Consortium GARR

THE ITALIAN EDUCATION & RESEARCH NETWORK

GARRbox

GARRbox
status and future directions

F. FARINA, P. VELATI, P. MANDATO – GARR

Rome, 29 January 2019

CS3 Workshop

Why GARRbox

Multi-year framework agreement between GARR and the Italian Ministry of Health

- GARR-X network high bandwidth connectivity
- Added value services: files sharing, HD-VCs, libraries & cloud storage

Researchers in bio-medicine, health, nutrition fields

- Not Universities, but small Organizations
- 200 researchers each institution on average
- GARR supports only R&E Organizations (58 over 81)

What GARRbox provides to researchers

- 20 TB aggregate storage, 20 GB personal quota
- No constraints on number of users – Organizations size may vary
- Support, authorization and management by GARR staff

GARRbox high-level architecture

Resources

- 3 RAID pools from different SANs
- 7 blades in different racks
- Production environment managed by VMware vCenter
- Pre-production environment on OpenStack + remote DR site

Ansible roles for subsystems

- Web: ownCloud, php-fpm, Nginx reverse proxy
- Storage: NFS, XFS, DRBD
- DBMS: Percona Cluster, MaxScale
- Primologin: custom AuthZ web service in Django
- Aux: Docker containers for caches, logging, monitoring, etc.

GARRbox high-level architecture

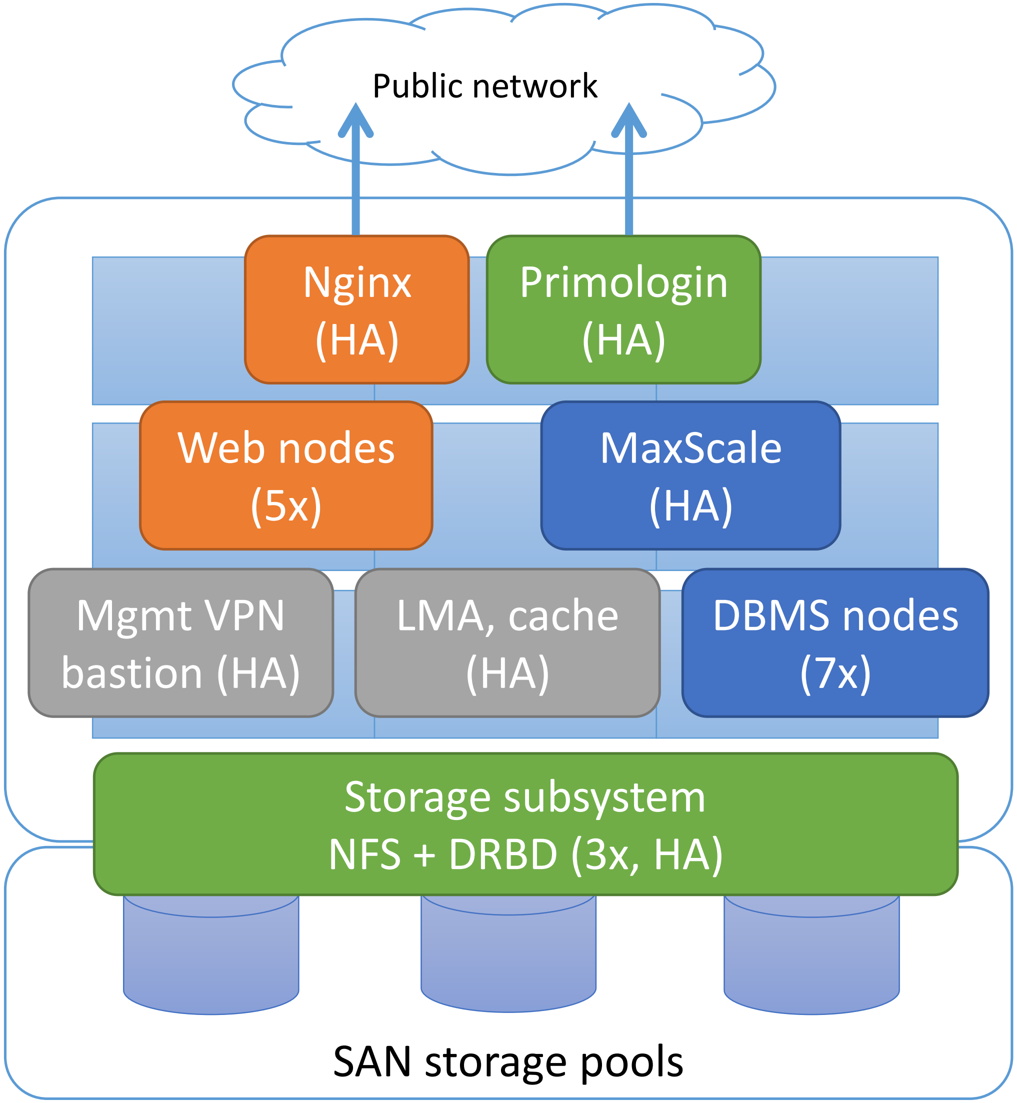

Service timeline

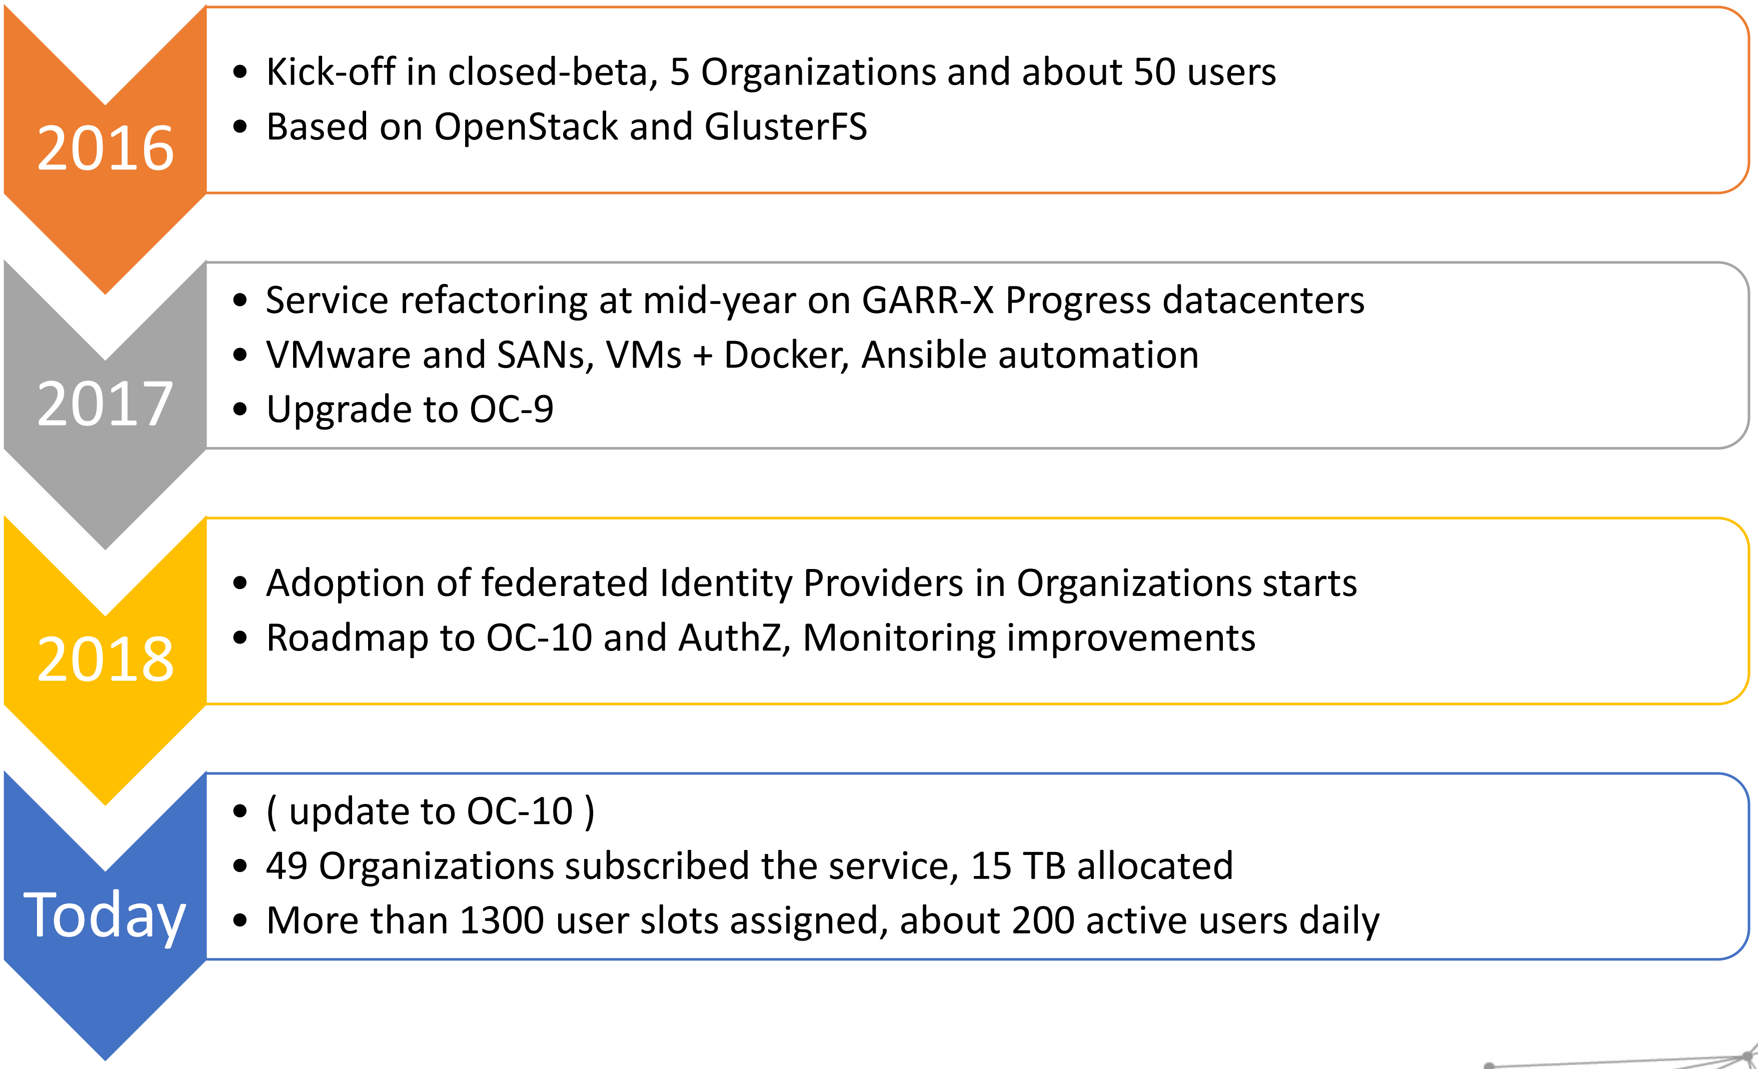

2016

- Kick-off in closed-beta, 5 Organizations and about 50 users
- Based on OpenStack and GlusterFS

2017

- Service refactoring at mid-year on GARR-X Progress datacenters
- VMware and SANs, VMs + Docker, Ansible automation
- Upgrade to OC-9

2018

- Adoption of federated Identity Providers in Organizations starts
- Roadmap to OC-10 and AuthZ, Monitoring improvements

Today

- (update to OC-10)
- 49 Organizations subscribed the service, 15 TB allocated
- More than 1300 user slots assigned, about 200 active users daily

AuthN/AuthZ

- Subscription by Organization, not single user: IDEM Identity Provider required
- Registration & password recovery → IDEM Federated Identity
- Access → OwnCloud local accounts and application tokens
- AuthZ enforcement → Primologin web service + post-login ownCloud App

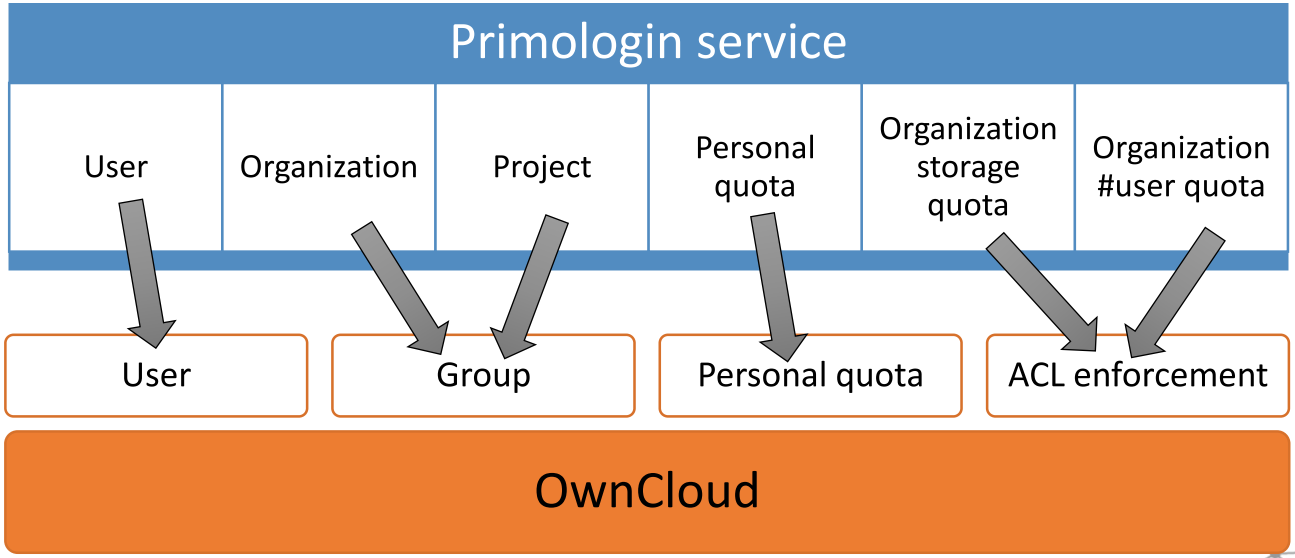

AuthN/AuthZ

- Domain Specific Language for Access Control

```
<allow|deny> if <attribute> [not] in [ "<pattern1>", "<pattern2>", ... ]

allow if email in [ "*@garr.it" ]                              # Access by email domain
deny if email in [ "*list*@garr.it", "*all@garr.it" ]          # Black-list lists

# Whitelist by username (ePPN)
allow if username in [ "user1@ente.it", "user2@ente.it", "user3@ente.it" ]
```

- Delegation models
  - Fixed – GARR manages directly policies and user support for subscribers
  - Flex – Delegate to Organization managers quota assignment and access control
  - Different subscription plans according the chosen delegation model

Future

Remove local user registration in favor of pure-SAML AuthN

- AuthZ post-login App refactoring was needed, Primologin upgrade
- OC-10 deployment

Enrich the AuthZ DSL features to face users’ new requirements

- Dynamic group assignment at login
- Quota dynamic update according to groups at login

On the user community

- awareness and adoption, slow progresses
- complete the Organizations opt-in process, related to IdP adoption

Future

GARR services central telemetry facility

- Follow-up of a dedicated PoC
- Elastic-stack for all GARR services and datacenter monitoring
- Multi-channel correlation for better understanding

Reduce further the operation effort

- On-going PoC
- Kubernetes and Helm, Docker for off-the-shelf components
- VMs only where strictly needed

Thanks! Questions?