

unipiBox owncloud @ University of Pisa



Simone Spinelli
simone.spinelli@unipi.it

University of Pisa

- Mid-big sized italian university:
 - 70K students
 - 6K employees
 - Hosts a PoP of GARR network
 - Surrounded by other research/educational institution
- Not a campus: offices are all around the city and connected with our own fiber optical network
- IaaS platform based on Openstack and CEPH already in production


Infrastructures@unipi

Storage:

- ~2PB CEPH Cluster:
 - 12 nodes (300TB) in production
 - 12 nodes (2PB) almost ready
 - 3 copies (1 per site)

Computing:

- Openstack infrastructure already in production
- Legacy Xen dom0 for centralized services



Sync&Share @unipi

For whom?

- For administrative offices
- For researchers
- For professors/students

For what?

- Substitution of small SMB shares
- Education (Moodle integration?)
- Sharing with others communities

Why not public cloud?

- Control on data
- Control on policies
- Integration with other in-house services
- Federations
- We already have the infrastructure (a part of it)

Sync&Share @unipi

Requirements:

- Scalability & Resiliency
- Quotas:
 - 10GB for Employees
 - 3GB for students
- LDAP authentication (SSO where/if possible)
- Leverage on the existent infrastructure

Previous experiences on Owncloud:

We have 2 small installation for initial tests already up&running, but:

- Local storage as backend
- Single machine installation
- ~100 users

owncloud@unipi

Proposed architecture:

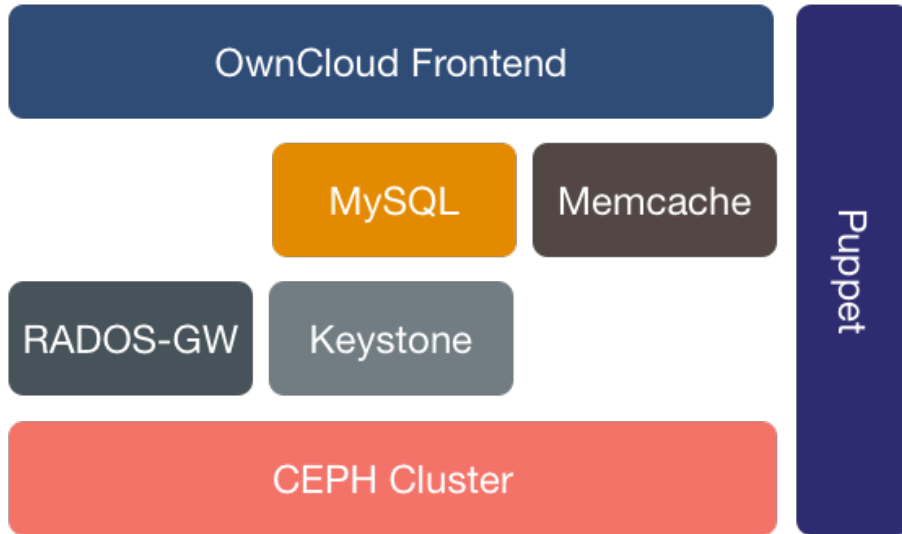
- CEPH/Swift backend
- Fully separated layers :
 - Frontend
 - DB
 - Storage
- No user sharding
- Components fully decoupled

UnipiBoxTrial:

Test the architecture with a relatively small amount of users and resources:

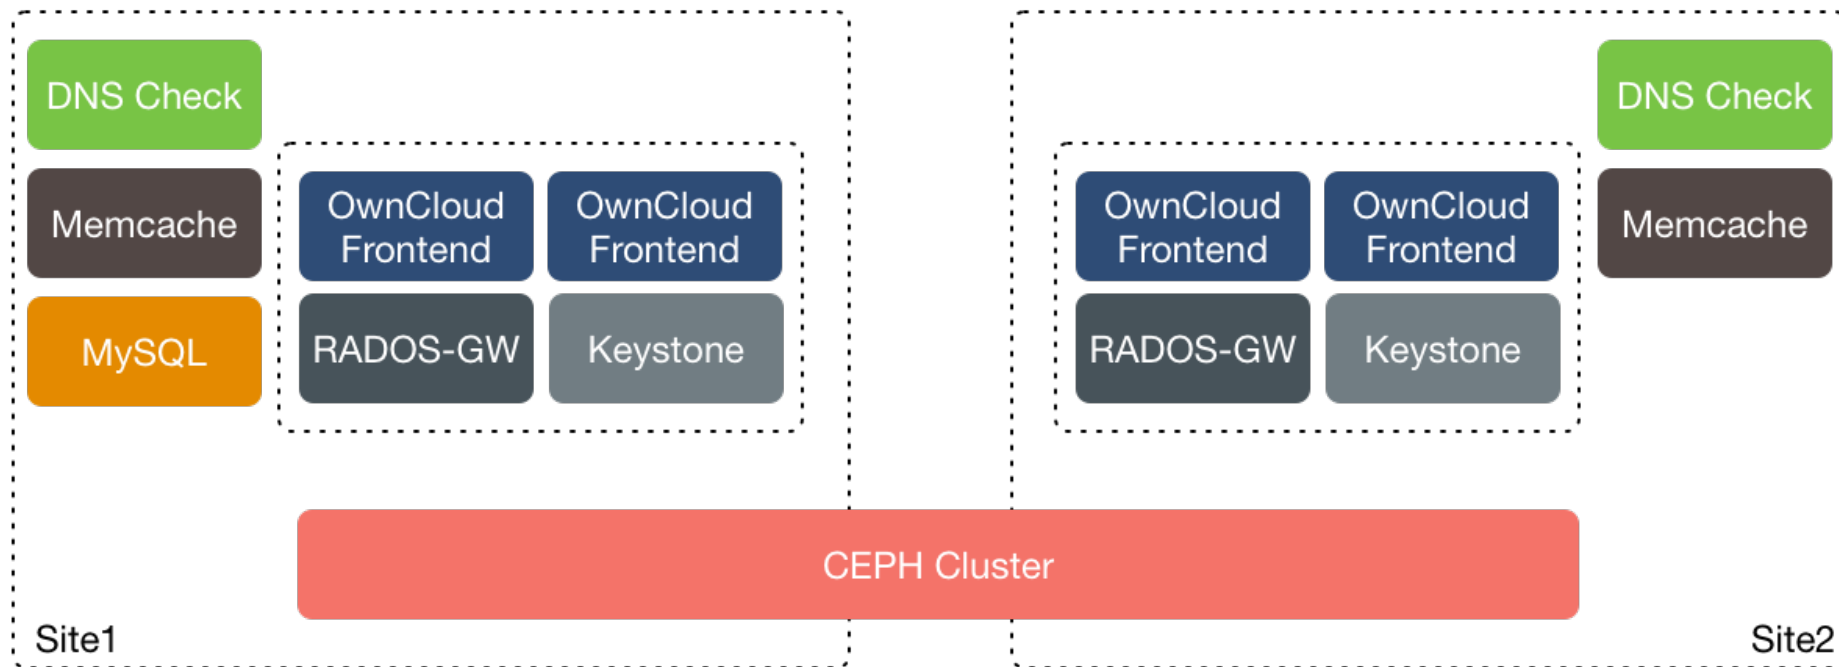
- 1600 students
- 300 employees
- 10GB quota
- 2 physical servers as “owncloud appliances”
- 1 physical server as DB

UnipiBox: the stack

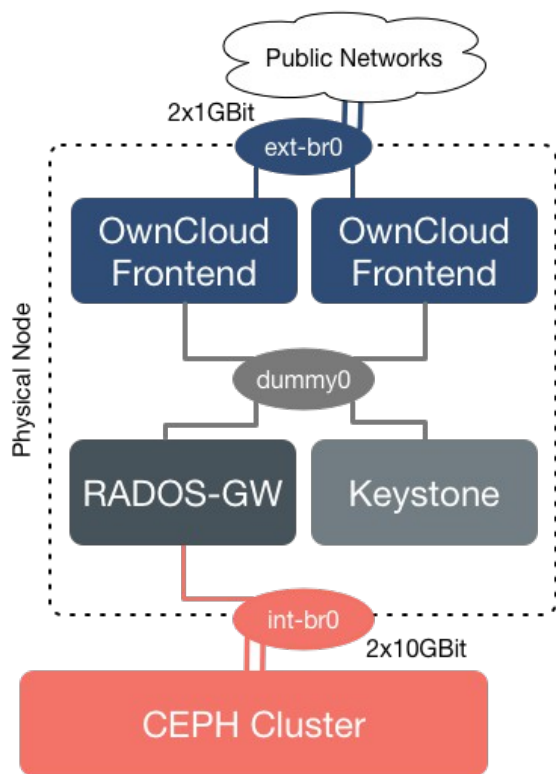


- Owncloud 8.2 CE with Nginx + PHP-FPM
- Memcache for shared sessions
- Mysql (no replica)
- CEPH as primary object storage in SWIFT flavor
- Radosgw + Keystone as SWIFT
- Puppet as Configuration Management

UnipiBoxTrial: an overview



UnipiBox: the “DIY appliance”



Xen PV virtual machines :


- Frontend node:
 - 4 vCPU
 - 16 GB RAM
 - 100GB HD (as upload buffer)
- Rados-gw node:
 - 8 vCPU
 - 16GB RAM
- Keystone Node
 - 2 vCPU
 - 1GB RAM

```
array (  
  'class' => 'OC\\Files\\ObjectStore\\Swift',  
  'arguments' =>  
    array (  
      'username' => 'OWC-USERNAME',  
      'password' => 'OWC-PWD',  
      'container' => 'OWC_CONTAINER',  
      'autocreate' => true,  
      'region' => 'regionOne',  
      'url' => 'http://192.168.252.249:5000/v2.0',  
      'tenantName' => 'OWC-TENANTNAME',  
      'serviceName' => 'swift',  
    ),  
)
```

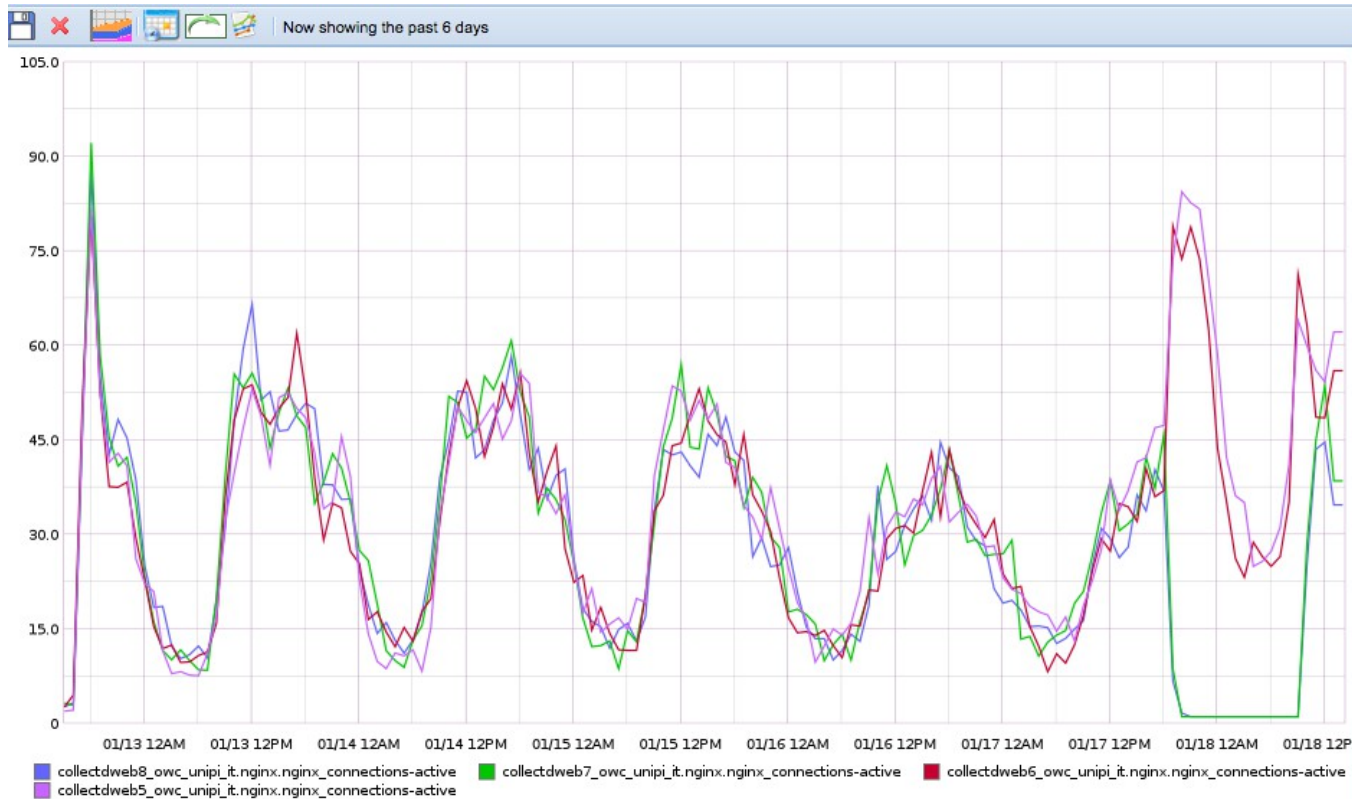
unipiBoxTrial: the opening

- People informed via email about the trial (Wed 13 Jan - 17:30)
- A sign-up website already open for 72 hours:
 - The student's accounts (1600) available ended in 1.5 h
 - The employees accounts (300) available ended in 2 h
- Actually we have more than 2000 users registered on LDAP (~100 manually added)
- Less than 20 trouble-ticket during the first 5 days (most of it to tell us they want more)

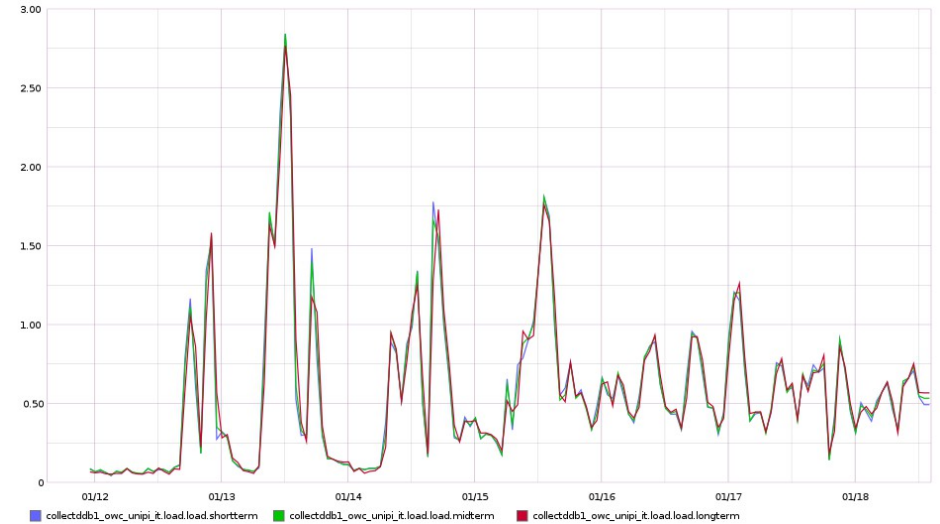

first results : CEPH-Objects ↔ oc_filecache



first results : frontends connections



first results : MySQL metrics



What to do now?

1. This is just a barebone: add functionalities
2. Users feedback
3. More monitoring and statistics
 1. long term metrics for capacity planning
 2. optimize VM sizing
4. Database is the most critical component for scaling (oc_filecache)

