

Oxford Site Report

HEPSYSMAN 2019

Pete Gronbech, Kashif Mohammad and Vipul Davda

Compute Nodes

- 138 Servers ~3460 Logical CPUs
 - 66 Supermicro
 - 36 Dell
 - 36 Lenovo/IBM
- CentOS 7

ARC-CE/HTCondor Manager Node

- CentOS 7

DPM Storage

- 33 SE pool nodes - SL6
- 1 Head node - CentOS 7



Fully automated (close, but not quite):

- Cobbler to install base OS
- Puppet (Version 4) - Configuration management
 - ~60 Modules
 - “Inhouse” modules are on local Gitlab server for source management.

Building

- Assign an IP address to a hostname
- Add MAC address, IP address and hostname to the DHCP server
- Add hostname to Cobbler - Name, Puppet profile, MAC address and IP address
- Boot

HTCondor with ARC CE

SL6

- Migrated all compute nodes to CentOS 7
- Retired SL6 ARC CE and 11 old compute nodes on 3rd May 2019

CentOS 7

- t2arc00: Production CE attached to ~3460 logical CPUs

DPM Storage ~800TBytes

- Head node running DPM version 1.12.0 on CentOS 7
- SE pool nodes running DPM 1.12.0 on SL6
 - There is still an issue with hanging semaphore
 - A weekly cron job to remove hanging semaphore

DOME upgrade is in progress

Cvmfs mount can go stale:

- Fewer jobs arrive, so the cvmfs file system needs reloading manually

Auto Check Script

- A python script running daily under cron, to check the cvmfs file system and to reload it if errors are detected

Inhouse python script using HTCondor Python bindings

Condor Jobs Overview on t2arc00.physics.ox.ac.uk

Wednesday, 22 May 2019 08:50:02

Jobs per VO

VO	Status	CPUs	#Jobs	#Cores	CPU Eff
group_ALICE_admin	running	1	45	45	94.60
group_ATLAS_multicore	idle	8	95	760	
group_ATLAS_multicore	running	8	97	776	67.44
group_ATLAS_pilot	running	1	135	135	76.30
group_ATLAS_production	running	1	4	4	38.20
group_CMS_pilot	idle	1	141	141	
group_CMS_pilot	running	1	499	499	92.67
group_DUNE_pilot	idle	1	33	33	
group_DUNE_pilot	running	1	99	99	186.00
group_LHCB_pilot	running	1	2	2	99.19
group_SOLID_pilot	running	1	1	1	99.15

CPU Summary

#WorkerNodes	#CPUs	#CPUsUsed	#CPUsFree	%CPUsUsed	LoadEff
136	3444	1562	1882	45	48

Jobs Summary

jobs	completed	flocked	idle	held	running	removed
1156	123	0	270	0	886	0

Jobs Exceeding CPU Threshold Of 120%

WorkerNode	JobID	Job Owner	CPUs	CPU Eff	Time
t2wn164	1705791.0	DUNE_pilot	1	225.85%	1 day 10:32:50
t2wn164	1706007.0	DUNE_pilot	1	226.40%	1 day 10:08:20
t2wn159	1706904.0	DUNE_pilot	1	245.96%	1 day 07:57:31
t2wn164	1706960.0	DUNE_pilot	1	219.11%	1 day 07:47:02
t2wn164	1706962.0	DUNE_pilot	1	231.79%	1 day 07:46:39
t2wn159	1707097.0	DUNE_pilot	1	243.57%	1 day 07:34:39
t2wn164	1707333.0	DUNE_pilot	1	226.97%	1 day 06:59:04
t2wn162	1707972.0	DUNE_pilot	1	206.03%	1 day 05:20:24
t2wn149	1708066.0	DUNE_pilot	1	182.58%	1 day 04:50:58
t2wn149	1708067.0	DUNE_pilot	1	179.32%	1 day 04:50:37
t2wn149	1708068.0	DUNE_pilot	1	181.96%	1 day 04:49:17
t2wn149	1708145.0	DUNE_pilot	1	190.42%	1 day 04:37:15
t2wn162	1708184.0	DUNE_pilot	1	246.86%	1 day 04:36:00

Compute Nodes

- SL6 - 37 physical servers
 - 29 Supermicro
 - 8 Dell
- CentOS 7 - 7 physical servers ~192 logical CPUs
 - 2 Supermicro
 - 5 Dell

HTCondor Manager Node


- CentOS 7 server

Storage

- Gluster 13 physical servers, CentOS 7 ~930TBytes
- NFS 13 physical servers, SL6/CentOS 7 ~450TBytes

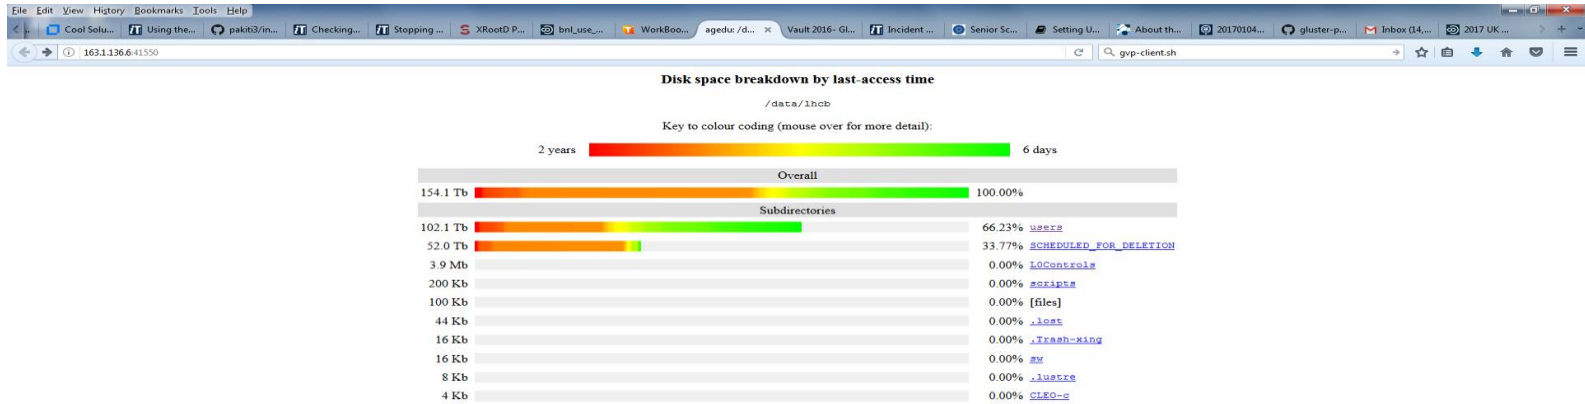
Interactive Servers

- Two SL6 interactive server
- One CentOS 7 interactive server



- Torque and Maui on SL6
 - The plan is to retire and migrate the compute nodes to CentOS 7
- HTCondor on CentOS 7
 - Fully configured by Puppet
 - Now there are few active users

- RedHat supported product
- RedHat developers are active on the mailing list
- Atlas has ~630TBytes storage and more than 100 million files
- LHCb has ~300TBytes storage
- Issues:
 - Large numbers of small size files slow down operation
 - Error logging is not helpful for diagnosis



Since there are millions of files on the gluster FS, agedu command takes weeks to complete.

- Storage/Interactive server
 - NFS ~55TBytes storage
- Two Worker Nodes
 - 32 Cores
- Slurm Batch System

Supported OS:

- Windows Desktops/Laptops and MAC Laptops
 - - supported by others in the group
- Ubuntu Desktops/Laptops

- Linux Desktop Flavours

- SL6 (2)
- CentOS 7 (<10)
- Ubuntu 18.04 - (~330)



- Desktop Environment

- Gnome



Building a desktop is fully automated

- Step 1: Register the MAC address and the hostname in the asset management
- Step 2: Add the hostname to the Cobbler
- Step 3: Boot

Step 1: Asset Management System

- Register MAC address and hostname
- Network subnet
- Operating system
 - Linux: Ubuntu (CPLXConfig3,UEFI)
- Location
- Owner
- Serial Number
- Warranty
- etc...

Department of Physics



IT Network Registration

- Add Registration
- Delete Registration
- Modify Registration

Registering a device on the network

MAC Address	<input type="text"/>
System Name	<input type="text"/>
Location	<input type="text"/>
Main User	<input type="text"/>
Owner	<input type="text"/>
Operating System	Linux: Ubuntu 16.04 (CPLXConfig3, UEFI) ▾
Network	172.18.0.0 ▾
IP Address	<input type="text"/>
Serial Number	<input type="text"/>
Appliance Number	<input type="text"/>
Purchase Order Number	<input type="text"/>
Purchase Date <small>(suggest dd/mm/yyyy format)</small>	<input type="text"/>
Warranty Type	<input type="text"/>
Warranty Expiry:	<input type="text"/>
Additional Comments <small>(eg. Name/Institution)</small>	<input type="text"/>
<input type="button" value="Submit Request"/>	

Wait for email confirmation

Step 2: System Provisioning - Cobbler

- Register the hostname in cobbler
 - Follow naming convention, pplxdt`nnn`
- Cobbler manages:
 - ISO images
 - PXE boot
 - tftp repo
- Cobbler only installs vanilla OS plus:
 - Deploys ssh keys to allow for node administration
 - Sets node's facts (role, sub department, and research group)
 - `query --ksmeta options from Cobbler`
 - Installs Puppet agent and repo
 - Runs Cobbler triggers for:
 - `copy hostkeys`
 - `copy puppet certs`
 - Clean up
 - `disable netboot`

Any other configuration aspect is handled by Puppet.

Remote reinstall

- Cobbler Koan - “kickstart over a network”
- When koan is run,
 - modifies the grub to netboot
 - it requests install information from the cobbler server
 - it then kicks off installations

koan works ~90% of the time

Two Fortinet 601E 10G/bit firewalls.

- to be installed and configured

Ten Dell PowerEdge C6420

- few issues when installing



GridPP

UK Computing for Particle Physics

Questions?