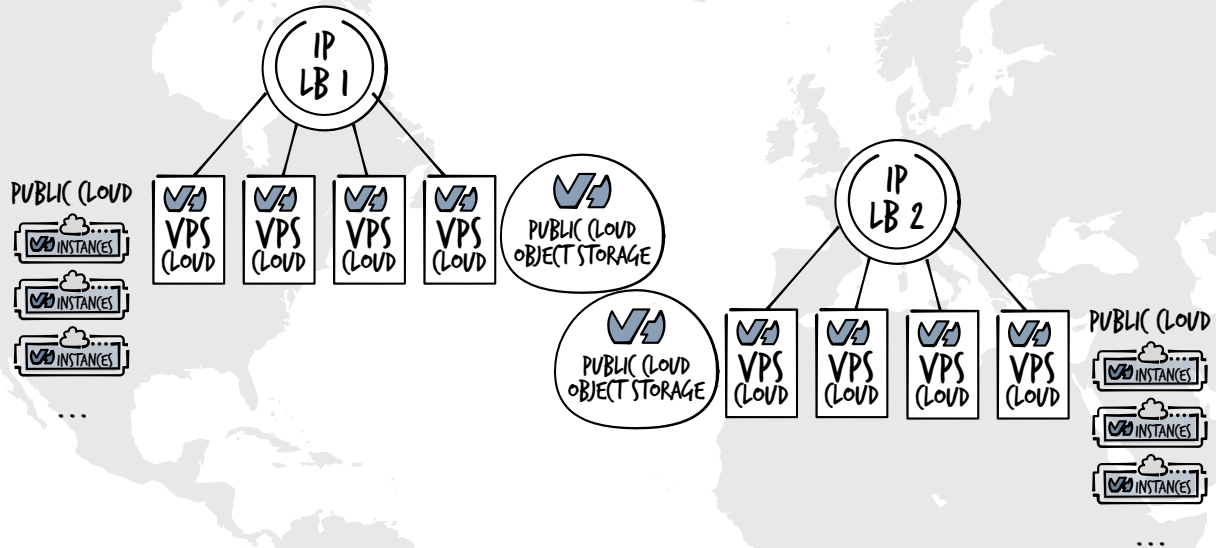
with the most fluid user experience possible. Giroptic opted for a VPS solution because it did not want to be preoccupied with the management of machine hardware. The VPS Cloud range includes a comprehensive SLA (99.99% with reboot times of less than 5 minutes in the event of hardware failure), which is perfectly compatible with the high availability requirements of the platform. In addition, the Cloud range opens access to advanced backup options (snapshot option or automatic daily backup).

Public Cloud Object Storage for Easily Scalable Storage Space

With their camera, users have the added benefit of storage space for backing up and sharing their 360° photos. As 360° videos files are quite large, Giroptic refers users to platforms such as YouTube, which announced compatibility with these types of files in the first quarter of 2015. The main issue for Giroptic, in addition to storing files, is that it doesn't know in advance how many users it service will have, let alone the volume of data uploaded by each. In addition, users have the possibility to store their files locally, on their own NAS, with a third party cloud provider, etc. Under the condition of not being able to define disk space requirements

Solution deployed

MEDIUM-TERM DEPLOYMENT PROJECT



in advance, consequently it is not possible to provision servers and disks accordingly. Public Cloud Object Storage brings with it simplicity, security and ability to better control costs. Employing Object Storage simplifies the administrator's life because data is uploaded to object storage via the OVH API and without the preoccupation of managing hardware or system files. The platform is more secure as all data pushed to object storage is accessible via https and with the possibility to configure access rules (ACL). Data is automatically replicated across three different racks: there is no need to worry about disk failures, reimporting data, reconstruction of RAID, etc. In addition, each object storage has a unique login. In theory, storage space is unlimited and billed monthly per GB with outbound bandwidth charges being billed simply for the amount consumed. Clearly, there is no need to risk budget on storage servers without knowing if or when there will be a return on their investment. Also, there is the more bearing effect with adding new disks, they take time to fill and to become cost effective. There are certain times throughout the year (after vacation, holiday season...) when users upload and consult images at a much higher rate than on average. The advantage of Public Object Storage is that it manages all traffic without

having to pass through Giroptic's servers by routing it directly to the OVH storage platform.

Projects: Deploy Additional VPS, Reinforce Infrastructure with Additional Public Cloud Instances and Object Storage over CDN

In the short term, it is possible to estimate the number of platform users. The number corresponds, more or less to the 3,500 people which registered for pre-sales of the camera via Kickstarter. In the mid-term and for the long term, it is impossible to imagine what sort of commercial success the product will encounter. This is why Giroptic has, from the point of design of its infrastructure, put thought into scalability. How do you go from a few thousand users to hundreds of thousands without having to rebuild the platform from scratch?

These concerns have of course guided the development of the platform. On the side of infrastructure, the response envisioned is the deployment additional Cloud VPS instances on each side of the Atlantic, in order to distribute the load across multiple servers simply using IP Load Balancing (horizontal scalability). In addition, there is the possibility to hot upgrade the power of the Cloud VPS (vertical scalability).



Benjamin CATELET,
Product Manager
at Giroptic

Finally, to deal with exceptional peaks in load, which are relatively limited in time, Giroptic's plans to exploit Public Cloud instances, billed by the hour, making it possible to add and remove machines from the IP Load Balancing back-ends using the OVH API. Finally, a mid-term project is to link Public Cloud Object Storage with the CDN system in order to reduce load times for remote users of the two anchor points of the infrastructure: Roubaix in Europe and Beauharnois (near Montreal) in Canada ■



Discover Giroptic's
Testimonial Video:
ovh.to/qZEoom



Infrastructure Components

2 Web Hosting "Performance" plans

1 IP Load Balancing

2 VPS

Cloud Range

Public Cloud

- Object Storage
- Instances



17 Datacenters Today...

Today, the 220,000 servers hosted, throughout the world, by OVH are housed amongst several locations: in Europe, Roubaix, Paris, Strasbourg and Graveline; in North America, Beauharnois (near Montreal). OVH's plan for tomorrow is to install datacenters in the United States and Asia.

#02

Implementation of an HA Infrastructure and DRP for Hosting a Critical Application (ERP)

One of the customers of a computer engineering service company includes a firm of 70 expert accountants. To host the firm's ERP, the initial choice was a Cloud VPS solution, because its SLA provided guaranteed maximum availability of the application. Facing an increase in the number of ERP users, the Cloud VPS's capabilities revealed to be insufficient. Migration to a dedicated server was required, but much thought had to be put into how to deal with potential equipment failure. Ultimately, in the end, a high-availability infrastructure was deployed, based on two identical servers, in addition to a backup server that was added to secure data as part of a disaster recovery plan (DRP).



Specifications

- ERP accessible from the Internet secured by an SSL certificate
- High availability of the enterprise's mission critical application
- Data security: Data is sensitive (including financial information) and there can be no risk of data leak or loss

Critical Application | Multi-Datacenter Infrastructure | High Availability | DRP
 IP Load Balancing | Dedicated Servers | Backup Server | VPS

This case was inspired by the infrastructure designed by Barne company for its client, Sextant Expertise and was awarded the OVH 2014 Infrastructure Trophy under the category "Enterprise Infrastructure (outsourcing)".



Key Points

70 daily users

2 identical dedicated servers, one hosted in Strasbourg and the other in Roubaix

-
- Backup on a So you Start server
 - VPS Cloud for HA monitoring



When we have to host a mission critical application, the first reaction is to rely on two servers: one being an active server (master) and the other a passive server (slave), capable of taking over in the event that the master should fail. This solution is perfectly suited as part of a disaster recovery plan (DRP) and built around two metrics, RTO (the maximum allowable interruption) and RPO (maximum period during which one can lose recording data). However, in the case of an infrastructure that is modest in scope, with a slightly higher budget, it is possible - and often desirable - to build an infrastructure which provides high availability that includes two active servers. The advantages of such a system are twofold: the second server can absorb part of the workload and in case of failure of one of the two machines, the application remains functional without any interruption to service. The VPS that was previously used for hosting the ERP can be repurposed to become a machine for monitoring.

Load Balancing

A Load Balancing (LB) IP is positioned upstream of the two servers to direct ERP user's HTTPS requests. The LB IP uses a round-robin algorithm and it is possible to add weight values to each back-end.

Here, more important weight is given to the server which contains the master database. As you will discover, if the two servers are active at the web and business layers, data is only written to one of the two databases (master) before being replicated to the other (slave). As a result, the server containing the local master database is more reactive than its counterpart, which experiences 10 ms of

latency between the SBG and RBX datacenters (private network that interconnects the two machines). The reason that more requests (60%) are sent to the first server to be treated is for the purpose of load balancing. The Load Balancing IP also plays a role in infrastructure availability: if one server is unavailable - verifying ping on ports, http (80), HTTPS (443), MySQL (3306) or PostgreSQL (5432) - traffic will not be directed to that machine until it becomes available again.

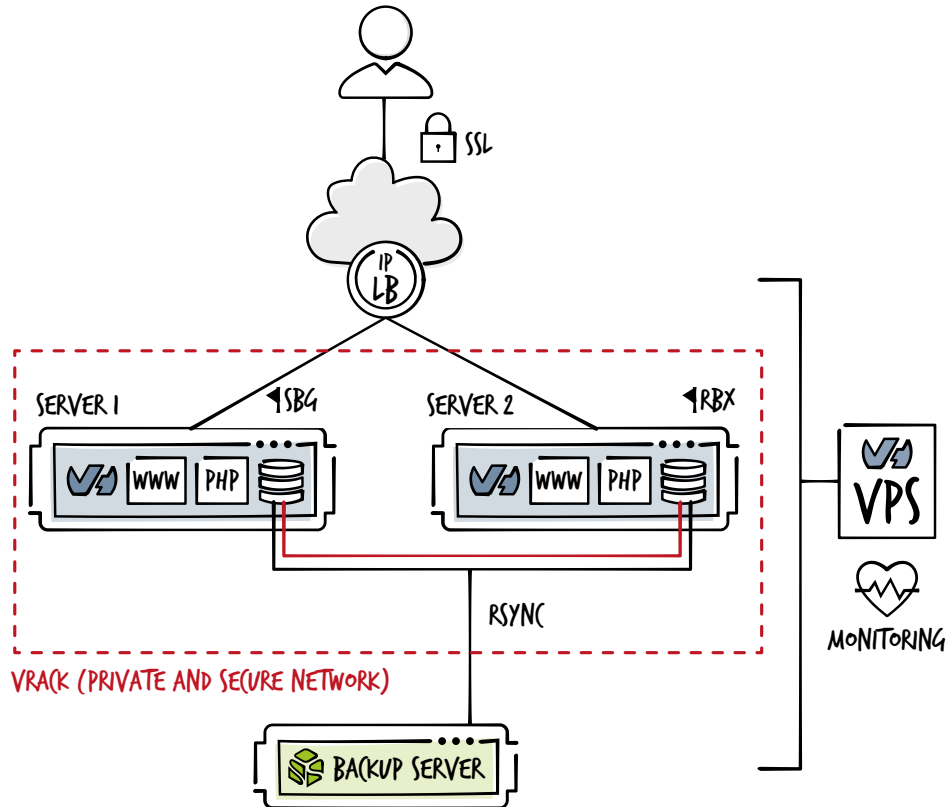
Finally, an SSL certificate is associated with the LB IP, securing user connections to the infrastructure.

The screenshot shows the OVH API interface for the 'ip' resource. The header includes the OVH logo and 'API.OVH.com' with a 'Login' button. Below the header, there is a search bar and a 'Filter' button. The main content area displays a list of API endpoints for the 'ip' resource, each with a method (GET, POST) and a description. The endpoints are:

- GET /ip: Your OVH IPs
- GET /ip/loadBalancing: List available services
- GET /ip/loadBalancing/{serviceId}/info: Get this object properties
- GET /ip/loadBalancing/{serviceId}/allowedBackends: List of backends you can attach to your IP
- GET /ip/loadBalancing/{serviceId}/backend: Backends for this IP load balancing
- POST /ip/loadBalancing/{serviceId}/backend: Add a new backend on your IP load balancing
- GET /ip/loadBalancing/{serviceId}/backends/{backendId}/info: Get this object properties
- POST /ip/loadBalancing/{serviceId}/backends/{backendId}/info: Add this object properties

The Load Balancing IP is configured via the OVH API.

Solution deployed





Bring high availability to
your infrastructure with
IP Load Balancing

www.ovh.com

A Distributed Infrastructure on Two Active Servers, Housed in Two Remote Datacenters

Each of the two servers contains the entire application: web layer, business layer and database (MySQL). The three-tier architecture is contained on the same server.

The first server is housed in OVH's Strasbourg (SBG) datacenter, while the second is located in Roubaix (RBX). Choice in server location can often bring the application closer to users, offering them the best access times. Here, the geographic dispersion of the two machines is designed to provide service redundancy in two separate areas, isolating any potential incidents from one another.

Under normal circumstances, the two servers share the load. In the event of failure, the functional server is by itself able to absorb user queries without affecting application performance. In the case of an application with more users, the application will function in a degraded mode until the situation returns to normal. The object is to assure continuity of business activity (BCP) by having a highly available infrastructure (HA).

Replication of Databases Between the Two Servers

The synchronous MySQL replication

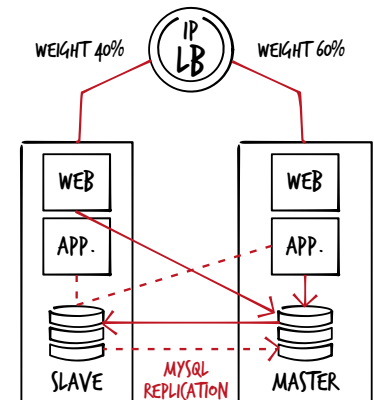
between the two master databases is a complex subject. Though the goal is straightforward (improve response times by distributing the read/write operations), implementation often gives system administrators a hard time, in particular when it comes to anticipating the problem of concurrent writes. This is probably the reason that the master/slave model continues to predominate hosted applications, like ERPs.

The solution adopted in this case is to fully exploit the resources of both servers for web and application layers, but only write to one of the databases. Operations are replicated using MySQL replication and synchronized to the slave, hosted on the other server, through the vRack. When requests are handled by the server containing the master database, writes are performed locally. When requests are sent to the second server, writes are performed on the master server through the vRack. Reading operations are performed on both databases.

OVH.com's vRack proves to be of great interest: securely connecting the two machines located in Roubaix and Strasbourg through a private network, while maximizing connectivity (10 ms latency between SBG and RBX,

with a bandwidth of 1 Gbps or more depending on the characteristics of the server's network card). Moreover, data does not pass through the public network. Located at the application layer is a script which manages write redirections to the active database. In case of failure of the server hosting the master database, configuration is transferred: the data on the slave becomes master and the weight assigned to each of the two servers in the LB IP configuration is automatically switched via the API. The transfer will stay in place until the next incident.

Monitoring is in place to alert the administrator of a loss of one of the two machines



while detecting the origin of the incident and reintegrating the machine into the infrastructure.

Backup to So you Start Server

A third server contains a backup of the database and is updated every hour (incremental backups) using the Rsync protocol over an encrypted SSH tunnel. As computing power is not a vital criterion when considering backups, a So you Start server was chosen as a matter of economics, cost vs storage.

The open source application Bacula is used to configure, verify and restore backups when required. The backup is configured so that the 24 hourly backups of the previous day are available as well as the hourly backups for the past seven days and weekly backups for the previous month.

These "snapshots" of the database are useful in many situations: simultaneous disk failure on the two servers, an attack or a software bug which corrupts the database or in the case of human error. These snapshots permit, within the time constraints, to recreate the infrastructure with the lost data. This is an essential element of the DRP (disaster recovery plan) of this infrastructure.

Monitoring

A VPS Cloud hosts the monitoring of the infrastructure: the resources of this virtual private server are more than sufficient and with its 99.99% SLA it guarantees high availability of monitoring.

Shinken software, whose interface is entirely configurable has been selected to oversee the infrastructure. Unlike other tools such as Zabbix, Shinken focuses on monitoring and not infrastructure metrology. In this case, the main goal of monitoring is to warn the IT provider of any incidents (notifications by e-mail or SMS). Concerning infrastructure metrology, this is to say, knowledge of the state of resources, we call on Graphite which ideally completes Shinken ■

This case was inspired by the infrastructure designed by Barme company for its client, Sextant Expertise and was awarded the OVH 2014 Infrastructure Trophy under the category "Enterprise Infrastructure (outsourcing)".



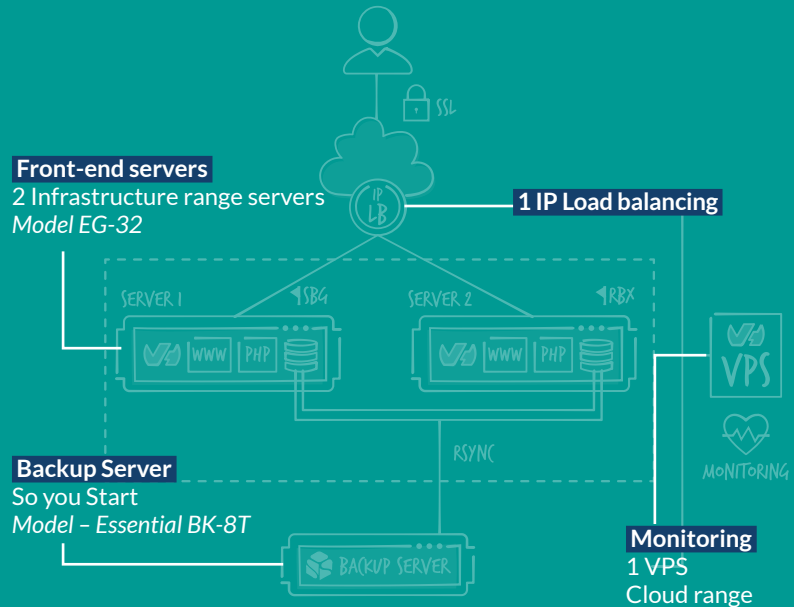
Laurent Barme
From the Barme
Company



OVH Advice

If the number of ERP users would grow exponentially, the infrastructure could evolve by separating the database from the application on the front-end servers. This would permit one to increase the number of database servers, in a manner to distribute the load over a larger number of servers. Moreover, the cost of backups can also be reduced by using Public Cloud Object Storage. The goal would be to only pay for storage that corresponds to the amount of data without having to provision hard drive resources and pay for a server whose CPU would be under used. With this solution comes added security, by default data is replicated three times on Public Cloud Object Storage.


Infrastructure Components



Server Private Interconnections

vRack

Service included with Infrastructure range servers



Intel
D3302S
VCI
PCI Express
Intel

3042011
200802

VCI

PCI Express

300 Servers per Day

OVH has its own assembly lines in France and Canada, outputting 300 servers each day. This industrialized production in conjunction with sophisticated logistics guarantees rapid server delivery and the ability to offer competitive prices.

The background image shows two men in a factory or industrial setting. They are wearing blue work jackets, clear safety glasses, and yellow earplugs. One man is in the foreground, looking intently at a piece of machinery. The other man is slightly behind him, also looking at the same area. The machinery is metallic and complex, with various pipes and components. The lighting is bright, typical of an industrial environment.

#03

High Availability and Collaborative Applications on Two Dedicated Servers

Edouard G, CIO of an industrial maintenance company, wishes to provide the company's mobile employees with secure VPN access to enable them to consult intervention schedules, access files, and review technical manuals in addition to providing them with the ability to work together in a collaborative work space.



Specifications

- Redundancy on two dedicated servers assuring high availability
- Master/Slave model with simplified data transfer
- Server synchronization via private network
- Monitoring to trigger activation of Master/Slave transfer
- Secure access (VPN service)

GMAO Extranet Collaborative Messaging High Availability Dedicated Servers
vRack (private network) VPN



Key Points

3 applications

(extranet, computerized maintenance management system (CMMS), collaborative messaging)

1 VPN IPSec service

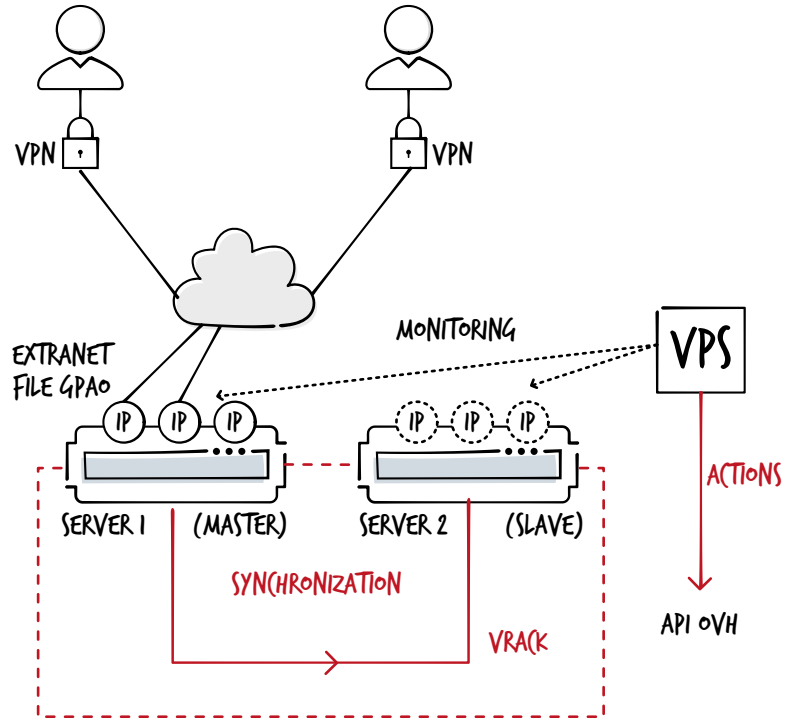
20 simultaneous
mobile users

Country
wide
access



Employees in the field access the company's main server using their mobile devices via VPN IPsec. Each service (extranet, file server, CMSS) has its own IP failover which can be moved from the master to the slave in the event of server failure. Upon failure of the master server, the monitoring VPS uses the OVH API to trigger the movement of the IP failover from master to slave. The slave server also has applications installed and the master's data is regularly synchronized with it over the private network. The VPS is also used to start/restart standby services on the second server ■

Solution deployed



“We opted for dedicated servers powerful in terms of CPU/RAM”

Edouard G. - IT Manager



Feedback

“My project was simple. I wanted to bring both dynamism and comfort to mobile users, while providing secure access. Availability is critical, the solution was to setup and ensure continuous availability of services during working hours. We opted for two dedicated servers from the Infrastructure range, powerful in terms of CPU/RAM and

comprised of a private network for synchronization. Having a large amount of memory permits scalability if other applications are deployed in the future. Today, my teams have access to their workspaces, files and our CMMS from anywhere in the world in an environment that provides high performance, security and redundancy.”

Infrastructure Components

2 Infrastructure Range Servers

Model EG-64

vRack

Service included with Infrastructure range servers

VPS

SSD range



SDG2

F

Strasbourg

In 2012, OVH implanted a datacenter in Strasbourg (East of France), not only enabling its customers to become closer to their Eastern European users, but also making it possible for them to build multi-sites infrastructures. An example requiring a multi-site infrastructure would be the implementation of a disaster recovery plan (DRP).

#04

Behind the scenes
of the infrastructure which
hosts the digital activities
of BFM, RMC and 01net

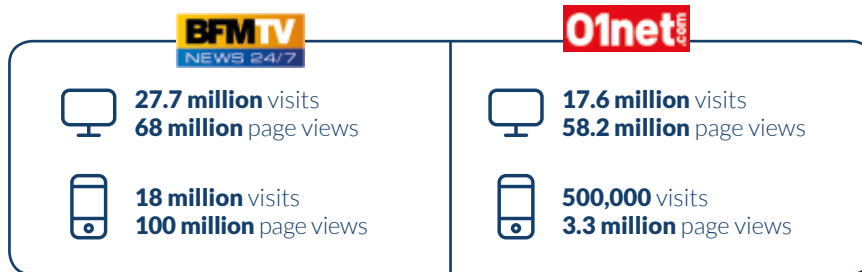
**Next
Radio
TV**

NextRadioTV is a French media group, present in the television, radio and digital landscape under the brands: RMC, BFM and 01net. To host the group's online media activities (web and mobile applications), NextRadioTV has deployed a high availability infrastructure, capable of absorbing the traffic of 9 million daily visitors. Combining dedicated cloud, public cloud and dedicated servers, NextRadioTV takes advantage of a hybrid infrastructure that is robust and quite scalable, which pools its resources to the benefit of its six websites and applications developed by the group. Discover how NextRadioTV designed this infrastructure which in 2014 replaced a fleet of more than 80 physical machines.

Specifications



- Load Peak – ability to handle an increase of 5 times the traffic in less than 15 minutes
- High Availability
- Platform Security – accessible to the public with the backoffice reserved to journalists



Hybrid Infrastructure | Load Distribution | HA | Scalability | Load Peaks
 Dedicated Cloud | Dedicated Servers | Public Cloud Object Storage | vRack

Audiences en avril 2015 (source Xiti)



Key Points

500 Mb/s bandwidth for a total of 5 TB of outgoing traffic for “normal” day

100 GB of stored images on Public Cloud Storage

60 GB database

50 journalists use the backoffice for content editing



The heart of the infrastructure deployed by NextRadioTV is comprised of a Dedicated Cloud equipped with seven XL Enterprise range servers (32 cores + 128 GB RAM) and eight datastores (6 x 1TB + 2 x 3.2TB) on which more than 80 virtual machines run. This dedicated cloud is completed with a dozen physical servers, used in part to form a Galera database cluster (10 machines) and secondly to monitor and analyze, in real-time, logs using ElasticSearch. Images are stored within Public Cloud Storage.

Load Distribution (DNS Round-Robin + HAProxy)

While performing DNS resolution for the domains hosting NextRadioTV applications, visitors are pointed to a scheduled type of round-robin (homogenous request distribution) on four machines running HAProxy. These four HAProxy virtual machines are each equipped with two

vCPUs and are running on four different hosts and four different datastores, by means of an anti-affinity rule created under vSphere (VMware hypervisor). In this manner, this level of the infrastructure can tolerate the loss of one host. These load balancers are only accessible on port 80 (http) and are the only machines of infrastructure exposed to the public network. They communicate with the VMs under them (Nginx servers) through the

vSphere 6.0 is coming soon and will eliminate the one vCPU per machine limitation in Fault Tolerance Mode (VMware function) allowing customers up to four vCPU per machine. NextRadioTV could then consider switching from four to two HAProxy VMs each with four vCPUs placed in Fault Tolerance mode, assuring high availability.

private vRack network.

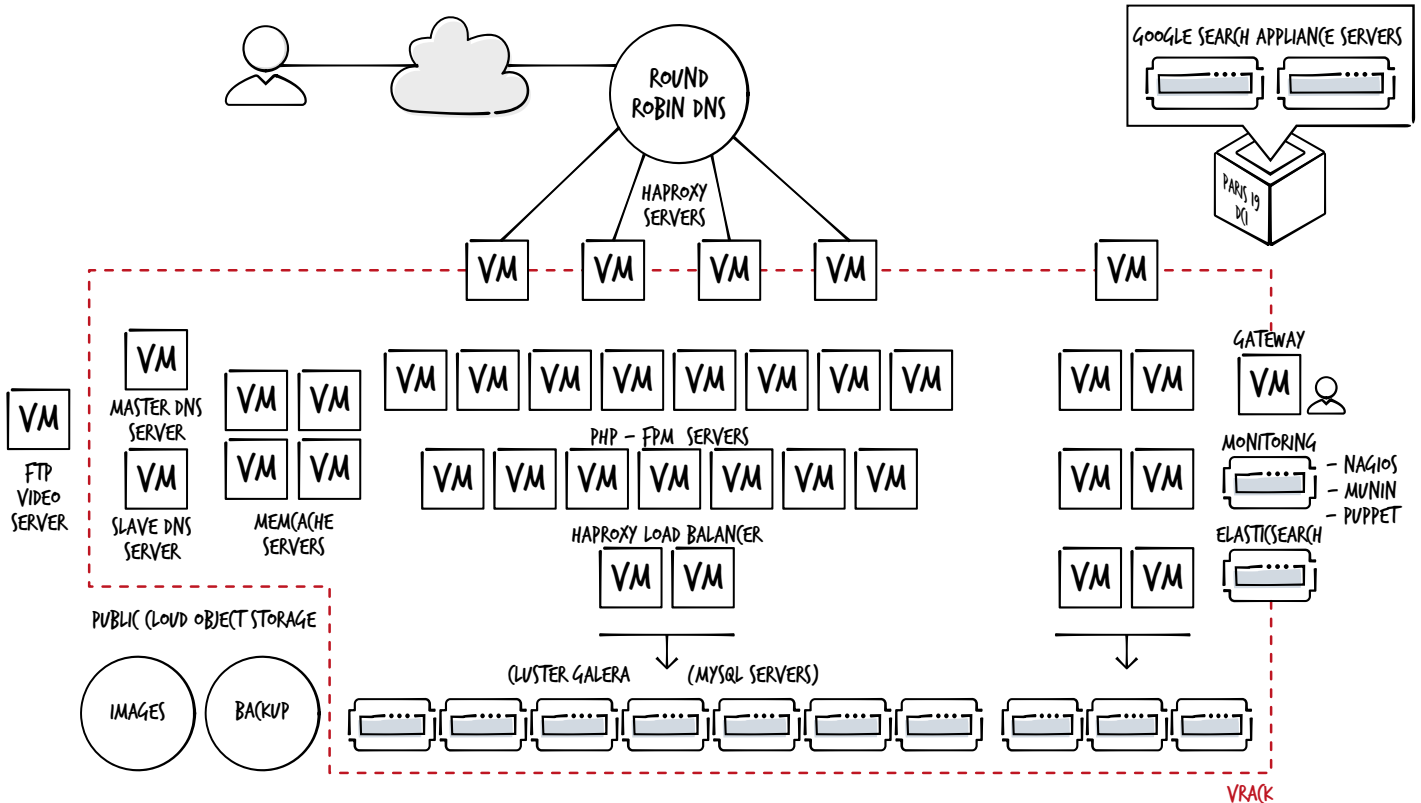
In the case of failure of one of the HAProxy virtual machines – which has yet to happen- the planned procedure will be to quickly associate the failed VM's IP with another VM. DNS round robin, which does not test page availability, will prevent one out of four requests from being directed to a 500 error page.

Web Servers (Nginx), Application Servers (PHP-FPM) and Caching Servers (Memcache)

There are eight virtual machines running Nginx which serve web pages and solicit a pool of seven PHP-FPM servers that are connected to a MySQL data base cluster. The Nginx servers cache all HTML pages (cache page FastCGI) in order to be able to serve pages, at a later time, without having to make additional requests to servers. The PHP-FPM servers, in turn, are connected to four caching servers with Memcache installed on them and are shared between all hosted applications. These Memcache servers conserve the different html blocks of the rendered pages (header, title, paragraphs of the article), allowing PHP-FPM servers to rebuild the pages as quickly as possible.

Image sources (uploaded by journalists)

Solution deployed



are stored in Public Cloud Object Storage. For the first call of an image, the web server retrieves the file using an Nginx module written in Perl via the OVH, API. As Public Cloud Object Storage is not yet compatible with vRack (it will soon be), a gateway server is used to exit the private network which isolates the infrastructure from the Internet. The Nginx web servers store image sources locally then resize images to meet the requirements of the web pages. The image source is cached and can be retrieved very quickly for the next visitors. Videos are registered on an FTP server then uploaded to a third party provider. NextRadioTV had a need for a turnkey solution (application layer) to provide video (store metadata, Flash/HTML 5 player, encoding and video storage and streaming).

On a final note, the use of two internal DNS servers are used for name resolution for the VMs and the resolution of external platforms solicited. The objective is twofold: accelerate DNS resolution, thanks to proximity of the DNS servers and eliminate generating traffic through the gateway server by caching DNS. The gateway server permits the virtual machines to access the Internet while allowing administrators to access the platform through a VPN.

Data base Cluster

Ten physical servers (infrastructure range, model EG-128) make up a Galera cluster of MySQL data bases. Physical servers are preferred over virtual servers due to the power that they are capable of delivering. The 10 machines of the cluster synchronize through vRack. Multi-master replication allows the simplification of code (not necessary to define which servers read and which servers write) and facilitates the addition or removal of servers from the cluster. Requests to these servers are directed by four virtual machines running HAProxy placed in Fault Tolerance mode, assuring high availability at the point of liaison with the data bases.

Two of the HAProxy virtual machines are dedicated to the public platform and distribute requests to seven of the ten database servers. The two other HAProxy VMs are dedicated to the journalist backoffice and direct requests to the remaining three servers of the Galera cluster. This separation permits the backoffice to remain operational even if the platform experiences an exceptional increase in traffic load.

Infrastructure Monitoring and Administration

A dedicated server (Infrastructure range, model EG-32) hosts Nagios and Munin, applications which monitor the system and network. This is in addition to Puppet, which makes possible configuration management of new servers (tool to deploy, in minimal time, new VMs at the points of infrastructure congestion). In the event of an incident, Nagios sends alerts to platform administrators by e-mail during the day and by SMS (via the OVH, API) during the night.

A dedicated server (Infrastructure range, model EG-32) collects data using a NoSQL data base and analyses in near real-time all of the infrastructure's logs (traffic, errors). This is made possible by ElasticSearch. In addition, NextRadioTV recently started using the IP-label service to receive different metrics from outside the OVH network. The New Relic application is being tested with the objective of optimizing code.

The administrators have access to the infrastructure via the gateway server (a

VM running pfSense) which they connect to securely through a VPN.

High Availability and Scalability of the Deployed Infrastructure

Currently, the infrastructure's seven hosts are each running at less than 50% of their full capacity. Thanks to VMware's Dynamic Resource Scheduler (DRS), the allocation of virtual machines is an automated process subject to the defined anti-affinity rules (for the HAProxy VMs downstream of the infrastructure, for example).

VMware's HA (High Availability) option is also activated, automatically allowing virtual machines running on a failing host to be restarted on a functional host (only possible if 50% of a host's resources are available). This coupled with the OVH guarantee to provide new hardware in 15 minutes or less, in the event of hardware failure, customers can rest assured that their infrastructures benefit with a very high level of availability.

Infrastructure scalability, with the capacity to absorb an increase in traffic, (peaks in load related to current events or a natural rise in traffic, in relation to the development of the digital activities of the hosted media) relies on its 3-tiered architecture and the possibility to create new virtual



machines within minutes. Configured with the aid of Puppet, virtual machines can be deployed rapidly at various points of infrastructure congestion (for example: Nginx web servers).

Today, NextRadioTV's Dedicated Cloud resources are sized to withstand peaks in load without having to add additional hosts. The goal, in the upcoming months, is to work on platform elasticity, making it possible to add additional resources, which are billed by the hour, in just five minutes. The idea is to "boost" the number of hosts on the platform, at the base as well as at the top, to optimize costs while being able to withstand even larger loads.

Infrastructure Security

As mentioned above, only the HAProxy virtual machines, assuring load balancing on the web servers, are exposed on the public network (with authorization only on port 80). They form the security entry point for the infrastructure, which resembles a demilitarized zone (DMZ). The goal is to secure the section of the infrastructure that is in production. Fail2ban is used to secure SSH connections and implementation of a system of prevention and intrusion detection is under study (IPS + IDS).

NextRadioTV purges backups using scripts. Even more simple: Public Cloud Object Storage makes it possible to define an expiration date for uploaded files. On the selected date, the files are deleted without requiring any other action. How do you configure the expiration to T+21 days and restore files?

```
$ curl -X POST <endpoint>/<path>/file.ext -H "Content-Type: application/json" -H "X-Auth-Token: <token>" -H "X-Delete-After: 1814400"
```

```
$ curl -X GET <endpoint>/<path>/file.ext -H "Content-Type: application/json" -H "X-Auth-Token: <token>" -I [...]
X-Delete-At: 1438584553
```

Critical Data Backup and Thoughts on Disaster Recovery

Initially, databases, infrastructure logs and critical tools (git/ticket system) were backed up to physical servers. Today these backups are made on Public Cloud Storage, mainly for economic reasons. The retention for backups is 21 days, with the exception of HTTP logs, which are held for 1 year. A

monthly backup of SQL data bases is also preserved.

To setup a disaster recovery plan (DRP) with the least amount of interruption to service (RTO) as possible, NextRadioTV is contemplating putting in place a second Dedicated Cloud which would be hosted in Strasbourg. Smaller in size than the Dedicated Cloud hosted in Roubaix, the second infrastructure would be identical in architecture and

synchronized to the first through vRack. It will be capable of taking over at any moment if an incident should occur in the RBX datacenter, placing the platform in degraded mode. While the main infrastructure is “up”, no requests shall be made to the backup infrastructure (the backup infrastructure will be placed in “sleep” mode).

Backoffice for Journalist and Pre-production

An infrastructure similar to the public accessible platform (but smaller in scale) has been deployed for journalist allowing them to enter their work into the CMS. Today this platform is secured by restricting and filtering IPs and soon by VPN access (Open VPN). This is a more practical solution for journalist in the field, allowing them to connect using different IPs. An eighth Host XL, (32 cores + 128 GB RAM) isolates the rest of the infrastructure and is dedicated to development and pre-production + Git.

Google Servers Search Appliance

NextRadioTV has installed two of their own Google Search Appliance servers in



NextradioTV group is known as BFM brands, O1net and RMC.

OVH's Paris datacenter (DC1). These two servers make it possible to offer visitors a search engine based on the Google algorithm. One machine is dedicated to the indexation of platform content while the other is used for search. The two servers

have been replicated on each machine. This allows either machine to take over both roles in the event one machine should malfunction ■



Feedback

“Previously, the NextRadioTV infrastructure was composed of a farm of 80 physical servers. To manage this many machines multiplies the risk of incident. In practice, we received alerts night and day, especially concerning hardware issues. We were forced to mobilize a team of system administrators to perform jobs that had little actual value. Today, the Dedicated Cloud solution allows us to delegate hardware management to OVH. Except for the Galera data base cluster and the two machines used for monitoring, we now use the infrastructure without worrying about the hardware layer. The Dedicated Cloud solution has simplified the multi-use of one platform to the benefit of all the group’s sites and applications. Recap: We have gained flexibility while reducing our costs by approximately 40 to 50%. It is now much faster to deploy a new VM rather than a dedicated server. The complexity of the project resides in the way that the migration of dedicated servers

to Dedicated Cloud is made without any service interruptions. For this, first we configured all the dedicated servers in the same vRack. Next we deployed the web servers on the Dedicated Cloud before configuring the DNS to point towards the new infrastructure. The migration of the MySQL servers to the Galera cluster was carried out by MySQL replication.”



Vincent Lae,

Manager of the IT infrastructure at NextRadioTV.

Infrastructure Components

Dedicated Cloud

8 hosts XL
Datastores 6 x 1 TB + 2 x 3.2 TB

Database (cluster Galera)

10 Infrastructure Range Servers
Model EG-128


Monitoring

2 Infrastructure Range Servers
Model EG-32

Public Cloud Object Storage


vRack

Service included with Dedicated Cloud and Infrastructure Range Servers



For their mission critical activities, thousands of CIOs, IT Managers, software developers and web agencies choose **Dedicated Cloud**

www.ovh.com



Water Cooling

Since 2003, OVH.com has relied on its exclusive water cooling system to dissipate 70% of the heat produced by its servers. In conjunction with datacenters designed to control airflow, the remaining 30% of heat is removed through the air. As of 2010, this system has made it possible for OVH.com to totally eliminate air conditioning from its datacenters and achieve a remarkable PUE of between 1 and 1.2.

#05



Virtual Desktop Infrastructure, Information System and Video Protection: The Urban Community of Bourget Airport Uses Dedicated Connect and Dedicated Cloud to Outsource Its IT at OVH.com



Specifications

- Completely separate infrastructure and Dedicated Connect links for video protection as required by regulations
- Redundant Dedicated Connect links using 2 different routes between the core of community's network and the OVH network in Roubaix
- Hosting in France

IT Outsourcing | Video Protection

Virtual Desktop Infrastructure (VDI)

Dedicated Connect | Dedicated Cloud



Key Points

1st Urban Community
in France to totally outsource their IT

80 Full HD video
protection cameras

10 Gpbs
Internet connection
Dedicated Connect links:
4 x 10 Gbps for data
4 x 10 Gbps for
video protection

150 Virtual Machines
for job oriented applications

1,000 Virtual Machines
for work stations
(650 active today, with 400 simultaneous connections every day of the week)

All hosted on

7 X Dedicated Cloud



The towns of Le Bourget, Drancy and Dugny and their 100,000 residents make up the Urban Community of Bourget Airport. This conglomeration of towns has a particularly forward thinking IT team, which in 2004 had high speed fiber installed throughout the community's public institutions (schools, libraries, municipal police stations, health centers...). Six years later, as head director of digital development for the Urban Community of Bourget Airport, David Larose undertook the project of virtualizing the urban community's servers using the VMware solution, ESX, with the goal of migrating the entire server farm – a little more than 1000 work stations, for nearly 2000 agents and 60 professions – towards thin clients. This is a solution which has permitted the optimization of the server farm, with machines running, on average, at 30% of their capacity.

In 2012, Mr. Larose took things a step further by launching a call for bids to outsource the entire server farm of the municipality with a

provider located in France. David goes on to explain, "On one hand, technology advances rapidly and hardware is quickly obsolete. On the other hand, a public authority does not have the budget to regularly upgrade its machines... nor the funds to recruit experienced system administrators, whose salary expectations are elevated due to the shortage of such professionals in the IT sector. Therefore, outsourcing has imposed on us the best solution."



David Larose, IT Manager.

OVH.com won the bid by proposing to directly connect the core of the community's network to a Dedicated Cloud infrastruc-

ture housed in Roubaix, France, through a solution called Dedicated Connect.

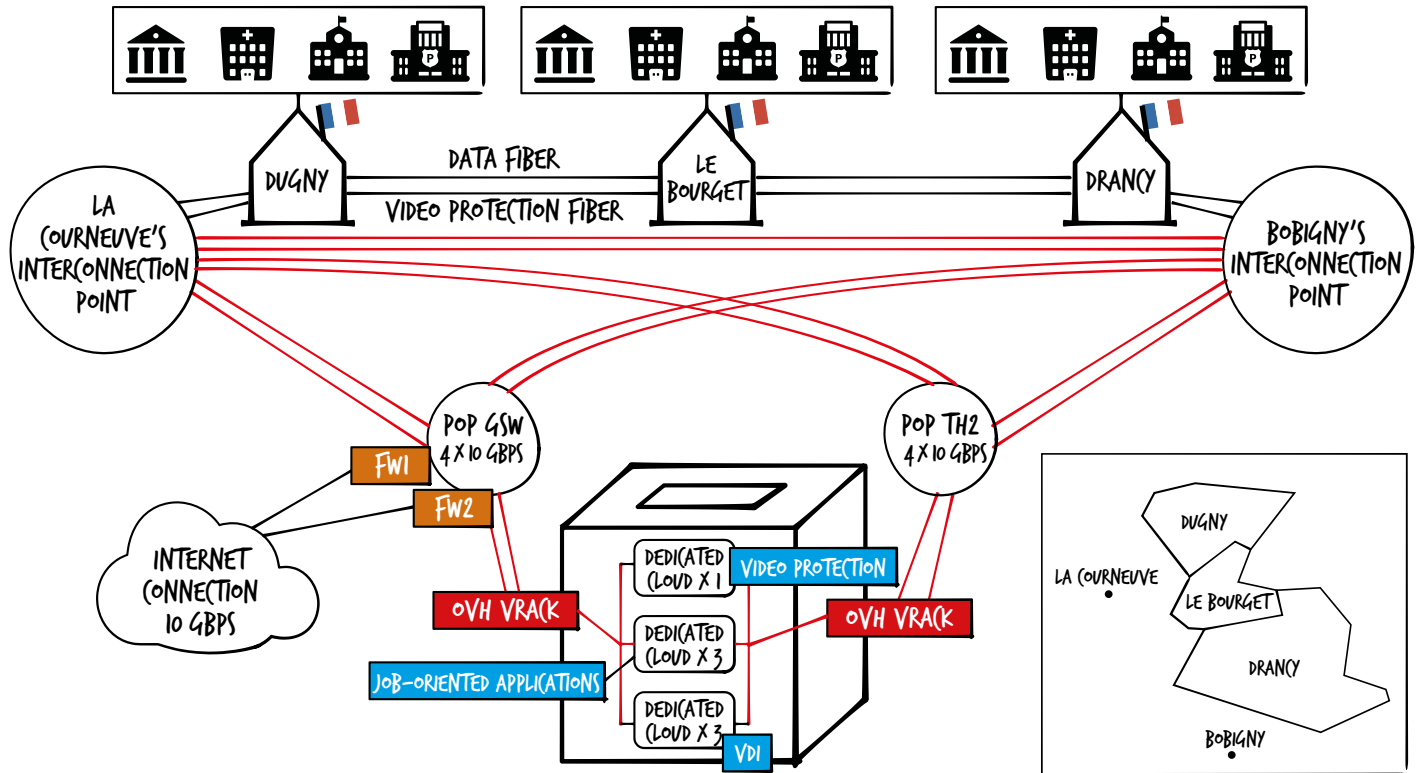
Deployment of fiber liaisons between the urban community's network and two interconnection points on the OVH network

In 2012, the urban community commissioned an important, civil project which established fiber optic connections between Darcy's town hall and an OVH interconnection point situated in Bobigny along with Dugny's town hall being linked to a second OVH interconnection point situated in La Courneuve.

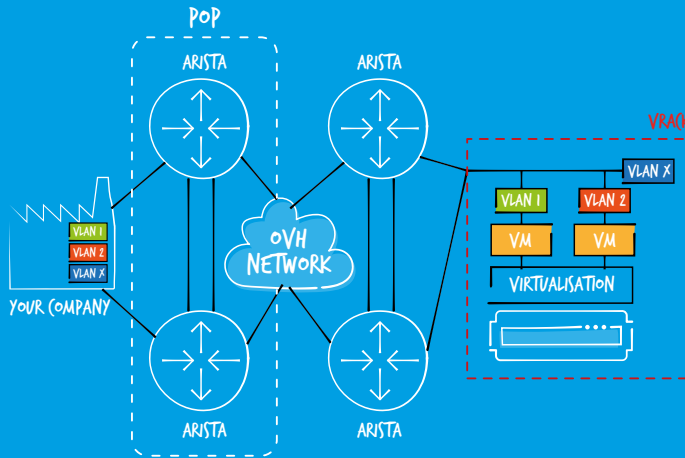
The goal: to benefit from a 2x10 Gbps fiber connection with the security of a second loop of the same capacity. On each loop, one fiber is dedicated to the data transfer of video protection only while the second fiber is used to transfer all other data (remote desktops and system information of the three communities).

One of the requirements of the call for bids included that the selected operator (OVH.com in this case, retained in 2012) extend fiber connections from a point of presence (PoP) to its backbone. In comparison to the cost of a 100 Mbps fiber connection offered by operators, "The investment of pulling fiber to the towns of Bobigny and

Solution Deployed



Dedicated Connect: How does it work?



With Dedicated Connect, OVH has its hand on the entire route, taking responsibility for availability and performance from the interconnection point all the way to the infrastructures (Arista is redundant at each MLAG-step). Another advantage: Dedicated Connect operates at the network level (L2), the equivalent of a direct physical connection between two devices.

In this manner, VLANs created within the urban community's network core are linked to Dedicated Connect and the outsourced infrastructures hosted at OVH.com. It is not necessary to rebuild an IP address plan: the OVH infrastructure and the municipality's machines are as if they are all connected on the same LAN.

La Courneuve, will see a ROI in 10 years, all while providing 100 times the bandwidth and the ability to communicate through a totally private network with our outsourced infrastructures.”

Dedicated Connect from OVH's Two PoPs in the Paris Region to Its Datacenter in Roubaix

From the interconnection points in Bobigny and La Courneuve, OVH extends the urban community's fiber network to its two PoPs located in the Paris area (Global Switch(GSW) and Telehouse(TH2)). Data is then channeled through four distinct fibers using two different routes with traffic terminating at the RBX datacenter, where the community's outsourced infrastructures are located.

At the Global Switch point of presence, the community has deployed two Cisco ASA firewalls coupled with a Sourcefire module. All of this provides the community's public employees with a secure and redundant 10 Gbps Internet connection.

Infrastructure Outsourced at OVH

The infrastructure deployed in Roubaix for the Urban Community of Bourget Airport is made up of seven Dedicated Cloud solutions.

One Dedicated Cloud is reserved solely for the task of collecting and processing video protection data and today there are 80 HD cameras within the infrastructure. The application Milestone is used for video management, with the software installed directly on the Dedicated Cloud. This application manages encryption as well as watermarking.

Besides not having to be preoccupied with hardware management and especially not having to add or replace disks, the main interest in outsourcing the infrastructure lies in the fact that Dedicated Cloud resources are elastic.

“Today’s regulations require that images captured from the video protection system be retained for 15 days. Due to this regulation, infrastructure storage capacity is allocated accordingly. If the law should change in the future and the required retention time increases, in just a few minutes, we could add an additional datastore at OVH. There is no risk on investment and we can be reactive in the event of regulatory changes.”

As for the remaining six Dedicated Cloud solutions, each of the three towns, which collectively make up the urban community, have been allocated with two Dedicated Cloud infrastructures each, to host business applications and employees’ remote desktops.

Three of the Dedicated Cloud infrastructures are used for hosting business applications along with a total of 150 virtual machines, on with which more than sixty business applications run and are linked to different public services managed by each town: school files, cafeteria, vaccinations, planning, roads, military census... Thanks to Dedicated Cloud, the platform is shared amongst the different applications, even though they operate within various environments. Data is secure as it is locally and redundantly stored on two datastores. This is in addition to being backed up by Veeam Managed Backup, the service offered by OVH.com to automatically backup critical VMs hosted on supplementary OVH infrastructures and is entirely managed by OVH.com.

The remote desktops of each municipality are accessible by public employees not only from the office but also from their home or anywhere on the go as virtual offices are also hosted on a Dedicated Cloud (one infrastructure per community) and are based on the VMware solution, Horizon. Combined, the Dedicated Cloud infrastructures have the capacity to run 1000 VMs hosting remote desktops (there 650 active today, with 400 simultaneous connections being made each day of the week). The arri-

val of these VDIs has significantly changed the daily routine of the urban community’s IT team.

VDI (Virtual Desktop Infrastructure) simplifies the following three areas:

1. template management/ images of different OS,
2. the rapid deployment of remote desktops
3. help desk support of these remote desktops.

These three areas can also help to define KPIs and should be examined prior to migration.

“Today we create templates for each type of work environment and deploy on demand. These templates contain software suites, business applications and are configured with the necessary permissions/rights to allow each employee to perform their job. Configuration, maintenance and updates, tasks which used to make up much of our activity, have now become much more



simplified. We are now able to devote ourselves to more interesting projects, for example, the deployment of interactive whiteboards in our schools or Internet access in our public libraries. Previously, making a work station available to our citizens was a real challenge. Now it is very easy. After each use, virtual machines are deleted and then recreated for the next user. Viruses are not persistent, bookmarks and user parameters are reinstalled... and we can offer each citizen the best experience possible."

Focus on the Virtual Desktop Infrastructure (VDI)

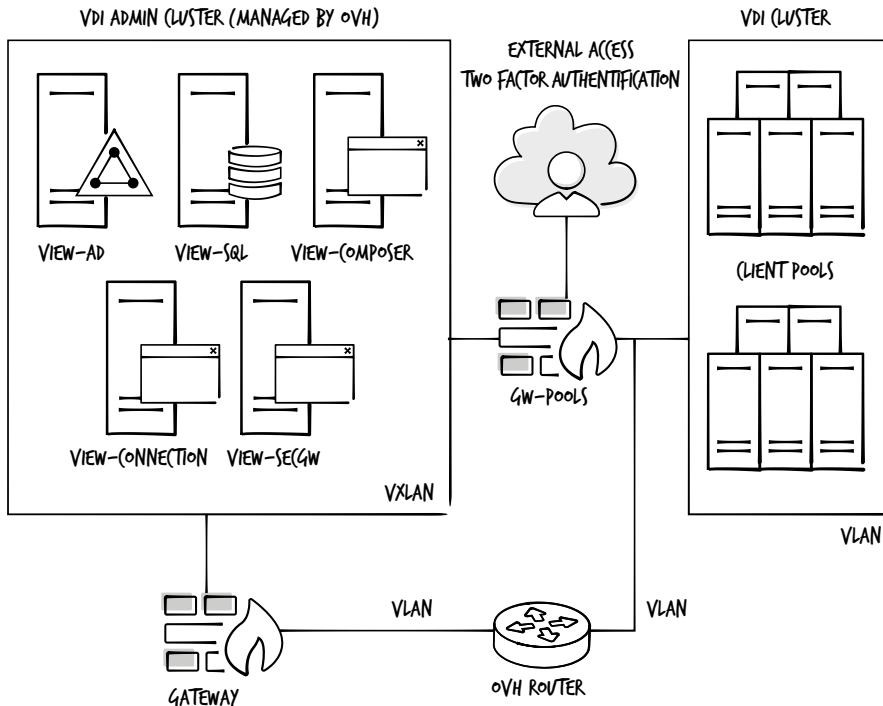
The platform put in place by OVH to provide the Urban Community of Bourget Airport with its remote desktops is particularly interesting due to the fact that it was used to develop the VDI solution as a service which OVH is preparing to bring to market. One part is a shared infrastructure (PaaS) and the other part, a form of dedicated infrastructure (IaaS) as is the situation in this use case study.

This platform relies on the VMware solution, Horizon. Two hosts of each Dedicated Cloud are used by OVH.com to form a cluster for administration, giving the urban community's system administrator read only access. On this cluster, different components of the VMware solution can be found: server connec-

tion, active server directory, central server, database server, and security server. On the resources cluster, administered by the urban community, pools of VDI clients are allocated on the hosts using VMware's DRS. Through dedicated VLANs, the VDIs have access to the application servers and their data or even the fourth Dedicated Cloud which is reserved for video protection, as is the case for the VDI of the municipal police. The entire route, from end user to the infrastructure hosted in Roubaix, is encapsulated in a private network, which is totally isolated from any exchange with the public network. The inter-VLAN is managed by one of the first versions of the virtual router (vRouter) that OVH developed, just like the firewall that enables a secure gateway (1 output for internal, to the administration cluster and one input/output for VDI, secured by a two factor authentication, Google Authenticator) ■

The Infrastructure of the Urban Community of Bourget Airport

A Close-up Look at the Dedicated Cloud devoted to the Towns' VDIs



Infrastructure Components

Dedicated Cloud

Dedicated Connect



OVH

an International Enterprise

OVH has a customer base that extends throughout 137 countries and a with its 16 subsidiaries physical presence on three continents, in June 2015, the enterprise reached a milestone when it hired its 1,000th employee.



#06

A High Availability
Infrastructure, Resistant
to Increases in Traffic,
Providing Background
Music for 7,000 Points
of Sale Around the World

radioshop

Based in Montpellier, Paris and Shanghai, RadioShop specializes in sound design. Its job consists of creating custom made playlists for outlets such as the Galeries Lafayette Paris Haussmann, Habitat retail stores, Yamaha concessionaires, Subway and Memphis Coffee restaurants. To match the "DNA" of a piece of music with that of a point of sale, then create and broadcast playlists that adapt in real-time to local weather conditions or even to the crowds of the stores, Radioshop relies on the Dedicated Cloud of OVH.com



Specifications

- Infrastructure simple to maintain, that is secure and scalable
- High availability
- Capable to handle peaks in load
- High bandwidth internationally

High Availability | Load Peaks | Infrastructure Security

Private Network (VLAN, vRack) | Dedicated Cloud



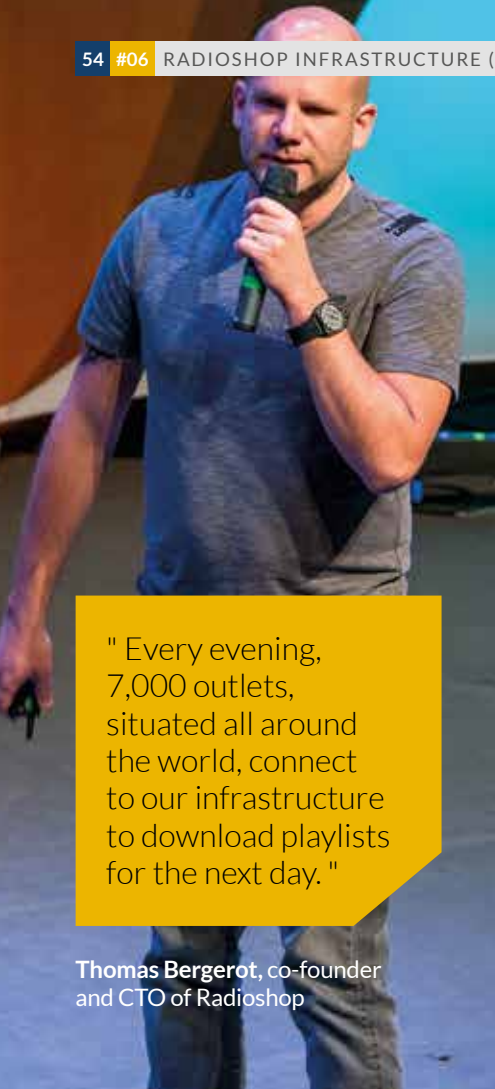
Key Points

7,000 points
of sale worldwide



500,000 music tracks
in the musical catalog

Up to **7,000**
simultaneous requests
on the infrastructure



" Every evening, 7,000 outlets, situated all around the world, connect to our infrastructure to download playlists for the next day. "

Thomas Bergerot, co-founder and CTO of RadioShop



RadioShop started its business with one dedicated server, then two, then three... and quickly the number of machines increased, causing a significant amount of administration and maintenance of the servers: management of hardware failures, OS updates, attack prevention, etc. However, as Thomas Bergerot, co-founder and CTO of the company, says, "Our job is not to manage servers. More precisely, as a computer scientist, my value is far greater when I am developing new services for our customers, in connection with creating in the sound environment." This is why RadioShop was interested in OVH's Dedicated Cloud solution. They were quick to delegate machine management to OVH and take advantage of an easily exploitable production environment that offers guaranteed high availability.

"Every evening, 7000 outlets, situated all around the world, connect to our infrastructure to download playlists for the next day. The stores' networks are generally less powerful and during the daytime are mono-

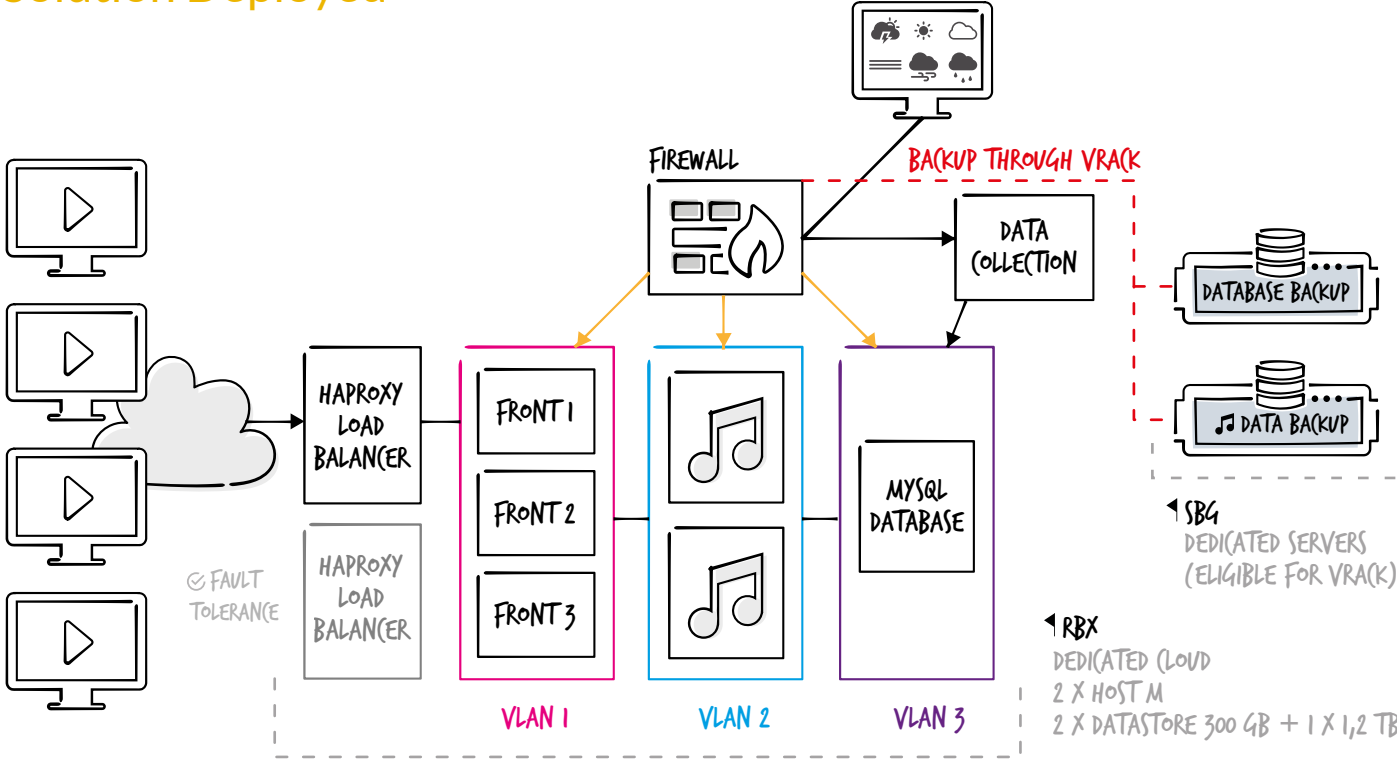
polized by electronic banking. Therefore, the time available to distribute playlists is relatively short. It is essential that our network is accessible during this short window of opportunity."

A Three-tier Architecture

RadioShop has built a three-tier infrastructure on the OVH Dedicated Cloud with each stage being encapsulated in a different virtual private network (VLAN). This secures back office access in the event of an attack. The first level consists of three web servers, in front of them an HAProxy server is in place to distribute the load evenly. The HAProxy machine is placed in Fault Tolerance mode (VMware functionality) as it is mission critical (it is the SPOF of the infrastructure). This means that the HAProxy VM is duplicated on another host and synchronized in real time. In this manner, it can take over immediately, without the loss of a single ping, in case of failure of the VM HAProxy "master" host. Moreover, it is the only machine of the infrastructure to be exposed to the public network.

The second level is comprised of two database servers, each containing the entire soundtrack catalog. Finally, the last level contains the MySQL data base, populated by data servers and another server (data collection server)

Solution Deployed



which collect, in real time, information from the web such as weather around a point of sale, in order to influence the playlists that will be generated.

Infrastructure Security Administration

To connect to its infrastructure, RadioShop implemented, on a virtual machine, a pfSense firewall configured to allow only specific IP addresses to connect. This VM has a public network card and several private network cards, in this way, having a foot in each of the infrastructures VLANs. This firewall also acts as gateway to allow the data collection server to access the Internet in a secure manner.

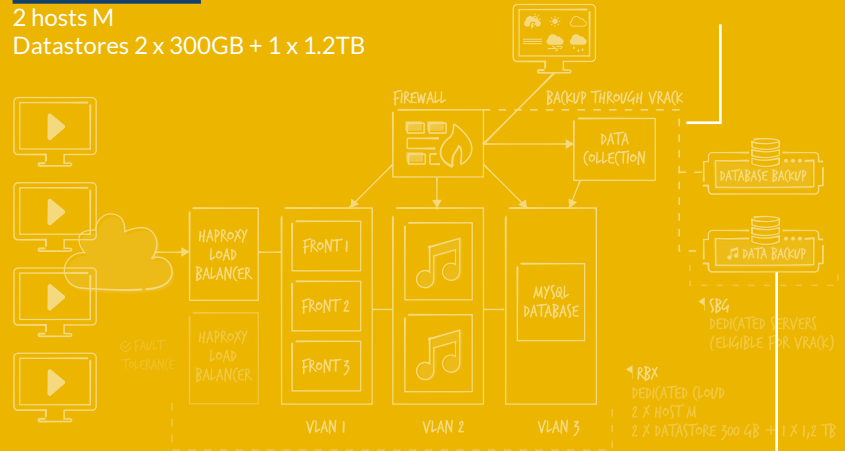
Backup Data to a Remote Site

The data base and music catalog files are backed up to storage servers which are located in Strasbourg (SBG datacenter). All data transferred between Radioshop's main infrastructures, located in Roubaix (RBX datacenter) and the backup servers is through vRack which provides a private network, assuring that all data remains secure. ■

Infrastructure Components

Dedicated Cloud

2 hosts M
Datastores 2 x 300GB + 1 x 1.2TB




vRack

Service is included with Dedicated Cloud and Infrastructure

Discover RadioShop's Testimonial Video:
ovh.to/YstRZYv

Backup

2 Backup range servers
Model FS-12T



Simplify Your Life and Get Creative!

Automate tasks using basic scripts,
and optimize your own functions. Create your own
application. Combine OVH APIs for amazing results!

api.ovh.com



OVH.COM


OVH.COM

Certified Recognition of OVH Expertise

OVH is engaged in a global certification strategy. OVH's Dedicated Cloud has already obtained ISO 27001, SOC and certification. OVH is capable of deploying private racks within its datacenters, to respond to even stricter data security requirements.

#07

Simplify Implementation of a Backup Strategy with Veeam Managed Backup



A university offers its students a “social network”, a platform to exchange accessible content with students, professors and administrative staff. In addition to instant messaging, this social network offers the possibility to create and join public or private groups as well as store and share content of all types. This service is accessible from any type of device (desktop, smartphone, tablet...) and users can connect 24/7/365. Let's focus on the backup strategy in place.



Specifications

- Disaster Recovery Plan (DRP) with Recovery Point Objective (RPO) of 24 hours. (RPO: It is the maximum targeted period in which data might be lost from a service due to an incident.)
- Control costs

Backup Strategy | Dedicated Cloud | Backup | Managed Veeam Backup



Key Points

20,000 users

13 VM hosted
on Dedicated Cloud

100 MB
records in the database

100 GB stored files



The university's social network is hosted on an OVH.com Dedicated Cloud and easily managed using VMware's vSphere hypervisor. The infrastructure is comprised of dozens of virtual machines: web servers (load balanced by an HAProxy VM placed in fault tolerance mode), a cluster of database servers, file server, cache server and, on the side, a development server (with a versioning system) in addition to a pre-production server that is used daily by the university's IT developers.

Definition of the Backup Strategy

In developing the backup strategy for this infrastructure, the university's IT department first identified what data was required to be backed up to restore service in the event of a problem. The Dedicated Cloud solution is designed to assure high availability, but this does not shelter the service from a hack, errors in manipulation or any temporary issues with infrastructure hard-

ware (disk failures, electrical problems...). This is why it is necessary, according to the importance of the application and/or the terms of the contract linking the platform to its users, to take necessary measures to be able to restart the service within the minimum amount of time and with the least amount of data loss.

In our case, the web servers containing the production version of the application, which are also present on a development server, as well as earlier versions are preserved by versioning system. In turn, the cache servers contain only volatile data: their role is to store the results of the most common queries to avoid unnecessary requests to the databases. Ultimately, only the file, data base and the system versioning servers contain critical data.

The file server contains the documents stored and shared by students across the social network. The data contained in the database VMs cannot be lost, otherwise the service could not be restored to the state it was in prior to any incident. The data on the development server represents many months of work and must be stored outside of the infrastructure. In addition, the backup of the VM used as the office VPN server, admin node and configuration management



server (under Puppet), is indispensable for quickly restarting the infrastructure if required.

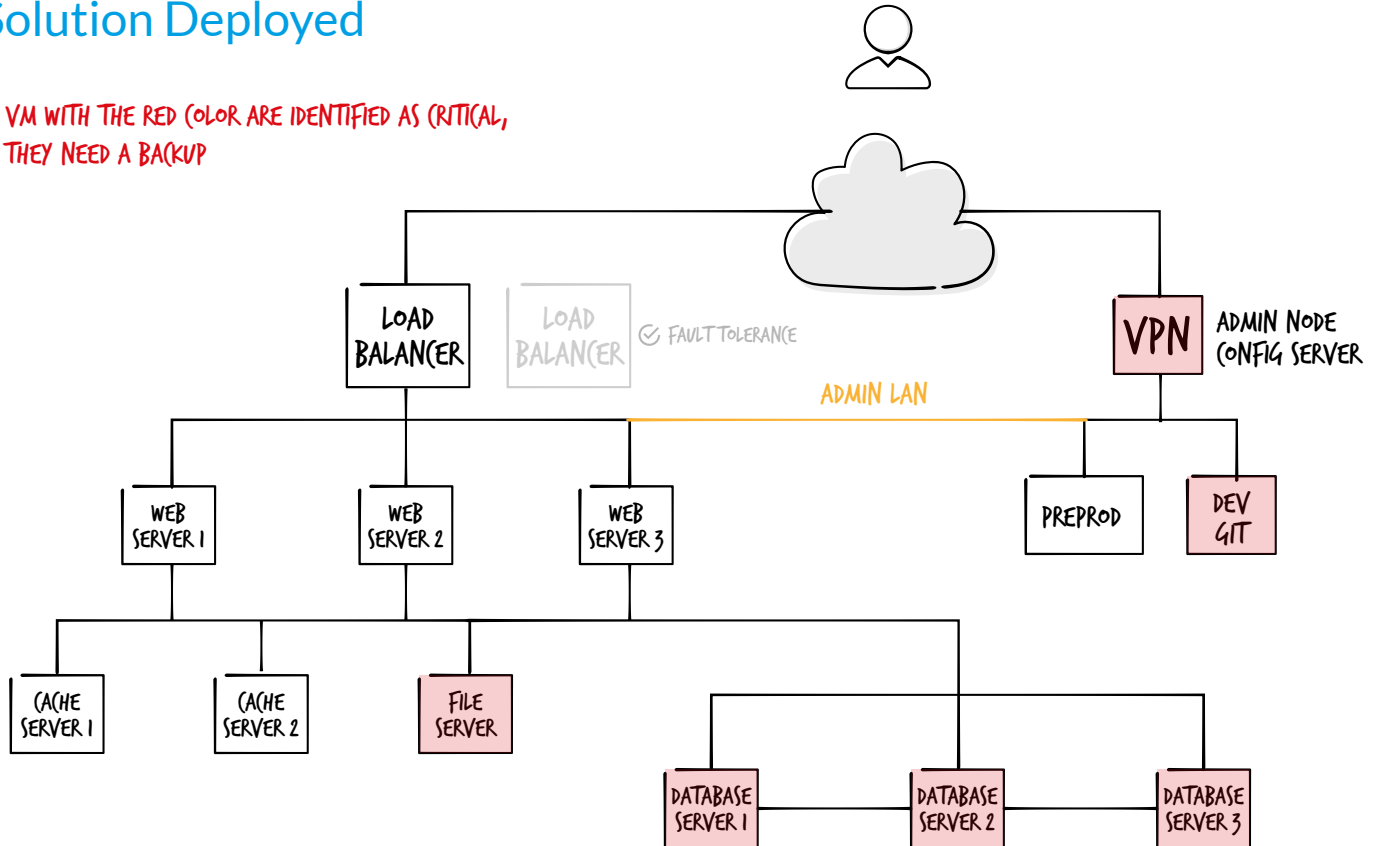
Implementation of the Backup Strategy with Managed Veeam Backup

It has been determined that a total of six VMs were identified as requiring regular backups. For reasons of cost and efficiency, the university's IT department has chosen to make these backups using the "Managed Veeam Backup" solution.

The alternative, considered more expensive and complicated to set up, was to opt for a storage server, connected to the infrastructure through vRack. An invest-

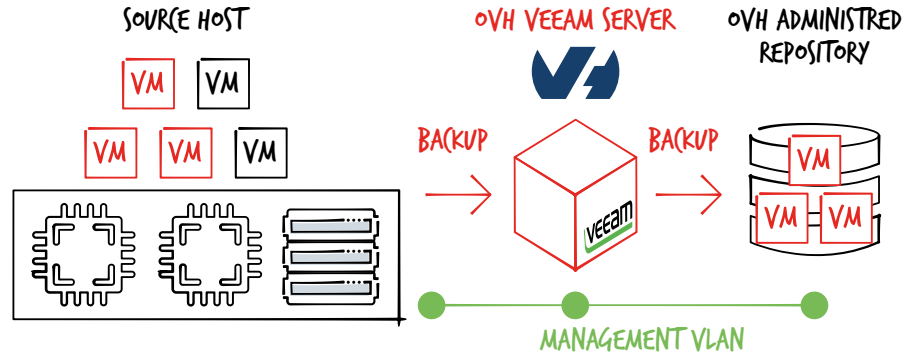
Solution Deployed

VM WITH THE RED COLOR ARE IDENTIFIED AS CRITICAL,
THEY NEED A BACKUP



ment requiring more time being spent, each month, configuring the machine by putting in place the Rsync protocol, verifying and testing backups, monitoring and maintaining the server and having to worry about possible hardware failures.

"Veeam Managed Backup" is a simple solution, being an option, which can be activated from the OVH customer control panel with the user only having to select the VM to backup by right-clicking within the vSphere hypervisor. Backups are made on a space which is independent of the of the user's Dedicated Cloud (mounted RAID repository, hosted and managed by OVH). In other words, there is a backup solution as a service. Data passes through a managed network (VLAN) with backups having no effect on network performance. Another advantage is that the space required to back up data isn't deducted from the Dedicated Cloud's storage space, where user data is stored (their datastores). The only thing that users see running on their Dedicated Cloud is a Windows VM, created by OVH, that consumes very little resources (1 vCPU). Billing of the service is transparent: \$8.50 excl. VAT per month per backed up VM, no matter the size. In this case, it would be \$51 per month compared to the cost of nearly \$60 per month for an OVH.com



Thanks to the OVHVeeamServer, managed backups, of selected VMs, are made daily. The infrastructure management vLAN permits the transfer of data to the repository which is RAID mounted and fully managed by OVH's teams. With this complementary storage, the user benefits from space used by its datastores.

dedicated backup server. Also, the price of a dedicated backup server doesn't take into account the required time investment, by the university's IT team, to configure and maintain such a server.

Backup Retention

A daily backup, with a snapshot of the VM which the automatic backup requested is sent to the repository. The backup retention period is between 14 and 20 days, allowing users to restore data to the same state as several days prior to any incident. For example, restoration to an earlier date is useful when a virus infects a service

with several day passing before administrators become aware of any issues. In practice, Veeam Managed Backup performs a complete weekly backup of a VM, plus 12-18 incremental backups. As a consequence, "full" backups take longer to perform (dependent on the amount of data the VM contains) whereas backups from subsequent days, only contain the changes to or addition of files since the last backup. Each day, a report is sent to the infrastructure's administrator to inform him of the successful completion of backups. In the event of an incident on the Dedicated Cloud infrastructure, the university's social network can restart with a maximal data loss

Infrastructure Components

Dedicated Cloud

Managed Veeam Backup

of 24 hours. This RPO (Recovery Point Objective) can be attained with the use of a configuration management tool, like Puppet. This allows for, after the restoration of vital VMs, the replication of the production's multi-tier architecture.

Restoring a Backed-up VM

Restoring a VM is just as easy as the activation of the automatic backup. A right click on the VM to be restored within vSphere launches the restoration job and a new virtual machine will be created from the selected backup. For example, "veeam debian (restored to 11-03-2015 22:00:58)". ■





OVH World Tour

Since 2014, the OVH World Tour has stopped in 29 cities throughout 10 countries, bringing over 12,000 customers face to face with OVH experts. This direct connection, peering between OVH and its customers, provides the perfect opportunity for end users to communicate exactly what they require from the OVH, in order to be successful in their professions. Listening to this feedback OVH.com is able to develop products, services and functions that are relevant and meet the needs of its customers.

#08

A Scalable Infrastructure Capable of Handling Large Peaks in Load

SO FOOT



SYSADMIN BADASS is the provider that designed the infrastructure for SoFoot.com, a site devoted to football news. Initially hosted on a single dedicated server, the site has experienced increased success since its inception in 2006. On the eve of the 2014 World Cup, it became necessary to put in place a more robust infrastructure and above all, scalable, in anticipation of the expected increase in visitors (more than one million page views per day with 500,000 unique visitors). The goal: zero interruption in service during the month-long competition.



Specifications

- Ability to handle increases in traffic
- Inability to overhaul the site code (very limited availability of the development team), the site was designed to run on a single machine (application not "cloud ready")
- Limited budget, cannot permanently over build the infrastructure
- 100% site availability

Scalability | High Availability | Hosting High Traffic Site

Dedicated Server | Public Cloud Instances | IP Load Balancing



Key information

85,000 articles

1,2 million
user comments

60 GB of static content
(images, videos, etc.)

13 million database
entries (5 GB)

The site had **3 times more**
visitors during the World
Cup competition



Load Distribution

An OVH load balancing (LB) IP was put in place at the top of the infrastructure to manage the HTTP requests of the two HAProxy dedicated servers. Considering the low CPU and RAM requirements, one might imagine that the load balancing of the cluster could be managed with a few virtual machines. In the case of a site with high traffic, a VPS's bandwidth limits (100 Mbps/VPS) would require employing a lot of virtual machines to handle such increases in traffic and this solution would prove to be uneconomical (due to the time spent having to manage the machines).

There are two reasons for the presence of load balancers at this stage: 1. to go further with the LB IP configuration and overcome the inability to route requests directly from the LB IP to OVH Public Cloud Instances*, which are used as the front-end, 2. to reinforce the three physical dedicated servers.

*Today, it's possible.

In this configuration, the LB IP performs basic dispatch requests (network level), while the HAProxy servers are able to sort more refined and targeted requests (applicative routing). For example, in certain cases, to manipulate response headers.

It must be noted that the LB IP plays a role in securing the infrastructure by only opening ports 80 (http) and 443 (https) thereby limiting the attack perimeter.

Additionally, the OVH.com load balancing service is based on two redundant physical load balancers (Cisco ACE), to provide a high availability service.

Front-end Scalability

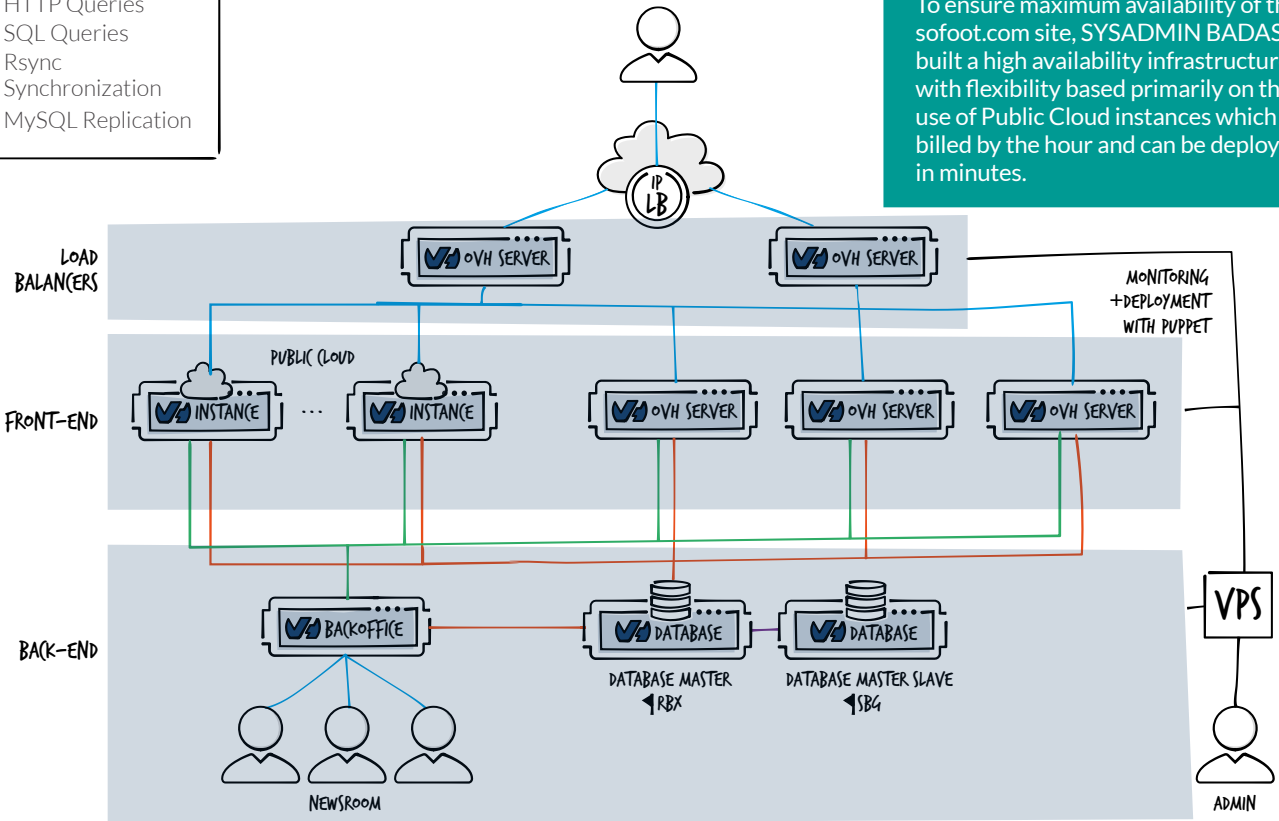
To absorb the expected traffic, a total of eight web servers were deployed: three physical OVH.com machines plus five Public Cloud Instances. Each of the web servers contains the sofoot.com website,

which is synchronized between the different members of the cluster via Rsync. The interest in the virtual Public Cloud machines is in the ability to add and remove instances as required. These resources are billed by the hour, allowing for temporary oversizing of the infrastructure to support increased traffic without incurring the costs of adding additional physical servers (rental for 1 week or 1 month minimum + possible installation fees). Furthermore, though the machines are virtual, Public Cloud instances offer high performance, which is a necessity in the case of the sofoot.com infrastructure. As previously explained, each of the front-end web servers includes the entire website (60GB of static content) and must be capable of offering users fluid navigation.



- HTTP Queries
- SQL Queries
- Rsync Synchronization
- MySQL Replication

To ensure maximum availability of the sofoot.com site, SYSADMIN BADASS built a high availability infrastructure with flexibility based primarily on the use of Public Cloud instances which are billed by the hour and can be deployed in minutes.





Database Replication

The site's massive database (13 million records in MySQL, 5GB) is hosted on a dedicated "master" server housed in a datacenter in Roubaix (RBX). This database is replicated to a second server configured as a "slave" and is able to take over in the event the "master" should fail.

The slave server is hosted in a datacenter situated in Strasbourg (SBG), assuring that the database is replicated in two incident zones which are isolated from one another. The two machines hosting the sofoot.com database take advantage of SSD disk technology, improving read/write performance. This is a judicious option that guarantees acceptable response times from such a large database.

Back office for Editing

Site editing (40 journalists worldwide) is made possible via the back office hosted on an OVH.com dedicated server. Authors manage their content (texts, photos, videos) in a secure environment (https connection and network restrictions) with all changes being propagated to the production environment database (web front end) via Rsync.

Automatic Deployment of Machines

A script permits the provisioning of Public Cloud instances directly from the OVH API. Provisioning is triggered manually or directly by the monitoring system when certain thresholds are reached. However, as a precaution, all VMs remained functional during the World Cup competition.

The application Puppet was chosen to automatically deploy the configuration of additional servers at the points of infrastructure congestion: HAProxy for load balancing and Public Cloud Instances at the webserver level. When a server is added to the cluster, Puppet automatically deploys the configuration setup onto a master server, automating the server's deployment.

Infrastructure Monitoring

The entire infrastructure is monitored with the assistance of Munin, installed on a VPS Cloud offering high availability. Munin generates a series of graphs from the information provided by the infrastructure servers: RAM usage, load average, disk consumption and network stats. Munin

functions on a server-client model. Munin-node, a small daemon, runs on each of the monitoring servers and provides the raw data as soon as it is requested by the application.

Backups

Backup Manager is used for regular backups, by means of the FTP protocol. The database as well as static content of the sofoot.com site are stored on the "Backup Storage" offered by OVH.com. Backup Storage is the backup space provided with each OVH.com server, by default 500 GB is included with each server and up to 10 TB of storage space is available as an option. Backup Storage can be accessed from any OVH.com server IP under the same customer account. This makes it possible for multiple servers to make use of the backup storage space. However, the rental of a backup server is being considered in order to centralize backups to one location ■



OVH advice

SYSADMIN BADASS has accomplished a real feat, having successfully built a highly available and scalable infrastructure, lowering costs while hosting an application that was not developed according to today's standards (cloud ready).

Possible optimization for this infrastructure would include the incorporation of the OVH CDN to store sofoot.com site's static content in order to reduce the volume of this data on its web servers, improving latency and, moreover, decreasing deployment time of any new server. However, this would require significant changes to the site, including the rewriting of all image links, as they would be hosted on a different subdomain.

Vincent Cassé,
Developer at OVH, specializing
in scalability issues.



Feedback

“The infrastructure handled the load from the first to last day of the World Cup and well after as this architecture was permanently adopted by sofoot.com.”

“The elasticity of the architecture that we imagined (Public Cloud Instances rented only for peak periods) was delivered at very a reasonable cost. Increases in traffic, while in production remained in the estimated limits with the maximum of 320,000 unique visitors/day, 2,300 requests/second and a tripling

of the load in less than 30 seconds (in general at the end of a match). During the two week period following the World Cup, sofoot.com established new records in site visits with peak traffic reaching 360,000 unique visitors/day, 3,000 requests/second and a tripling of the load which occurred in less than 20 seconds. This phenomenon was due to news about player trades and an influx of new visitors who became familiar with the site during the World Cup.



Record Visits

360,000 views/day

3,000 requests/sec.

3 times traffic
in less than 20 seconds

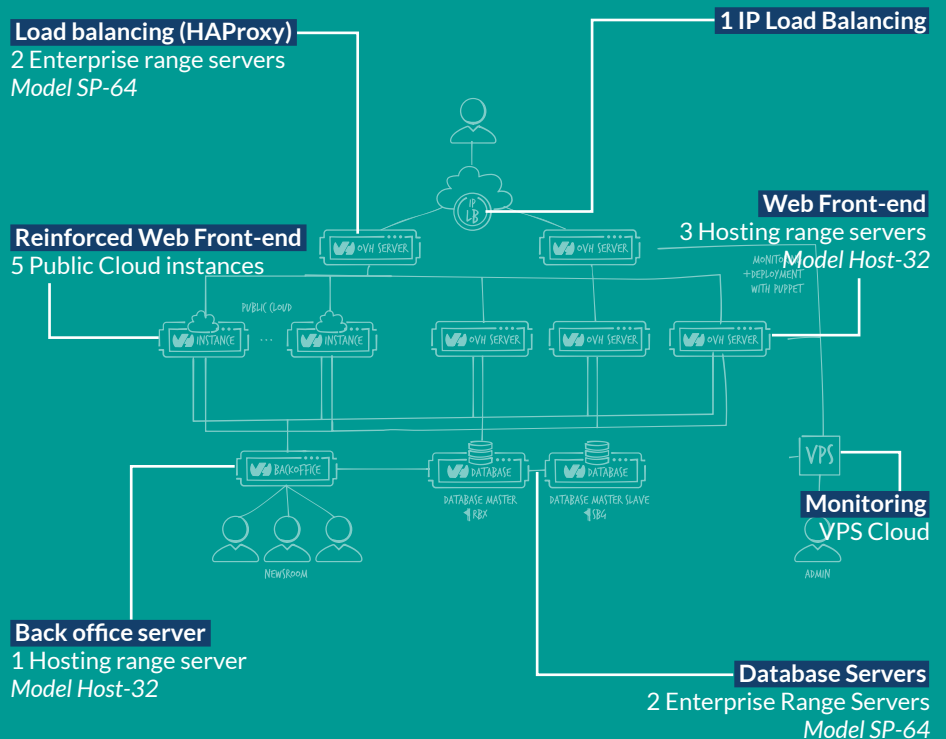
These peaks in traffic were fully amortized by the infrastructure in place, which is good news ahead of the Euro 2016 competition.

Today we are testing the automatic provisioning of additional virtual machines using the OVH API (based on thresholds in the monitoring tools), to benefit from a fully elastic and 100% self-administered infrastructure. Work is still required to simplify the infrastructure in a manner that reduces the time required to put a new machine in production from 20 to just 4 minutes.”



**Nicolas
CAPEYRON**
SYSMIN BADASS
Founder

Infrastructure Components





Training

At OVH, all new employees receive instruction through the “Training Center”. This center for internal training was established in August 2014 and is dedicated to providing knowledge. In the future, the goal is to offer, within the Training Center, certified training to OVH employees as well as customers and business partners.

#09

A Flexible Infrastructure for a Remote DVD Playing Service

vodkaster
MOVIES & TV SHOWS • BUY, WATCH, SELL

At the crossroads between VOD and DVD e-tailers, Vodkaster has been offering, since March 2014, a unique service to collect, dematerialize, remotely view and resale online its DVD collection. With more than 120,000 DVDs entrusted to it by its users -nearly 40% are for sale-, Vodkaster currently possesses the largest legal catalog of Over-The-Top (OTT) video. Today confronted with the challenge of storing large volumes of data and tomorrow with the mass distribution of videos, Vodkaster has built an infrastructure at OVH.com that is a combination of dedicated servers, Public Cloud (instances and storage on demand) and CDN. It is a robust and resilient infrastructure capable of handling increases in load which can easily scale in response to the success of the service that is supported by both private investors and public institutions.

Specifications



• Today

- optimize and control data storage costs.
- manage foreseen (end of the week) and unexpected increases in traffic and load (i.e. resulting from communication campaigns or mentions in the media).

• Tomorrow

- Broadcast large volumes of streaming video
- in DVD quality (with interactive menu, subtitles, bonus, etc.)
 - throughout the world
 - across all types of devices: PC, tablet, smartphone, TV set top box, and soon connected TVs.

Streaming Increased Traffic Scalability

Dedicated Servers vRack IP Load Balancing Public

Cloud Instances Public Cloud Object Storage CDN



Key Figures

1 million

visitors/month
on vodkaster.com

500,000

film and series
critiques

120,000

active user accounts

500,000

titles whitelisted
by users

3 million customer
evaluations

+ 120,000 DVDs

registered by users, dematerialized and viewable
online

+ 800 TB of data



Front-end Servers: Two Physical Machines and Public Cloud Instances to Reinforce and Manage Load Peaks

At the head of the infrastructure there is a load balancing IP in place evenly distributing users' requests between two EG-32 dedicated servers (Infrastructure range). These servers contain both web and PHP application layers, consisting of a social platform, an e-commerce platform and a platform for streaming with DVDs being viewed over HTTP using the VLC player.

These two servers handle load distribution. A Load Balancing IP, which tests server availability on a chosen port (in this case HTTP probe) before distributing requests, automatically manages the situation in the event that one server should go down by routing all traffic to the functioning server (load balancing fail-over).

The site is experiencing a steady growth in traffic along with, at times, sudden and major

increases in load. Some load increases are predictable while others are related to the nature of Vodkaster's business: on Friday, Saturday and Sunday evenings, these are periods that see a significant increase in the use of resources (users view their DVDs, buy new titles, etc.). Other load increases are more unpredictable or at least in terms of intensity. Load peaks are often a result of Vodkaster being mentioned in the media or follow viral communication campaigns launched by the company (for example: vodkaster.com/bienvenue-les-pirates).

To manage load peaks in an intelligent manner, that is to say without permanently increasing the size of the infrastructure, Vodkaster plans to complement its pool of two physical front-end servers with on demand Public Cloud instances. Instances will be deployed to counter against the

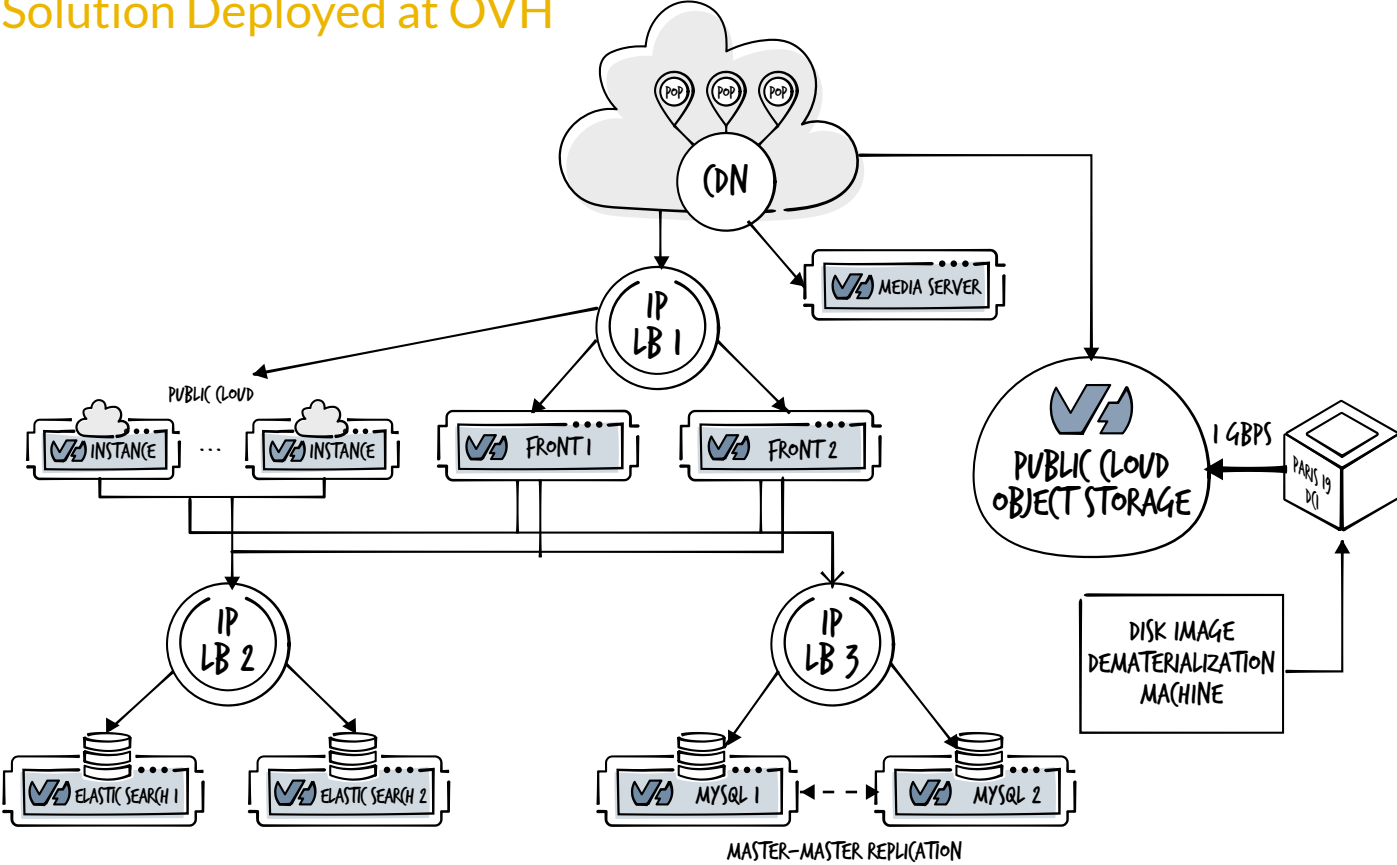
expected load increases during weekends or automatically for unexpected cases using a script triggered by monitoring, provisioning instances through OVH's API. Automatic provisioning, combined with a potential temporary reduction of intensive site features, will prevent the site from becoming unavailable due to a massive influx of simultaneous visitors. The configuration of instances deployed will be industrialized initially through customized scripts while later it is envisioned using a Docker image that provides the advantage of versioning.

Thanks to these automations and the possibility to add new back-ends to the Load Balancing IP via OVH's API, it is possible to deploy additional instances at the point of infrastructure congestion in just a few minutes. With billing by the hour, these cloud servers can be turned off once the demand

Supplementary instances can be quickly added via the OpenStack API, using the Nova client.

```
$ nova boot --flavor <flavor> -- image <image> MonInst-
tance1<flavor> : ID de flavor (ie: le gabarit d'instance)
<image> : ID d'image (ie: la distribution)
```

Solution Deployed at OVH



for resources has subsided. This will provide Vodkaster with significant savings in comparison to the costs associated with having to rent additional physical servers that include a monthly contract and are much of the time under-utilized.

Database: MySQL + Search Engine/ElasticSearch Cache

Two EG-34 Infrastructure range servers are used to host Vodkaster's MySQL database. A custom script performs the master-master replication of the two machines allowing both machines to handle requests from the front-end. Requests are dispatched by the second Load Balancing IP, configured with verification (probe) of availability of the servers via the OcO sensor. Information contained in the database is replicated on two other dedicated servers (SP-64 Enterprise range) in the form of a non-structured database (NoSQL) that takes advantage of Elasticsearch. This second database, synchronized with the MySQL database through the application layer, integrates Vodkaster's API (abstraction of objects abstraction of business data) allowing it to function like a cache of the main database and is primarily accessed by the front-end servers. If the requested information is not found or has not been updated in the

NoSQL database, requests are then sent to the MySQL database and the NoSQL database will be refreshed.

Storing Dematerialized Films in Public Cloud Object Storage

DVD disk images (bit by bit copy) are transferred to OVH.com's Public Cloud Object Storage then to Vodkaster's own server that has been installed (housing option) in OVH's P19 data center, located in Paris, and is connected directly to the OVH.com network through a 1 Gbps network connection. This storage solution was selected because of its simple billing (by GB stored and consumed outgoing bandwidth), unlimited storage capacity and the high availability of stored data (data is replicated three times across three server clusters that are capable of distributing their load). In comparison to the Gluster cluster initially implemented, which employed around thirty Kimsufi servers, the Public Cloud Object Storage solution allows Vodkaster to not have to worry about storage hardware issues (management of full disks, data restoration in the event of failure, etc.). To assure expatriate users optimal access to their dematerialized film collections, Vodkaster is currently considering putting a CDN solution in

place. However, due to the specific DVD format, this a real challenge if you wish to use the current offers on the CDN market.

Use of CDN WebStorage to Serve Static Content

Images, cover jackets, community user avatars... static content is managed by a media server, an EG-64 model from the Infrastructure range, which calculates the different image sizes required for the site's pages and pushes these files towards the OVH WebStorage CDN. The goal is to eliminate requests to the media server, whose resources are used primarily for quickly resizing media.

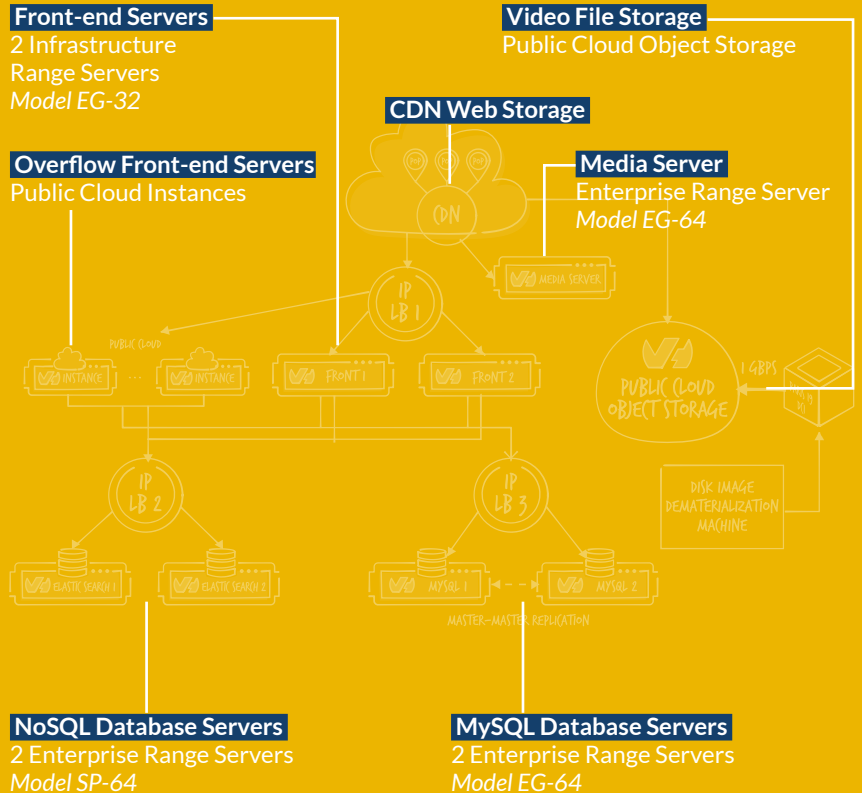
Project: Infrastructure Isolation within vRack

The main objective of Vodkaster's project is to isolate its infrastructure in order to provide maximum protection to the databases. Only the web servers will be exposed to the public network. This is a medium term project which justifies Vodkaster's choice in Infrastructure range servers due to their vRack eligibility. The feasibility of this project relies on the condition that two points on the OVH roadmap be completed. The first point is having the

possibility to add private IPs to isolated servers within the vRack, just like the backend Load Balancing IP, and the second point, having the possibility to connect Public Cloud Object Storage to vRack ■



Infrastructure Components





BHS, the Largest Data Center in the World

At the end of 2012, OVH put online its first data center overseas with a capacity to house, a record, 360,000 servers. “BHS”, situated in proximity of Montreal, Canada, is a natural choice for the company’s North American customers. Many European companies have also opted to host their servers in the “BHS” datacenter, assuring that the millions of customers, in the American market, have optimal access times to their applications.

#10



An Infrastructure with
High Availability Combining
Dedicated Cloud and
Dedicated Servers
through vRack

This online collaborative platform was launched in 2012 and as of today, has been adopted by more than 25,000 organizations. It aims to facilitate teamwork by offering an interface that incorporates all essential services for effective collaboration: synchronization and document sharing, (Wimi Disk), task scheduling, calendar sharing, instant chat, screen sharing and video and audio conferencing, (Wimi Air Time)... All associated services include management with appropriate access rights for professional use. To host this platform, the French start-up has built at OVH an infrastructure that is highly available, combining Dedicated Cloud and dedicated servers through the use of a vRack private network.

Requirements

- High availability of service
- Scalability: it is necessary that the infrastructure adapt accordingly to the number of connected users
- Data security: must be hosted in France
- Must provide users a secure connection (HTTPS) and an infrastructure for administrators (VPN two factor authentication)

Resistant to Increase in Traffic HA Storage Scalability Dedicated Cloud
Dedicated Servers vRack Load Balancing IP



Key information

78 servers:

8 physical + **70** VMs
spread out on OVH.com's Dedicated Cloud + an OpenStack cluster hosted locally

60,000 users worldwide

1,000 https requests/sec

4,000 SQL requests/sec

60 GB in databases

33 TB usable storage capacity



The high availability of the infrastructure is based on several elements. The first element is the choice of a native HA solution that Dedicated Cloud provides which includes redundancy of each component (hosts, power supply, and network connection). The second element is the division of the infrastructure by service. Each service is provided by a cluster of independent servers, guaranteeing resiliency. Within each cluster, a machine malfunction (or even multiple machines simultaneously) has no impact on service, functioning servers in the cluster temporarily absorb the additional load.

This configuration also guarantees the scalability of the infrastructure, having the ability to augment the capacity of the infrastructure by deploying additional machines at any points where bottlenecking may occur. This is an operation simplified by using Salt Cloud to automatically deploy and configure additional virtual machines.

Load Distribution with Two Load Balancing IPs

Users connect to their Wimi account via a web or mobile application and their connections are distributed between different web servers via the first load balancing (LB) IP. Download requests or uploads from the file synchronization client "Wimi Disk" or the web application are directed to specific web servers via a second load balancing IP. What is the advantage of the Load Balancing IP? In regards to routing hardware and its redundancy, this is managed and guaranteed by OVH.com. Concerning system administration, the configuration of adding or deleting IP addresses to machines in the cluster is minimal.

Front-end (Nginx + PHP-FPM)

The front-end servers are virtual machines created within the Dedicated Cloud which are administered with the VMware, vSphere hypervisor. The web servers (Nginx) encrypt the connection with an elevated level of security (HSTS + perfect forward secrecy) and render static data. To generate dynamic pages, FastCGI (PHP-FPM) is used for communication with the back-end servers. The purpose of this

separation is to be able to start quickly and efficiently new web or back-end servers, depending on the need.

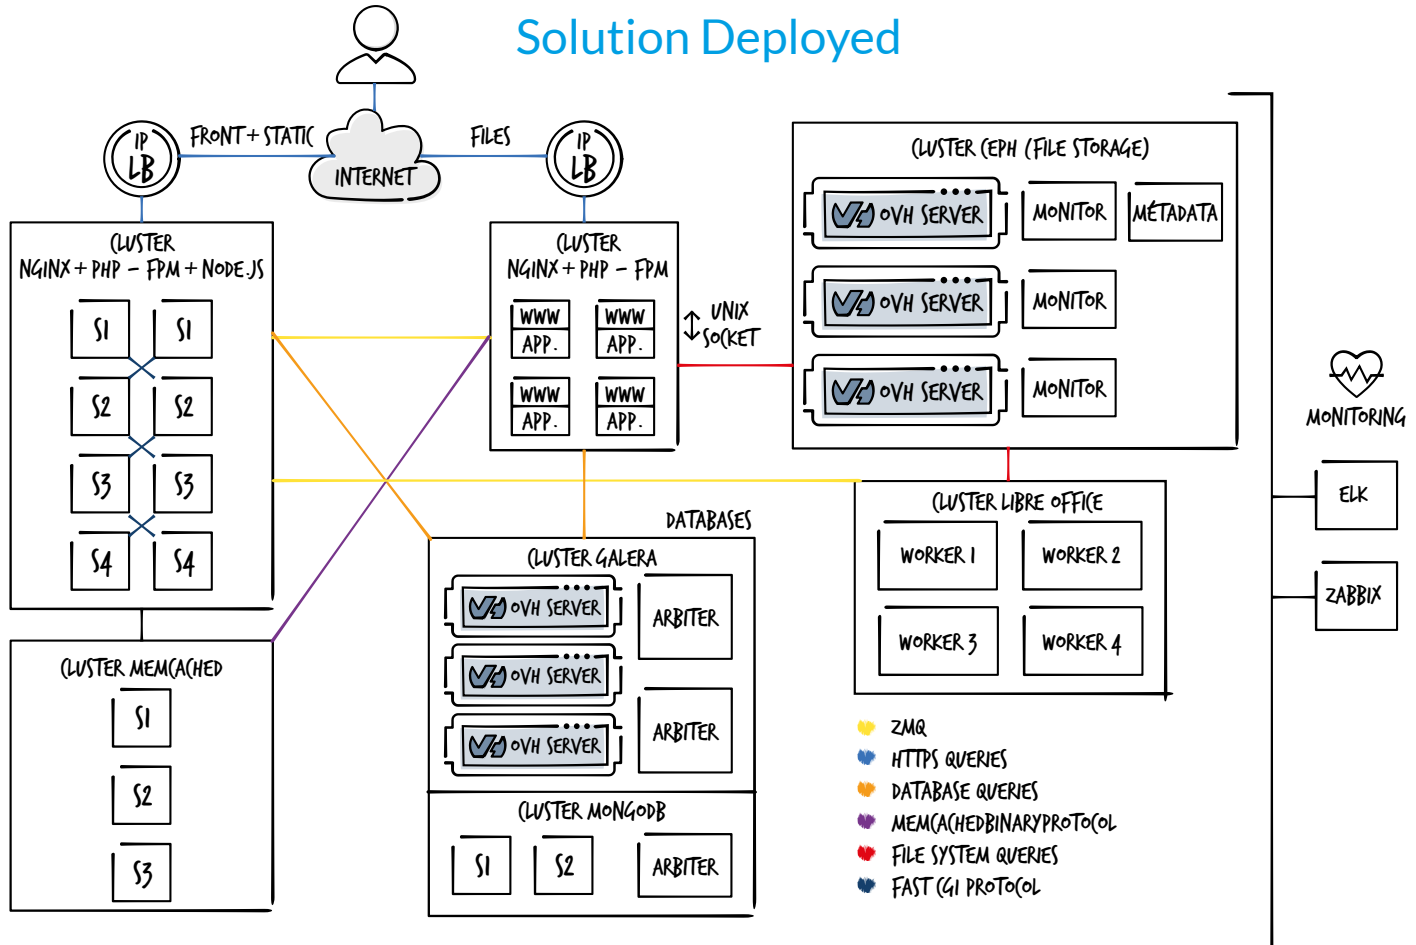
Back-end/Cache (PHP-FPM + Memcached)

The back-end servers, also hosted on VMs, are PHP-FPM and Memcached servers. Memcached is a distributed caching system amongst several machines. Its purpose is to reduce the number of queries on the database servers and PHP-FPM calculations thus reducing execution time. Only essential requests are sent to the database, pages recently served are cached. Regarding infrastructure evolution, the use of virtual-machines as back-end servers provides a double advantage. On one hand, upgrades can be made by hot-adding CPU and RAM (vertical scalability or scale up) and on the other hand, additional VMs can be added (horizontal scalability or scale out).

Databases (Galera cluster)

The principal database of the Wimi infrastructure contains 25 GB of data. A Galera database cluster was set up on the physical servers and interconnected by the vRack. This permits synchronized

Solution Deployed



replication between the three cluster machines, which are all active (master). The performance of the dedicated servers (infrastructure range servers) is an asset for database reactivity. Their connectivity (1 Gbps bandwidth within the vRack in the case of the EG-32 servers) reduces network latency generated by the Galera multi-master clustering.

Using Galera also serves in the quick provisioning of new database; horizontal scalability is always the goal. The two virtual servers hosted on Dedicated Cloud, play

the role as arbitrator, to reach a quorum of three functional nodes. This organization assures that the cluster remains functional even if up to two machines should fail simultaneously.

A second database was built under MongoDB, with a master server, a slave and an arbiter (three VM hosted on Dedicated Cloud). This database is used for the instant chat functionality, offered to Wimi users, and for the discussion wall. Its high availability and scalability is ensured natively by MongoDB.

File Storage (Ceph)

For storing files on Wimi Disk (and for which the last fifteen versions are kept), Ceph was the chosen solution. Hosted on several FS Max* servers interconnected via vRack, this storage system has been configured to replicate each file three times, on three different physical servers. A future goal of the project is to perform replication across three separate servers located in three different data centers. Ceph requires 'monitoring' and a metadata servers, which are hosted on Dedicated Cloud. The advantage of this distributed storage system is to offer high availability, scalability and maximum security of data.

To optimize the use of the FS-Max servers, the operating system was not installed on the server's disk. It's deployed at start-up via the vRack network using a PXE server (VM). This clever set-up allows for the addition of new servers and provides flexible configuration, that is similar to that of the cloud, while benefiting from all the performance of dedicated servers for storing data (tens of hard disks per machine).

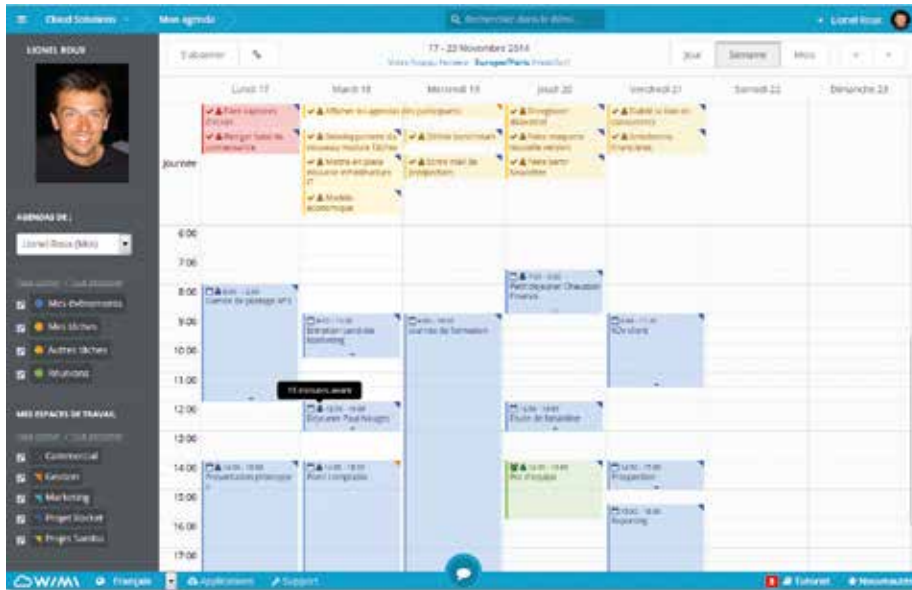
* FS- MAX servers have the ability to hot add (hotswap) SAS or SSD disks as needed, up to 36 TB/server.

"Peer to Peer" Communication Between Services

Initially, a "message queue" cluster, based on the RabbitMQ application, was employed, permitting asynchronous communication between servers but ultimately it provided a weak link between the different infrastructure clusters. With such a system, the message queue collected and stored the demands of the different servers, in a manner that they could continue working without waiting for the response from the servers which performed the required tasks. Today, RabbitMQ is no longer



Wimi offers access rights management adapted for professional use.



Shared calendar, an essential tool for easy collaboration.

used and has been replaced by ZeroMQ, which is much more flexible. ZeroMQ is a network library, which allows “peer to peer” communications between services. Communication is made directly from machine to machine, eliminating the need of the server acting as the “middleman”. The goal is to have services run in real time, with the websocket data push based

on ZeroMQ (for real time communication between the different services of the infrastructure) and Node.js for notification push to web clients. For example: the document preview is stored/shared on Wimi. Previously the preview was simplistic: all Office documents that arrived on the platform were transformed to .pdf format. During the day, there was a fairly

consistent queue, and the user never knew how long it would take before the preview would become available. As a consequence, many previews were made “for nothing” as they were never consulted by users. With development based on PHP – ZeroMQ – Node.js, now it’s the LibreOffice servers which manage client requests. When a customer requests a preview, if it does not already exist (in which case it does exist it is viewed instantaneously) the PHP node that received the request sends a ZeroMQ message to a LibreOffice server. Each of the servers can calculate a preview independently from the others. In this manner, when several users simultaneous make different preview requests, the load is distributed between servers and wait time is not multiplied. On the client side, an animated logo is displayed while the end user is waiting. Once the preview has been generated, a ZeroMQ message is sent to one of the front end node.js machines that manages sending messages via the websocket tunnel and displaying the preview for the client. If several clients request the same preview, at the same time, only one message is sent to one server but all clients go into waiting mode and then receive a real time notification once the preview is available. This results in the average wait time being a second between

the client clicking and the displaying of the preview, a major improvement compared to the several hour wait previously experienced by users!

Auto-Scalability with Salt Cloud

Today, Wimi is a tool that is used regularly and with predictable load increases, the infrastructure is able to scale accordingly. Mechanisms for provisioning, configuration and automatic deployment of additional resources have been put in place using Salt Cloud, a tool developed by SaltStack.

Through the vSphere API, Salt Cloud allows for the management of the complete life cycle of a virtual machine, from its creation on Dedicated Cloud to its deletion (an interesting point when the machines in question are created on an additional Dedicated Cloud host, billed by the hour). Virtual machines are provisioned automatically then copied and deployed from a common drive to all machines within the infrastructure. A template defines a VM's characteristics: number of CPU cores, RAM, network configuration... Next, the machine is named and configured by Salt Cloud, then added to DNS. All of this is accomplished in around 15 minutes and without any need for human intervention with the most amount of time being spent

making a copying of the template.

Salt Cloud, compared to Puppet, which Wimi previously experimented with, provides two major advantages. The first advantage is its compatibility with VMware's vSphere API, allowing the use of a single API to manage the entire infrastructure, where Puppet required custom scripting.

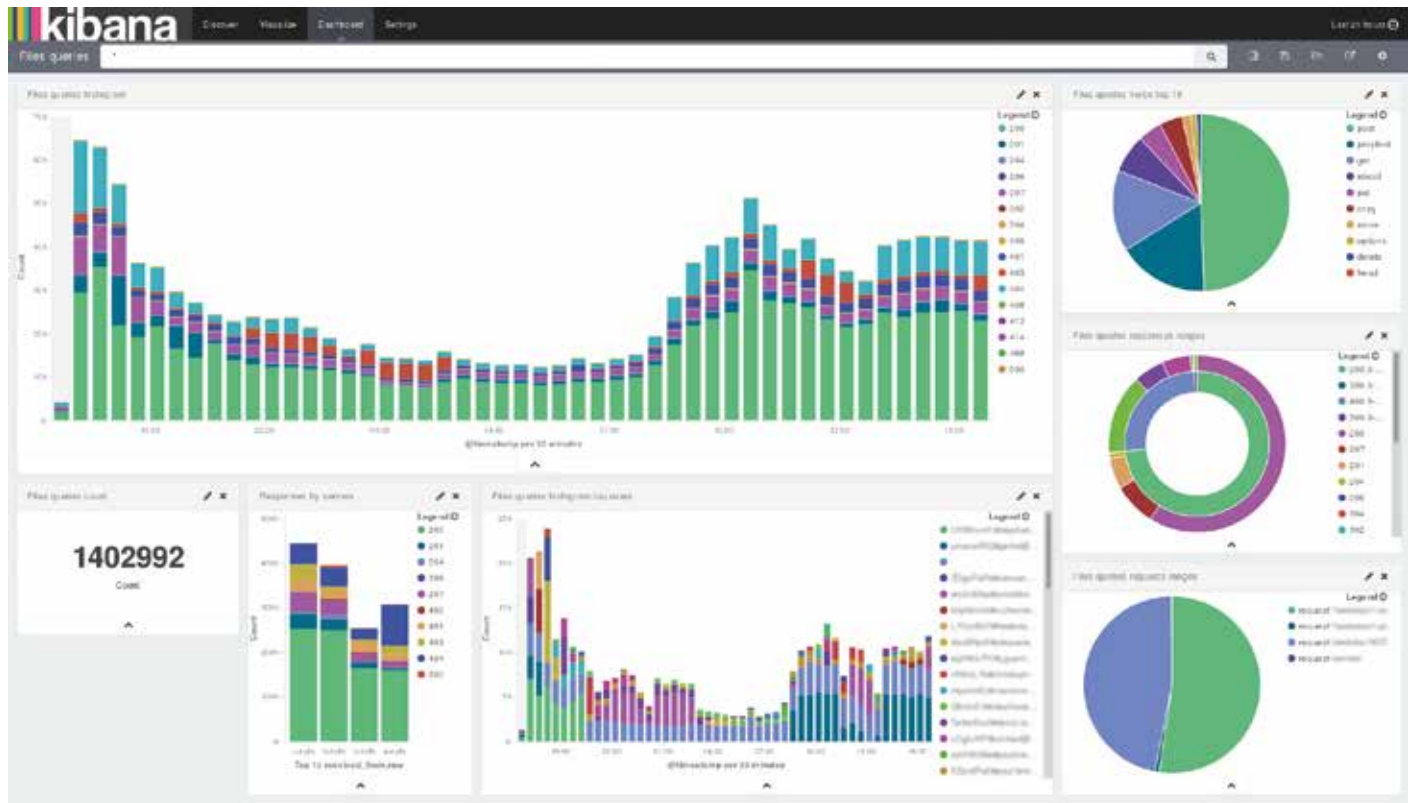
The second advantage is that Puppet is a process which is called by the client server at regular intervals (configurable), to perform a compliance audit in relation to the configuration of a master server. With Salt Cloud, the "minions" (the agents) installed on the client servers have a permanent connection with the master server and the master server can "push" the changes. Today, the addition of supplementary resources is triggered manually. To become fully automated, which Wimi will require in the future, the only thing that remains is to determine, service by service, the critical thresholds for scalability and to ensure that exceeded thresholds trigger the action of Salt Cloud. Salt Reactor is used to complement Salt and to react to the defined events which come from Salt or external sources, especially to automate the possible changes to be performed on the machines indirectly impacted by the actions of Salt Cloud. For example: at the end of deployment for a new php node, Salt

Reactor automatically calls the necessary processes to add the new node to the configuration of all concerned Nginx nodes and reloads services.

Monitoring and Performance Analysis

A virtual server running MongoDB collects and analyzes real time technical statistics (calculation time of the pages, request execution time...) of each of the user's services running on the infrastructure. An ELK stack (ElasticSearch, Logstash and Kibana) is used to retrieve, organize, store and view the logs in an effective manner. Kibana displays the logs in a user-friendly graphical dashboard to identify errors as quickly as possible. The ELK stack is one of today's most effective monitoring tools. It combines the power of Big Data (Logstash and ElasticSearch) with the customizable web interface of Kibana. In return, the deployment and maintenance of such a stack is a time saving investment.

Moreover, an architecture for supervision, based on Zabbix, has been deployed in the Wimi offices in Paris, at OVH and with another hosting provider. The open source solution, Zabbix provides status monitoring of services, system and network. Conveniently, through the use of a template system designed for each type



Kibana view after requesting files, server by server, showing different response codes and request types.

of supervised server, it has the ability to automatically display adapted probes for a machine's profile.

Backup

For backups, two dedicated servers were chosen from the Storage range (FS-72T). On the first server, an incremental backup of the databases is performed hourly via Rsync. And concerning the second server, a fourth copy of the files contained on the Ceph cluster are created, also via Rsync. To secure sensitive data on these machines, the servers are only accessible through backup applications (SSH access is not even possible). The use of ZFS disk partitions preserve different versions of the files (snapshots). It is possible to go back in time in the event that data is accidentally erased. The two backup servers are hosted a Strasbourg (SBG), a remote datacenter from in which the rest of the infrastructure can be found (Roubaix, RBX). The databases are stored in the Ceph cluster and the data is already redundant within the infrastructure. These backups are important, but would be only useful in the event of a major disaster: a plane crash into the datacenter, meteorite, tsunami, alien invasion...

Internal Network Infrastructure

It is difficult to detail the internal network in place in the Wimi infrastructure. There are some things to take into consideration. First, the vRack was put in place to allow Wimi to take advantage of the interconnections of the physical machines in the database and Ceph clusters, providing maximum connectivity between servers and the

private network, thus assuring perfect security. Thanks to the Cisco Nexus 1000V virtual router included in the Dedicated Cloud offer, Wimi successfully deploys a granular network, in which each service is insulated in a dedicated VLAN (which is a guarantee of security) ■



Wimi includes "AirTime": providing the possibility to launch audio/video calls and to share a screen directly through a browser.



Our Experience

"In 2012, when we were looking for a hosting provider to build our Wimi infrastructure, there were few choices available to us. Only the major players were offering virtualized infrastructure solutions combining flexibility and performance. But these players were American and we wished to reassure our users about the confidentiality of their data. We discovered OVH.com and their Dedicated Cloud offer during this time. In addition to the delegation of hardware management, the advantage of Dedicated Cloud lies in the fact that resources are dedicated and hosts and datastores can be added as needed. The arrival of the vRack was the decisive factor. It allowed us to combine the best of both worlds, the cloud for flexibility and dedicated servers for raw power. This is especially useful for hosting high volume databases. Without this private network, it would have been impossible to implement such technical choices such as the Galera cluster of MariaDB databases or the Ceph distributed storage. The interconnection



Thomas Lemarchand
Wimi System
Information
Manager.

of the machines would have been possible via their public IP, but that would have affected the overall performance of the infrastructure and raised security concerns. Now our ambition is to develop Wimi internationally and particularly in North America. This will require the deployment of a new infrastructure on the other side of the Atlantic, a project facilitated in part by the recent use of Salt Cloud and by the OVH.com's presence in Canada, with its Beauharnois datacenter (BHS) located near Montreal.

Infrastructure Components

2 Load Balancing IPs

Dedicated Cloud

4 hosts L

Datastores 2 x 300GB + 1.3TB

vRack

Service included with Dedicated Cloud and Infrastructure range servers

Databases (Galera cluster)

3 Infrastructure range servers
Model EG-32

File storage under Ceph

3 Storage range servers
FS-MAX (Pay as you grow)

Backup

2 Storage range servers
Model FS-72T



4 Tbps

To assure its customers maximum throughput, quality high speed bandwidth and the lowest latency times, OVH has made the choice to invest in deploying its own worldwide fiber optic network. Today it has a capacity of 4Tbps to the global Internet

#11

OVH.com Servers to Service the Connected Farm



Medria is a company that develops services for monitoring livestock. Today, there are up to 300,000 animals throughout Europe which are equipped with sensors fabricated by the company from Breton, France. These sensors make it possible for thousands of farmers to be alerted by SMS upon arrival of reproduction periods of their cows, imminent calving, dietary issues, or even for abnormal rises in temperature. To collect the 700 million monthly metrics and perform real-time analysis, Medria has built an infrastructure made of 18 dedicated servers at OVH.com.



Specifications

- High availability of the application (24/7); some of the services are critical for users (ex. calving alert)
- Collect, store and secure a large volume of data (long term conservation of strategic data for agribusiness research)
- Processing power: data is analyzed in real-time, making it possible to send SMS alerts when an event arises
- Infrastructure scalability: capability to retain even more data in the future without having to rethink the entire infrastructure

IoT (Connected Devices) High Availability Dedicated Servers vRack(Private Network)



Key Information

Between **100,000** and **300,000** animals are connected*

500,000 metrics are registered every 30 minutes, 700 million per month.



+ 160 SMS SMS sent every hour, 120,000 per month.

12 TB of data stored in database

18 OVH.com servers

* (depending on the period)



Medria designed an infrastructure capable of collecting a large volume of data (700 million measurements taken each month) and real-time analysis so that farmers can be alerted in case of an event (the start of the reproduction period, calving, dietary problems, rise in temperature). To accomplish this, Medria has put in place four levels of machines.

The first level is comprised of a server which manages routing, assuring that the data sent by the sensors (via a box installed on-site at each farm) is stored on the server containing the data base of the corresponding farm.

The second level is constituted of the machines that host the several thousand MySQL data bases, which total 12 TB in volume.

The third level, reserved for processing, is composed of servers which run proprietary software written in C/C++.

These processing servers have been virtualized using Xen 4.1, bringing with them the

advantages of flexibility in management and parallelization of calculations. The average CPU load of the processing servers is around 60%!

The fourth and last level of the infrastructure contains the alert server, which runs on proprietary software written in Java and manages the sending of SMS alerts via the API of Medria's chosen service provider. In this case, OVH Telecom is used as the backup operator.

To complete things, with the help of Nagios linked to a web interface based on the open source application Thruck Monitoring, a virtual server monitors the entire infrastructure. In particular, monitoring makes it possible to detect machine failures and the transfer of impacted services to a backup server. Backup servers are included at each level of infrastructure and are synchronized with the machines that are most susceptible. Dividing the infrastructure into functional levels ensures scalability. Additional machines can be added at the points of congestion of each level (in particular, data base servers and the servers used for processing data).

Isolation of Servers with vRack

Thanks to vRack, the private network offered by OVH.com, the infrastructure is almost

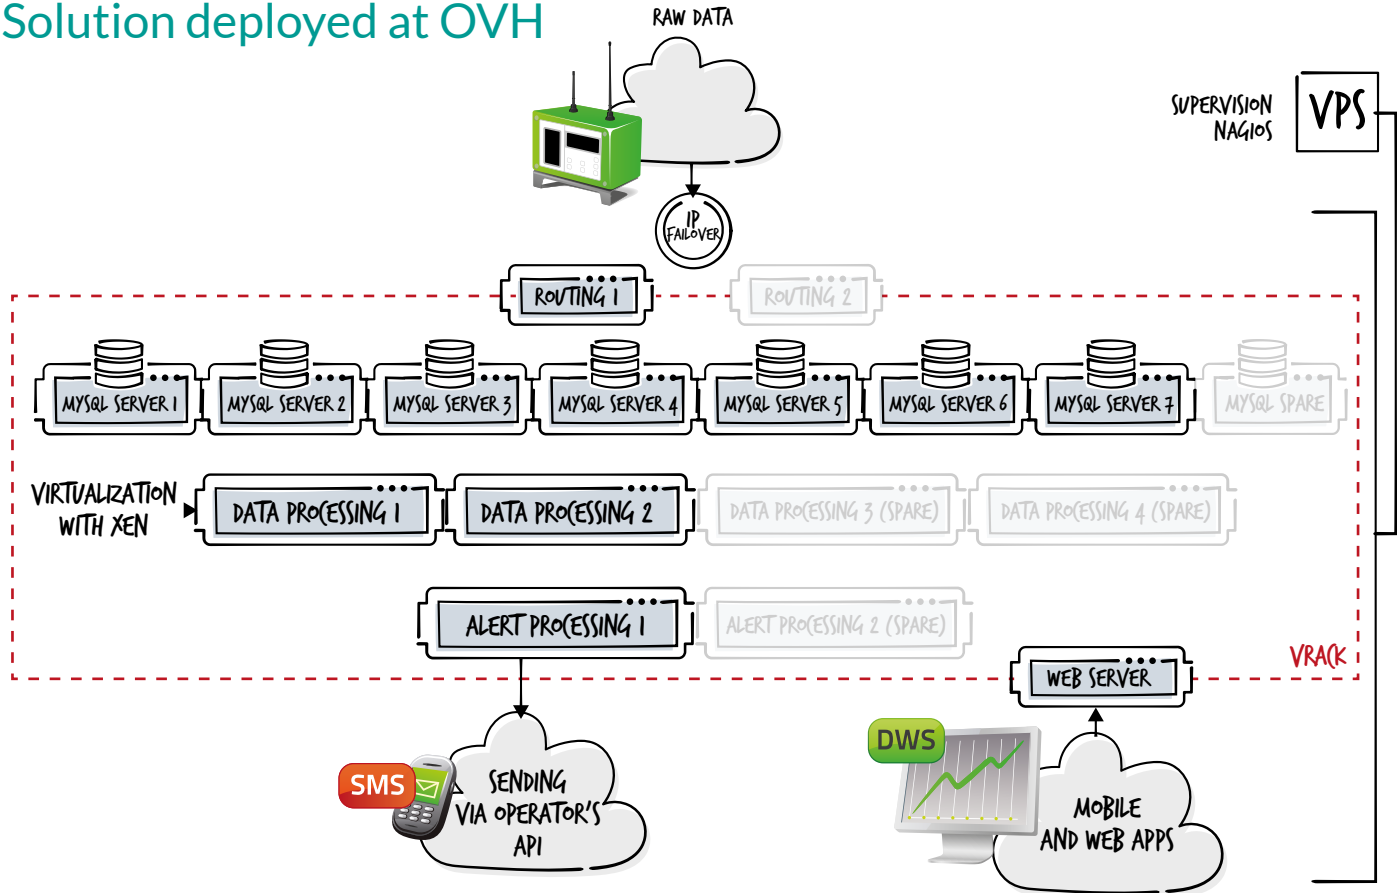
entirely isolated from the public network.

Only the routing and web servers (hosting the web platform and mobile applications and used for consultation of data) have access to the Internet.

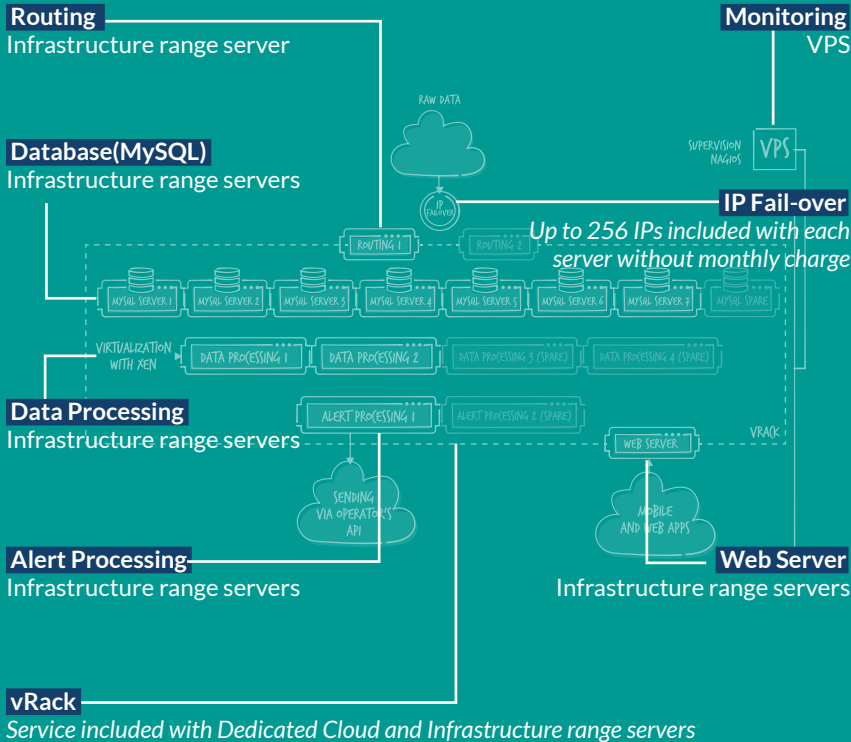
This isolation has three points of interest: 1. increases the resistance of the infrastructure to attacks, (requiring that attacks be more complex to implement), 2. accelerates communication between the servers and 3. prevents data from being transferred across the public network ■



Solution deployed at OVH



Infrastructure Components



OVH Advice

Projects linked to the IoT (Internet of Things) have multiplied in recent years and this is just the beginning. It is predicted by 2018, each person will own at least seven connected devices and this does not take into account any of the connected livestock, as demonstrated by Medria's business!

Of course, among the 50 billion connected devices, planned to be in use by 2020, there will be connected terminals (smart phones, tablets, e-readers...). The majority of the devices will actually be machine to machine related with sensors sending data to servers for analysis, producing alerts or even statistics and predictions. The volume of connected "things" is going to be a challenge at the network level. The shortage of IPv4 addresses, which is accelerating the adoption of IPv6, is only one aspect of the issues to come. To prevent network saturation, it will be necessary to increase the capacity of the fibers and cables required to transmit this enormous volume of new data. There

exist yet another challenge for service providers like OVH, to provide users with offers and infrastructures that are adapted to meet the needs of the IoT.

Medria notes that daily management of a server farm requires a certain amount of time be devoted to administration and requires specific skills. The IoT already calls for a large variety of skills, from design to data visualization by embedded electronics... In general, those who fail are those who have tried to do everything by themselves. This inspired OVH to quickly provide a SaaS platform dedicated to the IoT with a PaaS option for companies wanting to integrate their own business layer. This is a highly available turnkey platform, with data being pushed via the OVH API and automatically replicated. Platform billing will be

based on use. By nature the IoT is an event driven activity, generating spikes in traffic and load along with periods of machines sitting idle. The pooling of resources, as offered by a service provider the size of OVH, stabilizes costs and dramatically reduces the price for the end user. It will no longer be necessary to size the infrastructure according to maximum load peaks: the platform provides an infinitely scalable backend with hardware and software managed by OVH. For a company like Medria, the use of an IoT platform reduces the budget allocated to the infrastructure while freeing up human resources for missions that provide greater added value, such as development of algorithms. For any company about to launch, the use of a SaaS platform would strongly reduce time to market. Finally, even if the pooling of resources does not mean

the absence of isolation between users, we offer this same platform in an entirely dedicated mode, for projects that operate in a stricter and more regulated environment, as is the case for data linked to medical data.




Go to RunAbove

to beta test the platform



Sylvain Wallez,
Project Manager
OVH IoT Platform



More than 10 Tons of Materials

This is the approximate weight of the components that OVH receives each month to assemble its servers. Each server requires 15 minutes to build with an additional 25 minutes spent in the lab to perform quality control tests.

#12

A Highly Available, Secure and Load Resistant Infrastructure for Hosting a Virtual Key Application

“Keyzee” is a virtual key application that allows anyone to rent or share a vehicle. A box is connected to a vehicle’s OBD bus and communicates over a GSM network (also Bluetooth or NFC). Using a smartphone application, users are able to reserve, open, start or stop, and return a vehicle. This innovation not only interests a lot of businesses and car rental companies, for the optimization of fleet management, but also many individuals wishing to share their car using online car sharing applications. When you know that on average a car sits unused 92% of the time, you soon realize the potential of the project!



Specifications

- The service is mission critical, therefore the infrastructure must be available as close to 100% as possible.
- The infrastructure must be scalable and load resistant.
- Infrastructure security is a critical point, considering the possible consequences of a malicious hack.
- The infrastructure must be global for international users.
- It should be possible to use the data generated by the service at a later time (connection to a Big Data cluster).

High Availability International Users Multi-Datcenter Infrastructure
DRP – Security Dedicated Cloud Dedicated Servers



The Project in Numbers

In existence for **2 years**,
the application manages
300,000 vehicles.

**An award
winning
innovation!**

Fleet Innovation Award
2013 for Fleet Suppliers
in Belgium

Mobility Innovation Award
2014 in Luxembourg

Key Information



Keyzee is a project developed by **OTAKeys**, a joint venture between the multinational Continental and D'leteren, a Belgian importer of the brands Volkswagen, Audi, Seat, Škoda, Bentley, Lamborghini, Bugatti, Porsche, and Yamaha



OTAkeys is a client of **Vox Teno**, whose business revolves around three activities: the development of custom web applications, offshoring/nearshoring development projects and outsourcing services tailored to their customer's needs – with a majority of them being hosted at OVH.com. The company's 100 employees are spread throughout Belgium, Eastern Europe and Asia.



The infrastructure in production is based on Dedicated Cloud coupled with dedicated servers hosting the database and a NAS-HA containing files. The n-tiers architecture deployed assures both platform availability and scalability.

Infrastructure Security and Load Balancing

At the head of the bridge sits the VM on which the pfSense firewall runs, filtering all incoming and outgoing platform traffic. The incorporation of an intrusion detection system (IDS) within the platform completes the security plan. A level consisting of Nginx reverse proxies and HAProxy load balancers is in place to contribute to security and service availability. Given their position as an interface between the front and back-ends, the reverse proxies can certainly play a role in platform security by redirecting traffic. They are able to block certain types of web and application requests. Meanwhile, HAProxy has the function of evenly distributing requests to

machines of the next level - with the capability to stop distribution to any machine no longer responding.

Application Layer

Situated at a lower level are two types of virtual machines: the first type, a dedicated cluster that manages the virtual keys (Java application made available by Tomcat) and the other a cluster of web servers (Apache + PHP) which allow users to perform their actions within the mobile application.

Databases

A Percona cluster hosts the databases in a distributed manner. This cluster is attached to a highly available system that can tolerate the loss of a node without impacting the platform. Because of their superior performance, dedicated servers were chosen to host the databases rather than virtual machines. These servers are interconnected and linked to Dedicated Cloud via vRack (OVH's private network solution), guaranteeing security and rapid data transfers.

The Choice of NAS HA for Files

A storage server on a highly available network (NAS HA) is used to server files to the web

servers and to eliminate having multiple instances of the same content locally (with the obligation to propagate any changes to each machine). Centralization of files makes their exploitation much easier. In addition, this shared storage space is managed by OVH with the guarantee of a level of very high availability.

Disaster Recovery Plan (DRP)

A similar platform architecture to the one described (but smaller in size) has been deployed in the RBX datacenter, permitting service to function in degraded mode in the case of failure of the primary infrastructure which is physically located in Roubaix, France. This second infrastructure, which serves as the principal element of the disaster recovery plan or DRP, is synchronized with the main infrastructure via vRack, in order to be able to take over in the least amount of time.

Project: Globalization of the Infrastructure and Big Data

In the mid-term, the project is to have several similar infrastructures put in place in different areas around the globe (ideally an implantation in Europe, North America and Asia). With the managed fleets being separated, these infrastructures will function in an independent manner, providing users with optimal access

times following the anycast principle (users connect to the servers closest to them). The final piece of the project, the Big Data cluster, will in turn connect the entire infrastructure allowing for data collection. With data being returned from all the vehicles managed throughout the world, data will be centralized and analyzed with the goal of optimizing and enriching the service (data mining).



© Photograph OTAkeys

Infrastructure Components

Dedicated Cloud

Percona database cluster

Infrastructure range servers

vRack

Service included with Infrastructure range servers

NAS-HA





Tell Us About Your Projects!

Do you have an original project or an innovative use of OVH.com's services?

Provide us with a brief explanation about yourself and what you are doing. Prepare a schema of your infrastructure and contact us at the following link.

www.ovh.com/support/usercase/contact

Index

By Activity

Local Community

42-49

E-commerce

78-83

Application vendors and business software

14-21, 52-56, 86-95, 106-110

Education

60-65

Industry

24-27

Media

30-38, 68-75

Connected Objects (IoT)

98-103

Photography and Video

6-11

Online video

78-83

By Services Used

Dedicated Cloud

30-38, 42-49, 52-56, 60-65, 86-95,
106-110

Dedicated Connect

42-49

Web Hosting

6-11

IP Load Balancing

6-11, 14-21, 68-75, 78-83, 86-95

Public Cloud (Instances)

6-11, 68-75, 78-83

Public Cloud (Object Storage)

6-11, 30-38, 78-83

Dedicated Servers

14-21, 24-27, 30-38, 52-56, 60-65,
68-75, 78-83, 86-95, 98-103, 106-110

VPS (Virtual Private Server)

6-11, 14-21, 24-27, 68-75, 98-103

vRack (private network)

14-21, 24-27, 30-38, 42-49, 52-56,
86-95, 98-103, 106-110

By Issue

Virtual Desktop Infrastructure (VDI)

42-49

Total IT Outsourcing

42-49

High Availability (HA)

6-11, 14-21, 24-27, 30-38, 42-49, 52-56,
68-75, 78-83, 86-95, 98-103, 106-110

Monitoring

14-21, 24-27, 30-38, 68-75, 86-95,
98-103

Scalability

6-11, 30-38, 68-75, 78-83, 86-95,
98-103, 106-110

Storing Large Volumes of Files

6-11, 42-49, 52-56, 78-83, 86-95, 98-103

Backup Strategy and DRP (Disaster Recovery Plan)

14-21, 30-38, 42-49, 52-56, 60-65,
68-75, 86-95, 106-110

Video Protection

42-49



[linkedin.com
/company/ovh](https://www.linkedin.com/company/ovh)



[twitter.com
/ovh](https://twitter.com/ovh)



[youtube.com
/ovh](https://www.youtube.com/ovh)



[plus.google.com
/+OvhGroup](https://plus.google.com/+OvhGroup)



OVH.com