

Provisioning Other Services

Alexander Dibbo

Contents

- What do we mean by other services?
- What services we run?
- How do we run services now?
- Why?
- Changes that need to happen
- Use of Enterprise Virtualisation and Cloud
- How will we run services?
- Longer term plans

What do we mean by other services?

- Anything which is not a critical infrastructure component of the storage or batch systems
 - These will continue to be run on bare metal or Enterprise Virtualisation as appropriate
- Any general services we provide
- Any services we host for VOs
- Any services which support the rest of the infrastructure

What services do we run?

General Services

CernVM FileSystem (Stratum 0 and Stratum 1)
ARGUS (Site and NGI)
FTS
LFC
BDII (Site and Top)
ARC CEs
Squids
Uis
MyProxy

VO Services

Frontier
VO Boxes
Rucio (In testing with SKA)

Infrastructure Services

Database Cluster (MariaDB + Galera)
Database (MariaDB or MySQL)
DNS
Monitoring
Config Management
Load Balancers

How we run services now?

- Multiple infrastructures on which can be hosted
- Most services have individual endpoints or round robin DNS for high availability
- Some services have been moved to load balancers to provide high availability and scale out capability
- Some services are “Pets” with single instances

How do we run services now?

Bare Metal

Necessary for some services.

Can be hard to justify with squeezed budgets

Hyper V Cluster

In full Operation

Standalone

In full operation

VMWare

In full operation

Can support larger VMs (CPU + RAM not storage) than HyperV

Consolidation within SCD

Cloud

OpenNebula

In full operation but winding down

OpenStack

Shared

Storage

In full operation



How do we run services now?

Bare Metal

CernVM File System
(Stratum 1)

DNS

Hyper V Cluster

CernVM File System (Stratum 0)

DNS, Frontier DB Cluster
Config Man
CEs, Squids, myproxy
BDII, ARGUS, FTS

Standalone

Frontier DB Cluster
BDII, ARGUS, FTS
CEs, Squids

VMWare

Awaiting migration of services

Cloud

OpenNebula
Frontier

OpenStack
Shared Storage
Rucio



Why?

- Consolidation of expertise in SCD
 - There is more expertise in SCD on VMWare so that is the Enterprise Virtualisation we should run
- Reduce effort required to run the services we need to run.
- Cloud can be great
 - We should exploit the STFC Cloud where it makes sense to
- Make services more resilient
 - Services behind load balancers can have nodes fail without bringing down the service
 - This should reduce call outs

Changes that NEED to happen

- Migration of services on Hyper V cluster to VMWare cluster
 - Needs to happened by end of November
- Hyper V Cluster decommissioning
 - Shared storage will be out of warranty at the end of November
- OpenNebula decommissioning
 - User migration is in progress. Current target is end of October. Racks are needed for next year's procurements.
- Local storage hypervisors and flavors for OpenStack
 - Has already been tested and will be rolled out soon. First should be in place in October.
- Create a second VMWare Cluster
 - This will be the decommissioned hyper V cluster
- All appropriate services should be behind load balancers

Use of Enterprise Virtualisation

- Only production services will be allowed
 - Limited resource pre production environments will be allowed
- No development services

Use of Cloud (Tier 1 perspective)

- Any appropriately architected service allowed
 - Service owners should make sure VMs are distributed and make use of shared and local storage
- Any service which is managed externally
 - VO services
- Additional instances of any service for availability or performance reasons
- Bursting of batch farm

How will we run services?

Bare Metal

CernVM File System (Stratum 1)

DNS

VMWare Clusters

CernVM File System (Stratum 0)
DNS, Frontier
DB Cluster
Config Man
CEs, Squids
BDII, ARGUS, FTS
Frontier
MyProxy

Cloud

OpenStack
Shared Storage

Rucio

Frontier

VO Boxes

Bursting of services

Local Storage

Rucio

Frontier



Longer Term

- No current plans to change the operating model.
 - Minimise effort spent of R&D
- Exploit efforts around container orchestration where appropriate
 - Work is being done within SCD around Kubernetes, Rancher and OpenShift (STFC Cloud Team, Data Division and DAFNI)
 - Other efforts may be generalisable
- Exploit containerised software where easy and appropriate
 - If there becomes a simpler way we will use it

What about Other Other Services?

Facilities/Cloud/IRIS

- It is expected that STFC Facilities services which work with the Tier 1 will also follow this plan
 - Ever increasing use of STFC Cloud for service delivery
 - Data-Analysis-as-a-Service
- STFC Cloud team (supports SCD, STFC Facilities, ALC and IRIS) is expanding
 - Focus on enabling technologies (Self service virtualisation and networking, orchestration of VMs and Containers, Load balancers)
 - Give “users” what they need to run their own reliable services
- Create Digital Assets where appropriate to support user communities
 - i.e. multi tenant Rucio deployment
- If anyone wants to know more about these then give me a shout.



Science & Technology
Facilities Council

UK Research
and Innovation

Any Questions?

alexander.dibbo@stfc.ac.uk