



Integrating Array Management into Lustre

System Fabric Works

Roger Ronald / Kevin Moran

LUG 2014



Topics

- Lustre Adoption Barriers
- Wider Lustre Adoption
- Integration of Lustre with Other Components and Functions
 - Integrating NetApp's E-Series Disk Arrays / Santricity with **Intel® Enterprise Edition for Lustre** (Intel EE Lustre)
 - Screen Shots
- Summary



Lustre Adoption Barriers

- Until recently, Lustre has been solely focused at supercomputers and large-scale tech-savvy enterprise customers
 - Lustre performance was wasted on smaller configurations
 - Lustre deployment/operation cost was high
 - Lustre feature set was missing key capabilities for general usage
 - Native clients for “other than Linux” operating systems
 - Effective support for tiered storage (HSM, SSDs, tape)
 - Lustre was not simple to deploy and operate
 - Command line biased interface
 - Experts required
 - Lustre wasn't a fully integrated storage solution
 - Lustre uses storage hardware, Linux servers for OSS and MDS functions, and switches
 - Separate tools needed to monitor each component

Lustre has historically been a very-high performance niche file system



Wider Lustre Adoption

- Recent developments have added to Lustre's functionality
 - New features (notably HSM) create good reasons to deploy Lustre even when performance requirements are modest
- New technologies and tools are making Lustre easier/less expensive to deploy and operate
 - Virtualization makes it practical to build smaller Lustre configurations
 - New Lustre releases (notably Intel's Enterprise Edition) are focused on improving ease of Lustre deployment and usage
 - Possible to use a GUI for most actions
 - Vendors are integrating monitoring of other components and functions into Lustre tools
- More systems and applications require high performance file systems

The obvious barriers to Lustre adoption are being addressed



Integration of Lustre with Other Components and Functions

- “Single pane of glass” is beneficial for system monitoring
 - Key requirement is to alert and inform system operator
 - Errors are easier to notice using a single tool
- “Good” Lustre operational monitoring doesn’t need to replicate every low level capability for components
 - Experts on individual components (storage, server, network) can use their own tools to address problems
 - Goal is not to replace all tools for all uses - goal is to provide sufficient visibility to:
 - Alert user
 - Assist in determining when additional steps are needed
 - Assist in showing where to start/who to call

“Single pane of glass” provides a better operations perspective



Integrating Santricity with Intel EE Lustre

- NetApp E-Series storage has a long history with large Lustre deployments (e.g. LLNL)
 - Management/monitoring of NetApp arrays traditionally performed using NetApp's Santricity software
- NetApp now provides an Intel EE Lustre plug-in for monitoring the health and performance of NetApp E-Series storage
 - NetApp originally built the plug-in for Whamcloud Chroma
 - NetApp funded SFW to update the plug-in for Intel EE Lustre
 - SFW work just completed, plug-in available
 - NetApp plug-in monitors the health/performance of storage arrays and volumes (screen shots follow)

NetApp E-Series supports an Intel EE Lustre management plug-in



NetApp Intel EE Plug-in Screen Shots

intel | Dashboard | Configuration | Alerts | History | Logs | Help | 2.0.1.0 | admin | Account | Logout

Storage

Storage Resources

+ Add Storage Device

Filter: netapp_storage_plugin-StorageArray

Show 25 entries

Name	Address 1	Address 2	
Lustre-SFW-1	netapp	netapp	✓

Showing 1 to 1 of 1 entries



NetApp Intel EE Plug-in Screen Shots

Storage Resources

+ Add Storage Device

Filter:

Show 25 entries

Name	Size	Filesystem type	Raid type	Cache	State	Primary controller	Volume group	Storage array	Wwn
OST1	2.0TB	None	DDP	read write	normal	Controller 2	Disk Pool 1	Lustre-SFW-1	60080E50001F6B6000001403527A17
OST2	2.0TB	None	DDP	read write	normal	Controller 2	Disk Pool 1	Lustre-SFW-1	60080E50001F6B6000001404527A17
MDT	100.0GB	None	RAID 1	read write	normal	Controller 2	META	Lustre-SFW-1	60080E50001F6B6000001416527AA7
MGT	100.0GB	None	RAID 1	read write	normal	Controller 2	META	Lustre-SFW-1	60080E50001F6B6000001418527AA7

Showing 1 to 4 of 4 entries



NetApp Intel EE Plug-in Screen Shots

The screenshot displays the NetApp Intel EE Plug-in interface. The top navigation bar includes 'Dashboard', 'Configuration', 'Alerts', 'History', 'Logs', and 'Help'. The user is logged in as 'admin' with version '2.0.1.0'. The main content area is titled 'Storage' and shows 'Storage Resources' with a table of storage devices. A 'Volume' configuration window is open for 'OST1', showing details such as Cache (read|write), Filesystem type (None), Primary controller (Controller 2), Raid type (DDP), Size (2.0TB), State (normal), Storage array (Lustre-SFW-1), Volume group (Disk_Pool_1), and Wwn (60080E50001F6B6000001403527A1755). Performance graphs for 'read_iops', 'write_throughput', and 'write_iops' are also visible.

Storage Resources Table:

Name	Size	Filesystem type	Raid type
OST1	2.0TB	None	DDP
OST2	2.0TB	None	DDP
MDT	100.0GB	None	DDP
MGT	100.0GB	None	DDP

Volume Configuration (OST1):

Cache	read write
Filesystem type	None
Primary controller	Controller 2
Raid type	DDP
Size	2.0TB
State	normal
Storage array	Lustre-SFW-1
Volume group	Disk_Pool_1
Wwn	60080E50001F6B6000001403527A1755

Performance Graphs:

- read_iops:** Line graph showing read operations per second over time. Y-axis ranges from 0 to 0.05. X-axis shows times 22:50 and 22:52.
- write_throughput:** Line graph showing write throughput in KB/s over time. Y-axis ranges from 0 to 1,500. X-axis shows times 22:50 and 22:52.
- write_iops:** Line graph showing write operations per second over time. Y-axis shows a value of 0.4.



NetApp Intel EE Plug-in Screen Shots

Storage Resources

+ Add Storage Device

Filter:

Show entries

Name	Storage array	Used space	Capacity	Raid type	Disk count	State	Wwn	
Disk Pool 1	Lustre-SFW-1	4.0TB	6.3TB	DDP	20	complete	60080E50001F6B6000001402527A16E7	✓
META	Lustre-SFW-1	200.0GB	465.3GB	RAID 1	2	complete	60080E50001F6B6000001415527AA6FA	✓

Showing 1 to 2 of 2 entries



NetApp Intel EE Plug-in Screen Shots

Intel Dashboard Configuration Alerts History Logs Help 2.0.1.0 admin Account Logout

+ Add Storage Device

Filter: netapp_storage_plugin-DiskDrive

Show 25 entries

Name	Storage array	Tray	Drawer	Slot	Capacity	Spindle speed	State	Media type	Interface type	Volume group	Manufactu
99_21	Lustre-SFW-1	99	-	21	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_22	Lustre-SFW-1	99	-	22	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_8	Lustre-SFW-1	99	-	8	465.3GB	7200	optimal	hdd	SAS		SEAGATE
99_14	Lustre-SFW-1	99	-	14	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_19	Lustre-SFW-1	99	-	19	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_10	Lustre-SFW-1	99	-	10	465.3GB	7200	optimal	hdd	SAS	META	SEAGATE
99_24	Lustre-SFW-1	99	-	24	465.3GB	7200	optimal	hdd	SAS		SEAGATE
99_18	Lustre-SFW-1	99	-	18	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_4	Lustre-SFW-1	99	-	4	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_20	Lustre-SFW-1	99	-	20	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_2	Lustre-SFW-1	99	-	2	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_1	Lustre-SFW-1	99	-	1	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_6	Lustre-SFW-1	99	-	6	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_16	Lustre-SFW-1	99	-	16	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_12	Lustre-SFW-1	99	-	12	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_13	Lustre-SFW-1	99	-	13	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_17	Lustre-SFW-1	99	-	17	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_5	Lustre-SFW-1	99	-	5	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_15	Lustre-SFW-1	99	-	15	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_23	Lustre-SFW-1	99	-	23	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_11	Lustre-SFW-1	99	-	11	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_9	Lustre-SFW-1	99	-	9	465.3GB	7200	optimal	hdd	SAS	META	SEAGATE
99_3	Lustre-SFW-1	99	-	3	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE
99_7	Lustre-SFW-1	99	-	7	465.3GB	7200	optimal	hdd	SAS	Disk_Pool_1	SEAGATE



NetApp Intel EE Plug-in Screen Shots

intel | Dashboard | Configuration | Alerts | History | Logs | Help | 2.0.1.0 | admin | Account | Logout

Storage

Storage Resources

+ Add Storage Device

Filter: netapp_storage_plugin-HostInterface

Show 25 entries

Name	Storage array	Controller	Port/Channel	Type	Max Rate (Gbps)	Current Rate (Gbps)	State	
HostInterface (Fibre Channel, 1)	Lustre-SFW-1	Controller 2	1	Fibre Channel	8.0	8.0	down	✓
HostInterface (Fibre Channel, 2)	Lustre-SFW-1	Controller 2	2	Fibre Channel	8.0	8.0	down	✓
HostInterface (Fibre Channel, 3)	Lustre-SFW-1	Controller 2	3	Fibre Channel	8.0	8.0	down	✓
HostInterface (Fibre Channel, 4)	Lustre-SFW-1	Controller 2	4	Fibre Channel	8.0	8.0	down	✓
HostInterface (SAS, 5)	Lustre-SFW-1	Controller 2	5	SAS	6.0	-	optimal	✓
HostInterface (SAS, 6)	Lustre-SFW-1	Controller 2	6	SAS	6.0	-	optimal	✓
HostInterface (SAS, 7)	Lustre-SFW-1	Controller 2	7	SAS	6.0	-	optimal	✓
HostInterface (SAS, 8)	Lustre-SFW-1	Controller 2	8	SAS	6.0	-	optimal	✓

Showing 1 to 8 of 8 entries



NetApp Intel EE Plug-in Screen Shots

The screenshot displays the NetApp Intel EE Plug-in interface. At the top, there is a navigation bar with the Intel logo and menu items: Dashboard, Configuration, Alerts (highlighted), History, Logs, and Help. A red indicator light is visible next to the Help icon.

A "Notifications" window is open on the left side, showing a list of alerts. The first alert is "Battery removed (Battery Battery (99, 1,...))" with a timestamp of "2014/03/17 22:45" and a green checkmark icon.

The main content area shows a table of alerts:

Entity	Message
Battery (99, 1, 1)	Battery removed (Battery Battery (99, 1, 1))
Finished 2014/03/18 01:29	Entity Controller 2 Message Controller Firmware Level is unsupported (Controller Controller 2)

Below the table, there are expandable sections for "History" and "Commands".



NetApp Intel EE Plug-in Screen Shots

The screenshot displays the NetApp Intel EE Plug-in Configuration interface. The top navigation bar includes the Intel logo, a search icon, and menu items for Dashboard, Configuration (selected), Alerts, History, Logs, and Help. The user is logged in as 'admin' with options for Account and Logout. The main navigation pane on the left shows Servers, Volumes, Power Control, MGTs, File Systems, Storage (selected), and Users. The 'Storage' section is active, showing 'Storage Resources' with an 'Add Storage Device' button. A filter is applied to 'netapp_storage_plugin-Battery'. A table lists two battery entries:

Name	Tray	Slot	Position	Storage array	Manufacture date	Vendor	Serial number	Battery type	State	
Battery (99, 1, 1)	99	1	1	Lustre-SFW-1	01-01-70			dual, individual FRUs	removed	!
Battery (99, 2, 1)	99	2	1	Lustre-SFW-1	01-01-11	VN LSI	SN SY10200107A1	dual, individual FRUs	optimal	✓

Showing 1 to 2 of 2 entries



Summary

- Lustre historically has been limited to the niche role of a very high-performance file system
- Lustre will expand into other roles as barriers to adoption are addressed and functionality is added
- Intel EE for Lustre Plug-ins address a significant adoption barrier by improving ease of use
 - SFW has implemented a NetApp plug-in for Intel EE Lustre
 - Additional plug-ins for storage, networks, and servers are encouraged

Lustre plug-ins connect deep component-level reporting with the Intel EE Lustre dashboard for consolidated, smart monitoring