



Science & Technology Facilities Council

e-Science

Virtualisation & Cloud Computing at RAL

Ian Collier- RAL Tier 1

ian.collier@stfc.ac.uk

HEPiX Prague 25 April 2012



Virtualisation @ RAL

- Hyper-V Services Platform
- E-Science Cloud
- EGI Federated Cloud Task Force
- Jasmine/CEMS
- Contrail



Virtualisation @ RAL

- **Hyper-V Services Platform**
- E-Science Cloud
- EGI Federated Cloud Task Force
- SCT Jasmine
- Contrail



Hyper-V Platform

- Development & testing use for over a year
 - Local storage
 - Small test batch system
 - Examples of all grid services nodes
 - Key in testing/rolling out EMI/UMD middleware
 - Test castor head nodes
 - etc, etc.
- Progress of high availability platform (much) slower than we'd have liked
 - Technical issues with shared storage
 - Took a **long** time to procure Equallogic arrays after successful evaluation early last year
 - Just arrived a couple of months ago
 - But the 10gig interconnects are incompatible



Hyper-V Platform

- Recently moved first non-resilient external services in to full production
 - fts, myproxy - argus coming
- Also internal databases & monitoring systems
- Move to production very smooth
 - Team familiar with environment



Hyper-V Platform

- 18 Hypervisors deployed
 - 10 R410s & 510s w 24GB RAM & 1TB local storage
 - 8 New R710, 96GB RAM 2TB local storage
- All gigabit networking at present
 - Migrating to 10 gigabit over coming weeks/months (interconnect compatibility issues between both hosts-switches and storage-switches)
- ~100 VMs – nearly all Linux
 - 10% production services 90% dev. & testing



Hyper-V Platform

- However, Windows administration is not friction or effort free (we are mostly Linux admins....)
 - Share management server with corporate IT – but they do not have resources to support our use
 - Troubleshooting means even more learning
 - Some just ‘don’t like it’
- Hyper-V continues to throw up problems supporting Linux
 - None show stoppers, but they drain effort and limit things
 - Ease of management otherwise compensates for now
- Since we began open source tools have moved on
 - We are not wedded to Hyper-V



Virtualisation @ RAL

- Hyper-V Services Platform
- **E-Science Cloud**
- EGI Federated Cloud Task Force
- SCT Jasmine
- Contrail



E-Science Cloud

- Prototype E-Science Department cloud platform
- Initially for internal test & development systems
 - Aim to provide resources across STFC
 - Both scientific computing & ‘general purpose’ systems
 - Potentially federated with other scientific clouds
 - Based on StratusLab – moving target as project develops
- Work done by graduate on 6 month rotation
 - They’ve moved on
 - Now waiting for new staff to continue work

StratusLab





E-Science Cloud

- Very fruitful security review
 - For now treat systems much like any Tier 1 systems
 - Monitor that eg central logging is active, sw updates are happening
 - Cautious about user groups we open things to
 - Will need work before we can take active part in federated clouds
 - Need better network separation – coming to StratusLab
- Mostly developed using old (2007) WNs
 - At end of 1st phase deployed 5 R410s with quad gigabit networking
 - Enough to
 - Run a meaningful service
 - continue development to cover further use cases
 - Still evaluating storage solutions



Virtualisation @ RAL

- Hyper-V Services Platform
- E-Science Cloud
- **EGI Federated Cloud Task Force**
- SCT Jasmine
- Contrail



EGI Federated Cloud Taskforce

- Colleagues in SCT working on accounting
- EScience Cloud tracking work, not quite ready to take active part (see security/policy discussion above)
- Dedicated talk on Friday



Virtualisation @ RAL

- Hyper-V Services Platform
- E-Science Cloud
- EGI Federated Cloud Task Force
- **Jasmine**
- Contrail



JASMIN/CEMS

The JASMIN super-data-cluster

- UK and European climate and earth system modelling community.
- Climate and Environmental Monitoring from Space (CEMS)
- Facilitating further comparison and evaluation of models with data.

4.6 PB Storage Panasas at STFC

- Fast Parallel IO to Compute servers (370 Cores)





Science & Technology Facilities Council

e-Science

JASMIN/CEMS



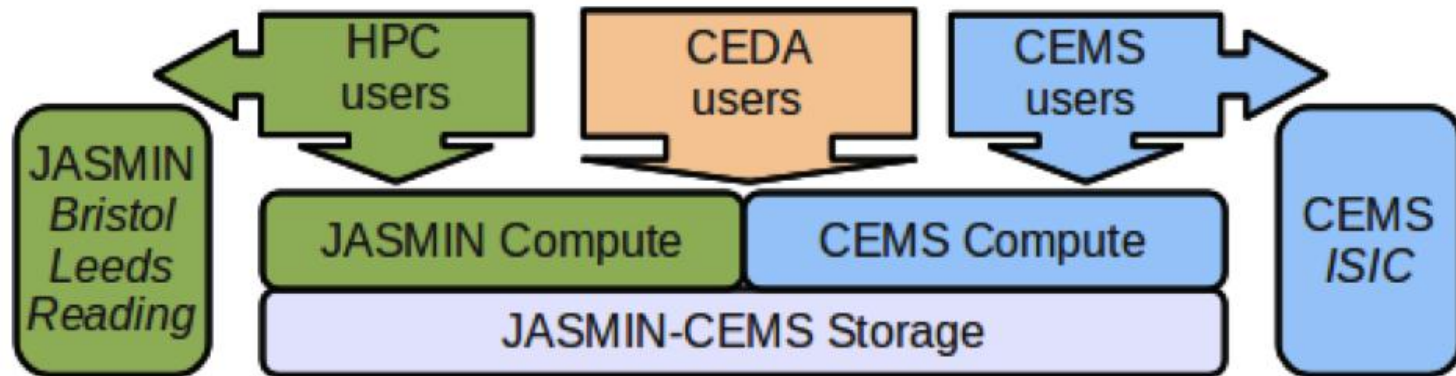


JASMIN Super Data Cluster

JASMIN	3.5 PetaBytes Panasas Storage 20 x Dell R610 (12 core, 3.0GHz, 96G RAM) 1 x Dell R815 (48 core, 2.2GHz, 128G RAM) 1 x Dell Equallogic R6510E (48 TB iSCSI VM image store) VMWare vSphere Center 1 x Force10 S4810P 10GbE Storage Aggregation Switch
CEMS	1.1 PetaBytes Panasas Storage 7 x Dell R610 (12 core 96G RAM)Servers 1 x Dell Equallogic R6510E (48 TB iSCSI VMware VM image store) VMWare vSphere Center + vCloud Director



JASMIN Super Data Cluster



JASMIN provides three classes of service:

- Virtualised compute environment (not strictly a "private cloud").
- Physical compute environment.
 - No private data connection
- HPC service ("Lotus").
 - Separate data connection.
 - Not easily reconfigurable to JASMIN cloud.



JASMIN Super Data Cluster

Two distinct clouds

- One supports manual VM provisioning by CEDA and the climate HPC community
 - Configuration controlled at site
 - Therefore greater trust and greater network access
- One supports more dynamic provisioning by the academic users in the CEMS community.
 - Users provision own VMs
 - Access to Panasas
 - Otherwise less trusted
- So, they have different vCentre server installations.



Virtualisation @ RAL

- Hyper-V Services Platform
- E-Science Cloud
- EGI Federated Cloud Task Force
- Jasmine
- **Contrail**



Contrail

- Integrated approach to virtualization
 - Infrastructure as a Service (IaaS)
 - Services for federating IaaS Clouds
 - Platform as a Service (PaaS) on top of federated Clouds.
- STFC e-Science contribution
 - identity management
 - quality of service
 - security



Virtualisation @ RAL

- Many strands
 - Hyper-V Services Platform
 - E-Science Cloud
 - EGI Federated Cloud Task Force
 - Jasmine/CEMS
 - Contrail

