

Migrating to new Hosting provider

Hetzer due diligence

Question: Does it make sense for us to move to Hetzner? What would a setup of ours look like there?

General questions:

- The FLEXI pack (13€/month) is needed for any sort of customization of a machine. This is quite expensive. Can we avoid this for most of our machines? Private machines only need one connection and if we rent a switch, can we skip FLEXI pack?
- In case of hardware failure, how long is the delay to get a replacement machine ?
- What about the replacement of hard disks in RAID arrays ?

Machines for quote:

Questions for Hetzer

- Need room to grow
- Need private switch
- Each server connected to private switch

- 2 gateways: EX-40 SSD
- 2 database servers: PX-91-SSD + 2 additional SSD drives + 1 RAID controller
- 10 app server: PX-61

OVH due diligence

[thoughts on OVH go here]

Online.net

Online.net was suggested as a possible provider.

Services

Services that need to be moved first:

1. Gateways
2. MusicBrainz app servers
3. Postgres & new replicas
4. Redis/memcached
5. FTP
6. CAA (needs the same db as mb)
7. Stats (not really moved, new instance)

Services that should be moved next:

1. Discourse -> Google Cloud
2. Wiki
3. CI (Google Cloud may be too slow, MB build takes 19mn on hobbes, self-contained)

Services that can move later:

1. CB
2. Jira -> Google Cloud
- 3.

Hardware

RAID Issues

Software raid is needed for some machines, **but not all**.

Software: App servers, SOLR based search servers

Hardware: Databases, gateways <- not even sure, they mostly write logs, read almost nothing, software RAID + SSD and moving old logs to an archive server would be better.

Hardware RAID 4 ports card cost 15€ per month ... that's 180 € / year / server -> x 16 (minimum number of servers) = 2880 € / year if all servers use such card

If we want to reduce number of cards, they are mandatory are only database servers, ftp server, we can go software for other servers. = 4 or 5 * 180 € / year = 900 € / year

NOTE: hardware RAID cards require FLEXI pack

(https://www.hetzner.de/fr/hosting/produkte_rootserver/raid-controller)

Redundant power supplies ?

Not sure this is something we need to care about, especially with OVH vRack

Machine types

Of the different

App servers:

MBS web servers don't need RAID, and HDD is enough, fast CPU, a lot of cores, and decent RAM

Gateways:

SSD is a must if they log to disk, else HDD, RAID to be safe
Very fast CPU and network cards, decent RAM

Database master: Powerful enough to handle lots of writes.

Database read-only slave: Most powerful machine to handle lots of reads. Any read-only machine should be able to take over for the master. So, read-only should be more powerful than master.

SOLR servers:

Both db servers need a lot of RAM, and big and fast SSD drives with RAID, decent CPU with many cores, fast network cards

Catch-all VM machine:

KodeStar suggested <http://www.proxmox.com/en/>

Memcached: Do we need separate memcached machines? Many machines have more RAM than we need, so we can setup memcached instances for various projects on things like the gateways.

Influxdb / stats : machine has to be fast on write, SSD is a must

Zas propose to use 2 redis servers, getting rid of memcached (at least for mbs), and set up something like described at [How to setup HAProxy and Redis Sentinel for automatic failover between Redis Master and Slave servers – All Help & Support](#)

Since redis is mostly a RAM user, and influxdb mostly a CPU/disk user we may use 2 redundant machines with:

- HAProxy + Redis Sentinel, one master, one slave
- Influx db + influx relay

Web servers / gateways: they don't need much disk space (apart for logs, but we don't need to store a lot of them), CPU power + a lot of RAM + SSD

[what else?]

We need to find a way to centralize logs, it could be done with a SSD+HDD machine, using remote syslog. Though daemontools logs cannot easily be made remote.

Main advantage is that most machines are currently using disk space only for logs (web servers, gateways), so it would allow us to get low cost non-RAID small disks for those instead (and replace machines on failure)

About number of machines, we will go for almost the same setup as we have now, with slight changes:

- 2 gateways (HA using corosync + pacemaker)
- 6 web servers (5 + 1)
- 2 db servers (master + slave)
- 2 redis & stats servers (master + slave)
- 3 catchall VM servers (3 for proxmox HA cluster ?)
- 2 search servers (perhaps 3 if not enough)
- 1 log collector (?)

Rack questions

How much space do we want to rent? Renting space, safeguards future growth. We'll need to rent a switch.

OVH vRack seems to solve this point.

Bandwidth questions

Each machine comes with some amount of bandwidth. For us, only two machines are ever going to use this bandwidth. How does that work?

Backup questions

Each machine comes with some amount of backup. Some machines need very little backup, some others quite a lot. How do we balance this?

We need to rethink our backup system, Wiley is old and dying -> save to cloud ?

Security questions:

What about firewall?

Our current gateways are providing a complex firewall, based on linux iptables.

If we go to another architecture, firewalling would have to be redesigned.

Current load:

Do we have the current load statistics somewhere (RAM,CPU, disk I/O and consumed HDD space)?

<http://stats.metabrainz.org/> has graphs for almost everything.

Various

- Be sure that all servers have separate /tmp, /var/log partitions. This has to be done on initial setup.
- IP multicast is required for corosync (proxmox)
- Physical machines on the same physical network vs vRack (does corosync work with vRack, what about performance?)