

Lustre at Harvard

Teresa Kaltz, PhD Research Computing

Faculty of Arts and Sciences

- FAS governs Harvard College, Graduate School of Arts and Sciences, School of Engineering and Applied Sciences, and Continuing Education
- Over 1000 faculty





Research Computing Community

- Some write sophisticated parallel codes that scale to thousands of CPU's
- Others use packaged applications and treat HPC resources as black-box



Early universe simulation



Varied Resource Requirements

- Throughput engines on a grid
- Analysis of instrumentation data
 - Telescopes
 - Microarrays
- >100,000 files in a single directory!



ATLAS grid jobs simulating events at the LHC (CERN)



Consolidation of HPC in FAS

- Created "Research Computing" group in 2008
- Faculty purchase hardware with grant \$\$
- Research Computing (RC) provides:
 - HVAC, network, data center space
 - Systems design and administration
 - Software and algorithm/programming support
 - PhD level domain expertise
 - \$/TB scale-out tiered storage model
 - Differential price for backed up vs non-backed up
 - Currently no explicit SLA for uptime or performance



First RC Cluster "Odyssey" (April 2008)





Odyssey Compute Cluster

- 512 Dell dual socket blades
- 2.33 Ghz quad core E5410 (Harpertown)
- 8 x 4 GB 667 MHz FB-DIMM per blade
- Mellanox ConnectX DDR mezz card
- Two IB switches per chassis
- Two Cisco 7024 root switches in redundant configuration
- #60 on Top500 in June 2008, #80 Nov 2008



Research Computing Operations

- Started in 2008 with 5 people with root privs
- Grown to eight as of January!
- World class organization
 - 6 science/engineering PhD's
 - Corporate, academic and research institute backgrounds
 - Leverage FAS IT networking, infrastructure, unix teams
- How are we doing?



First the Good News...

Users are very happy with the compute resources!!





Compute Resource Details

- Original Odyssey Infiniband cluster
- Cerberus Infiniband cluster
 - Same config as Odyssey but with 192 blades
 - Could not expand Odyssey IB fabric due to space constraints
- IB/GbE Clusters, GPGPU, fat SMP nodes



TECHNOLOGY







Software Stack and Environment

- RH 5.1-based image
- kernel 2.6.18-92.1.17.el5
- OFED 1.3.1
- LSF
- Over 150 applications, tools and libraries available
- Lustre
 - 1.6.6 on servers
 - 1.6.7 on clients (patchless)



Research Computing Infrastructure

Datacenters: 2 On Campus 1 Downtown Boston





But Success Has Its Price

Storage has been a real challenge...

- Rapid growth has prevented strategic planning until recently
- Dynamic requirements from research community
- Constant evolution requires extreme flexibility in our storage architecture



-l-u-s-t-r-e-

- Obvious choice for up-and-coming HPC site like Harvard
- Established in HPC community
- Open source / open storage model
- Single namespace for very large capacity
- High performance / low cost
- Command line/scripting perfectly ok (who needs a GUI?)



Lustre at Harvard

- Considered professional services option but could not justify the cost
 - RC has limited operational budget
 - RC to provide \$/TB advantage
- Decided against Lustre appliance as we needed to develop in-house expertise to provide the most flexibility moving forward
- Although we had no filesystem specialist we had a high quality team



Capacity Filesystem: circe

- Deployed August 2008
- General user filesystem ("free")
- Located on campus (remote from IB clusters)
- Bonded 4GbE
- Increase capacity opportunistically •
 - Expect a mix of h/w types
 - Grow from 150 TB to petabyte?
- No failover/redundancy
- No backup service





Storage Building Block: Sun x4500



- 48 x 1TB SATA disks w/6 Marvell controllers
- 2 dual-core AMD Opterons w/16GB RAM
- 4 x 1Gb bonded interfaces
- Some have 1 SDR Mellanox HCA (PCI-X)



Metadata Hardware Details

- MGS/MDS
 - Dell dual socket 2950
 - Quad core E5410
 - 16GB RAM
- MDT
 - Dell MD3000 storage array
 - 15k 146GB SAS drives





Evolution of Circe Architecture



TÅS

INFORMATION TECHNOLOGY



Capability Filesystem: aega

- Deployed March 2009
- Single research group
- Static configuration at 250 TB raw
- No quotas/backups/failover
- Located in downtown Boston with Odyssey/Cerebus clusters (rented space)
- Mount over IB to Odyssey
 - Mount GbE on other systems



Aega Architecture (167 TB)





Some Aega Performance Results

Rates were measured during production Write likely bottlenecked by software RAID 5

	# Clients	Write	Read
Infiniband	32	2.1 GB/s	2.6 GB/s
Benchmark: IOR Default striping		blocksize: 32GB xfersize: 256MB	



Harvard's Lustre Experience

- We love Lustre for...
 - Price/performance
 - Single namespace
 - Ease of creating and expanding filesystem
- But we have had some heartburn...
 - Stability Problems
 - Administration Difficulties

What follows is a selection of the stability problems Harvard has experienced



Duplicate Inode Throwing LBUG

- Bugzilla 17485
- After client eviction ost stripe object is assigned to different files
- Recently got re-opened!

Dec 11 21:58:58 kernel: LustreError:

19050:0:(osc_request.c:2872:osc_set_data_with_check()) ASSERTION(old_inode->i_state & I_FREEING) failed:Found existing inode ffff8100559ff6f8/120835506/1537734185 state 0 in lock: setting data to ffff8100559ff3b8/120835505/1537734184 Dec 11 21:58:58 kernel: LustreError: 19050:0:(osc_request.c:2872:osc_set_data_with_check()) LBUG



Unprovoked Client Evictions

- We were seeing these frequently on 1.5.6.1
- Bugzilla 17379: stack overflow on patchless client
- Bugzilla 17631: Unexpectedly long timeout in ptlrpc_unregister_bulk()
- Fortunately upgrading to 1.6.7 on clients (not servers!) has improved the issue



OSS/OST Hangs

- This turned out to be a problem with the bonding driver in Redhat
- Redhat Bug 457300 hang in ad_rx_machine due to second attempt to lock spin_lock
- Software RAID rebuild times prohibitive
- Rebuilt kernel module with patch since public Lustre roadmaps did not indicate when patched kernel would be supported



Administration Issues

- Best practices for backups
 - Currently using LVM snapshot for MDT
- Enabling and using quota
- Determining disk usage for large filesystems
- Monitoring and alerts
- Understanding Lustre error messages
- Filesystem diagnostics and troubleshooting



Troubleshooting Lustre Issues

- Determining whether a problem is a known issue can be difficult
- Searching lustre-discuss is extremely helpful
- Searching bugzilla is not always so useful
 - Bugs often do not have clear summaries of symptoms (error message strings we can search for)
 - Bugs are marked as duplicates of existing bugs that are private!



Not helpful...

	/bugzilla.lustre.org/show_bug.cgi?id=16427	☆ ▼) ° (G▼ Google	٩)	
st Visited 👻 Getting Started Latest Headlines 🔊	BMC Remedy Gmail			
Bugzilla – Access Denied		Report bugsila bugs lo <u>lustra-bugsila-admin@sun.com</u>		
ome I <u>New</u> I <u>Search</u> I Fin	I Reports I Requests I New Account I Log In			
You are not authorized to a an account with the approp	access bug #16427. To see this bug, you mu priate permissions.	ust first log in to Creating an acco	entation	
ease press Back and try again.				
ctions: _{Home} <u>New</u> <u>Search</u>	Find Reports Requests New Account Log	<u>In</u>		
e		bugzilla	lustre.org 🔒	

VERI

TAS

INFORMATION TECHNOLOGY

Final Thoughts From a Newbie Site

- Lustre is a great technology but requires lot of sweat equity
 - Does Lustre require a dedicated FTE?
- Lustre knowledge transfer model
 - − Source code \rightarrow free
 - 1-800-HELPME → \$\$
 - − Best practices \rightarrow ??

Harvard is currently evaluating how to best utilize Lustre in our storage environment



Architectural and Technical Information

 Harvard Research Computing Web Site http://hptc.fas.harvard.edu/



Acknowledgements

- Larry Levine
- Nancy Maul
- Jeremy Bloxham
- Lars Hernquist
- Jacques Laflamme
- David Sullivan
- Paul Arena
- Brian Mantanuto
- Nick Amento

- James Cuff
- Suvendra Dutta
- Jerry Lotto
- Mike Ethier
- Ford Fay
- Chris Walker
- John Brunelle

