

Manchester Site Report

Alessandra Forti, Andrew McNab, Robert Frank

Overview

- Foreman Migration
- Procurement 2016/17
- Hardware
- CE / Batch System
- VAC
- Current and Future tasks

Foreman Migration

- Finalised Foreman set up in December.
- Started to migrate to Foreman in January.
 - Batch system WNs done.
 - Most hypervisors done.
 - But only some VMs done.
 - Disk servers are causing problems.
- Certificate issues with Hammer CLI.
 - No verification of server certificate.
 - No support for client certificate authentication.
 - Both fixed now.

Procurement 2016/2017

- Spent most of the money already.
- Compute and storage, both Dell.
 - 6 PowerEdge C6320, E5-2620v4, 64GB
 - 15 PowerEdge R730xd, 12x10TB
- Storage arrived within a week.
 - Put in racks the next day and cables the day after.
 - Migration to Foreman delayed installation.
- Compute took longer to arrive.

Procurement 2016/2017

- Compute nodes with add-on 10GbE cards.
- PXE booting in BIOS mode only from on-board cards → had to use UEFI.
- Figure out how to do UEFI boot with Foreman.
 - Grub UEFI for SL6, Grub2 UEFI for CentOS7.
 - Can't use boot partition on MD raid on CentOS7.
 - OS installation changes UEFI boot order.
 - Change it back in post install.
 - BIOS setting to prevent access to UEFI variables.

Hardware

- Compute nodes
 - 20 Viglen HX425T2i quad nodes
 - 7 Dell PowerEdge C6100 quad nodes
 - 11 Viglen HX525T2i quad nodes
 - 1 OCF GPU node (4 Tesla K40c GPU)
 - 4 XMA HX625T2i quad nodes
 - 6 Dell PowerEdge C6320 quad nodes

Hardware

- Storage
 - 19 Viglen HX425S-36i
 - 14 with 2 TB disks
 - 1 with 3 TB disks
 - 4 with 4 TB disks
 - 22 Dell PowerEdge R730xd
 - 7 with 8 TB disks
 - 15 with 10 TB disks
 - Total maximum capacity: 3363 TB



Hardware

- Servers
 - Viglen IX2300R
 - 12 Dell PowerEdge R610
 - 1 Dell PowerEdge R630
 - Being used as KVM hypervisors to run service VMs, including Perfsonar.

ARC/HTCondor

- Moving away from Cream/Torque
 - HTCondor-CE very appealing
 - Reduces amount of configuration needed
 - Same log files and tools when debugging
 - Not integrated enough in Europe
 - ARC-CE/HTCondor
 - 5.3.1/8.6.0
 - CentOS7 WNs
 - UMD4

Puppet changes

- Some changes included in github: to integrate more recent versions of HTCondor and in CentOS7 [issue #58](#) [issue #59](#)
- Some other modules/classes written to integrate in our puppet system
 - `create_dirs` to create directories for every module where they might clash
 - `grid_accounts` to create and delete pool accounts, `gridmapdir`, `grid-mapfile`, `groupmapfile` (is it really still needed?)
 - `voms_client` to replace the CERN voms module dependency

CentOS7 WNs

- Tested UMD4 middleware
 - Had to downgrade HTCondor to 8.6.0 due to clashing dependencies [MWREADY-135](#)
 - This has been solved now
 - Still using YAIM underneath but replaced already a number of functions with puppet classes
 - voms_client for example.
 - CREAM written puppet modules for CE and WNs but use puppet 4.0

Migration Plans

- Currently three CEs, ce01, ce02, ce03.
- Most compute nodes moved from ce02 to ce01.
- ce02 to become the first production ARC-CE.
- Move compute nodes from ce01 to ce02.
- ce01 to become second ARC-CE.
- Move compute nodes from ce03 to ce01.

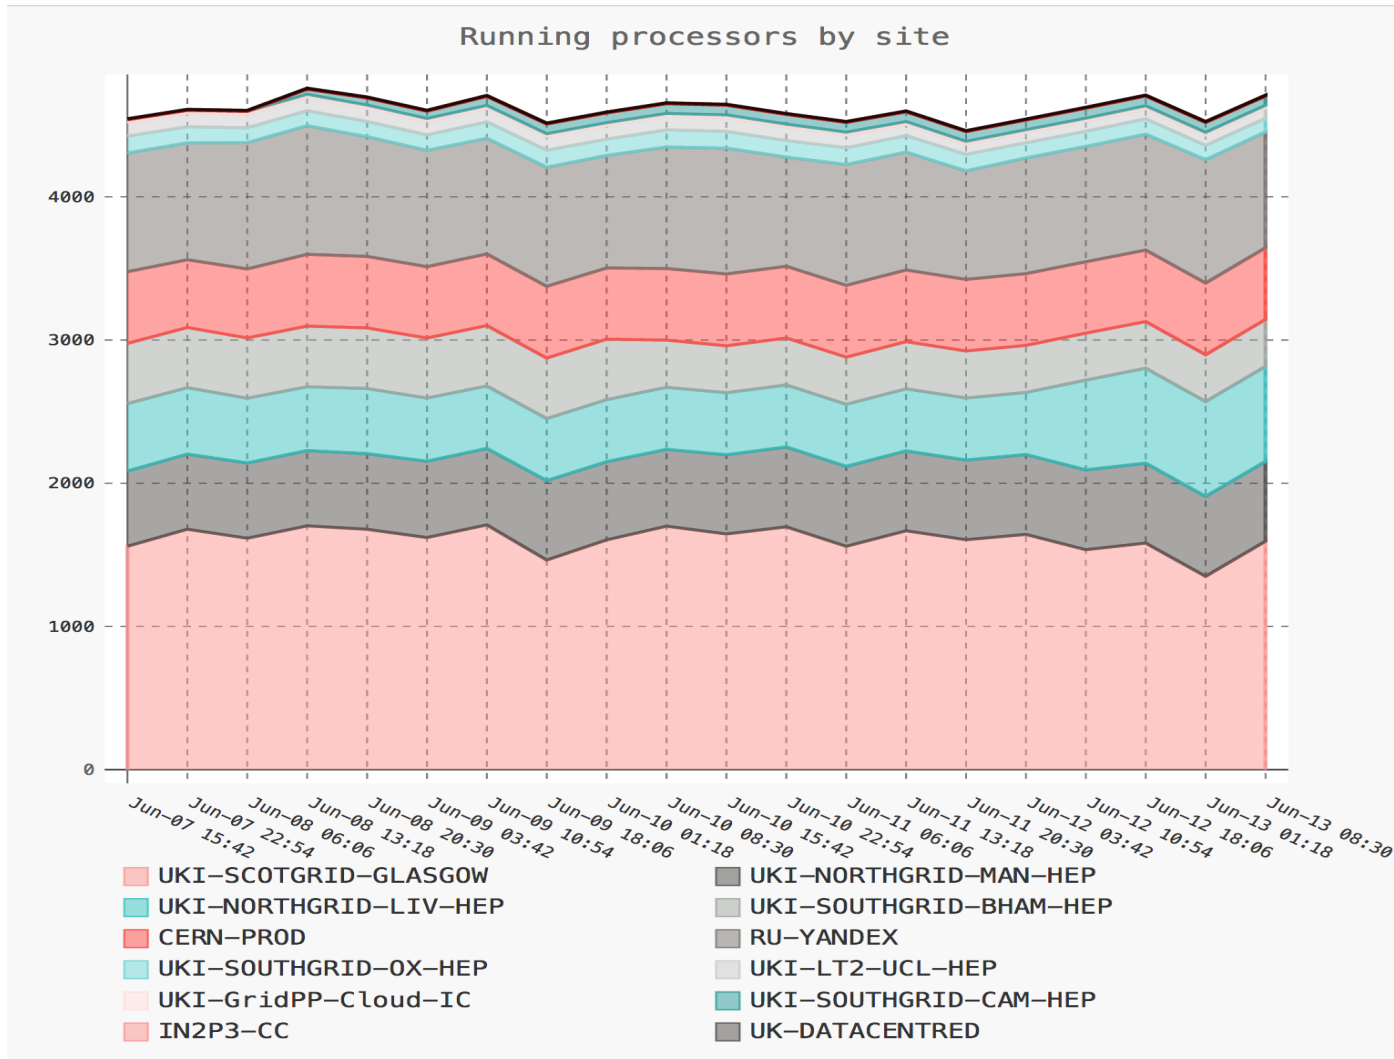
Vac at Manchester

- Three vac spaces (=clusters)
 - Each 1 rack
 - vac01 12 hosts x 24 processors each. Vac-in-a-Box
 - vac02 12 x 24. Managed by Puppet (2 dead – HW failure)
 - vac04 10 hosts x 4 processors each
 - Vac and Vac-in-a-Box development
 - 568 processors in total
 - Squid caches on each host, with failover to site squid caches
 - Very little OS/Software maintenance. Effort goes into HW

GridPP Vac/Vcycle services

- repo.gridpp.ac.uk
 - Software and VM definition distribution
- depo.gridpp.ac.uk
 - Log file repository
- viab.gridpp.ac.uk
 - Vac-in-a-Box dashboard and repository
- vm21
 - GridPP Vcycle service and development
 - Manages VMs at CERN and DataCentred OpenStack sites
- vacmon.gridpp.ac.uk
 - VM and VM-factory monitoring with ElasticSearch.

Example of Vac monitoring graphs



Current and Future Tasks

- Move to ARC-CE/HTCondor.
- Move remaining machines to Foreman.
- Upgrade DPM to 1.9 (still on 1.8).
- Upgrade to Puppet 4.
- Move Nagios.

? Is there enough time left for questions ?