GridPP

UK Computing for Particle Physics

Virtualisation at the RAL Tier 1

Ian Collier

STFC RAL Tier 1

HEPiX, Annecy, 23rd May 2014

Science & Technology Facilities Council

e-Science

GridPP
UK Computing for Particle Physics

Virtualisation @ RAL Tiere 1

- Context at RAL Tier 1
- Hyper-V Services Platform
- Scientific Computing Department Cloud
- Dynamically provisioned worker nodes

Science & Technology Facilities Council
e-Science

Context at RAL Tier 1

- Historically requests for systems went to fabric team
  - Procure new HW – could take months
  - Scavenge old WNs – could take days/weeks – ***and*** they are often unreliable
- Kickstarts & scripts needed customising for each system
- Not very dynamic
- For development systems many users simply run VMs on their desktops
  - Hard to track
  - Often not well managed - risky

Evolution at RAL Tier 1

Many elements play their part

- Configuration management system
  - Quattor (introduced in 2009) abstracts hardware from os from payload, automates most deployment
  - Makes migration & upgrades much easier (still not completely trivial)
- Databases store hardware info – scripts feeding configuration management system
  - Provisioning new hardware much faster
  - With Aquilon it is easier still

Virtualisation & Cloud @ RAL

- Context at RAL
- Hyper-V Services Platform
- Scientific Computing Department Cloud
- Dynamically provisioned worker nodes

Hyper-V Platform

- Over last three years
  - Initially Local storage only in production
  - iSCSI (EqualLogic) shared storage in production 1 year ago
  - ~250 VMs
- Provisioning transformed
  - Much more responsive to changing requirements
  - Self service basis – requires training all admins in using management tools – but this proved not so hard

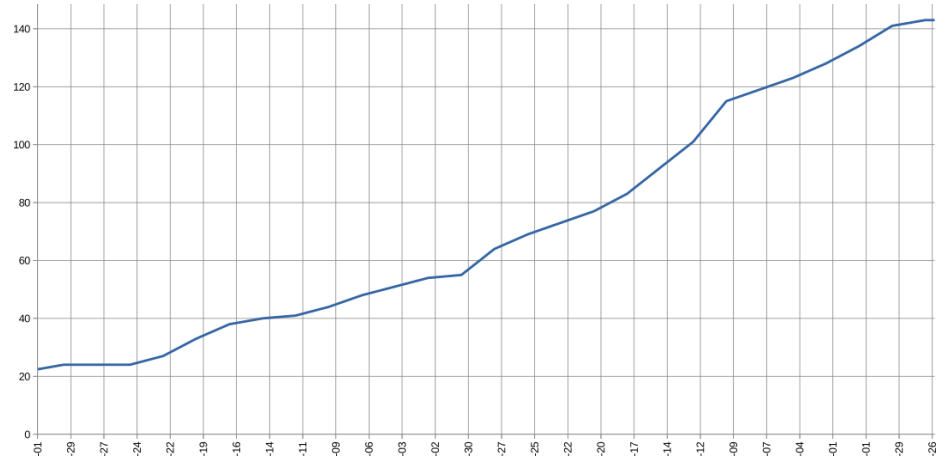

Virtualisation

- Three production clusters with shared storage, several local storage hypervisors
  - Windows Server 2008 + Hyper-V (2012 being tested now)
  - Clusters distributed between two buildings
- However, saw issues with VMs
  - Migration problems triggered when we needed to shut down cluster for power intervention
- Re-building the shared-storage clusters from scratch
  - New configuration of networking and hardware
  - Windows Server 2012 and Hyper-V
  - Currently migrating most VMs back to local storage systems

Hyper-V Platform

- Most services virtualised now
  - Exceptions: top bdii, ganglia, Oracle
- Internal databases & monitoring systems
- Also test beds (batch system, CEs, bdiis etc)
- Move to production mostly smooth
  - Team had good period to become familiar with environment & tools
- Shared storage placement & migration managed by production team.

Hyper-V Platform

- When a Tier 1 admin needs to set up a new machine all they have to request is a DNS entry
  - Everything else they do themselves
  - Placement of production system by production team
- Maintenance of underlying hardware platform can usually be done with (almost) no service interruption.
- This is already much, much better – especially more responsive – than what went before.
- Behaved well in power events

Hyper-V Platform

- However, Windows administration is not friction or effort free (we are mostly Linux admins….)
  - Troubleshooting means even more learning
  - Share management server with STFC corporate IT – but they do not have resources to support our use
  - In fact we are largest Hyper-V users on site
- Hyper-V presented problems supporting Linux
  - None show stoppers, but they drained effort and limited use
  - Ease of management otherwise compensates for now
  - **Much** better with latest SL (5.9 & 6.4)
- Since we began open source tools have moved on
  - We are not wedded to Hyper-V
  - Realistically will run it for a while

GridPP
UK Computing for Particle Physics

Virtualisation & Cloud @ RAL

- Context at RAL
- Hyper-V Services Platform
- Scientific Computing Department Cloud
- Dynamically provisioned worker nodes

Science & Technology Facilities Council
e-Science

SCD Cloud

- Prototype E-Science/Scientific Computing Department cloud platform*
- Began as small experiment 2 years ago
- Using StratusLab
  - Share Quattor configuration templates with other sites
  - Very quick and easy to get working
  - But has been a moving target as it develops
- Deployment implemented by graduates on 6 month rotation
  - Disruptive & variable progress
- Worked well enough

*Department changed name

SCD Cloud

- Resources
  - Initially 20 (very) old worker nodes
    - ~80 cores
    - Filled up very quickly
    - 2 years ago added 120 cores in new Dell R410s – and also a few more old WNs
    - Also added half a generation of retired WNs ~800 cores
    - Retired disk servers for shared storage
  - This has been enough to test a number of use cases

SCD Cloud

- Last summer established principle that we were ready to move from best effort prototype to a service
- Use cases within and beyond Tier 1 coming in to focus
  - Exposing cloud APIs to LHC VOs
  - Platform for SCD projects (eg H2020 projects)
  - Compute services for other STFC Departments (ISIS, Diamond)
  - Self service IaaS across STFC
- Agreed budget for 1FTE to April 2015 (end of GridPP 4)
- In February we found someone great
  - At the start of this month they took a different job ☹
- In March we managed to 'catch' £300K of underspend
  - ~1PB of disk (for Ceph backend)
  - ~1000 cores of compute ~3.5TB RAM

SCD Cloud Future

- Develop to full supported service to users across STFC
- IaaS upon which we can offer PaaS
  - One platform could ultimately be the Tier 1 itself
  - Integrating cloud resources in to Tier 1 grid work
- Participation in cloud federations
- Carrying out fresh technology evaluation.
  - Things have moved on since we started with StratusLab
  - Currently favour Opennebula

GridPP
UK Computing for Particle Physics

Virtualisation & Cloud @ RAL

- Context at RAL
- Hyper-V Services Platform
- Scientific Computing Department Cloud
- Dynamically provisioned worker nodes

Science & Technology Facilities Council
e-Science

Bursting the batch system into the cloud

- Aims
  - Integrate cloud with batch system, eventually without partitioned resources
  - First step: allow the batch system to expand into the cloud
    - Avoid running additional third-party and/or complex services
    - Leverage existing functionality in HTCondor as much as possible
    - Should be as simple as possible
- Proof-of-concept testing carried out with StratusLab cloud
  - Successfully ran ~11000 jobs from the LHC VOs

Power management in HTCondor

- Existing functionality for powering down idle machines & waking them when required
  - Entering a low power state
    - HIBERNATE expression can define when a slot is ready to enter a low power state
    - When true for all slots, machine will go into the specified low power state
  - Machines in a low power state are “offline”
    - The collector can keep offline ClassAds
  - Returning from a low power state
    - condor_rooster daemon responsible for waking up hibernating machines
    - By default will send UDP Wake On LAN
      - Important feature: this behaviour can be replaced by a user-defined script
  - When there are idle jobs
    - Negotiator can match jobs to an offline ClassAd
    - condor_rooster daemon notices this match & wakes up the machine

Provisioning worker nodes

- Our method: extend condor_rooster to provision VMs
- What we did
  - Send an appropriate offline ClassAd to the collector
    - Hostname used is a random string
    - Represents a class of VMs, rather than specific machines
  - condor_rooster
    - Configure to run appropriate command to instantiate a VM
  - When there are idle jobs
    - Negotiator can match jobs to the offline ClassAd
    - condor_rooster notices this match & instantiates a VM
  - VM lifetime managed by HTCondor on the VM itself
    - START expression modified so that jobs can only start for a limited time
    - HIBERNATE expression set to shutdown the VM after being idle for too long

Observations

- Our VMs were almost exact clones of normal WNs
- Including all monitoring
- Condor may deal well with dynamic resources
- Nagios does not
- When we thought about it we realised almost all the WN monitoring was unnecessary on virtualised WNs.
- Really just need the health check hocked in to startd
- With minimal tuning efficiency loss was 4-9%
- All in all a very successful first step.

GridPP
UK Computing for Particle Physics

Virtualisation & Cloud @ RAL

- Context at RAL
- Hyper-V Services Platform
- Scientific Computing Department Cloud
- Dynamically provisioned worker nodes

Science & Technology Facilities Council
e-Science

Summary

- Using range of technologies
- Many ways our provisioning & workflows have become more responsive.
- Private cloud has developed from a small experiment to the beginning of a real service
  - With constrained effort
  - Slower than we would have liked
  - The prototype platform has been well used
- Have demonstrated we can transparently and simply expand batch farm into our cloud.
- Ready to start making much larger changes.