

whamcloud

The logo for Whamcloud features the word "whamcloud" in a bold, dark grey, lowercase sans-serif font. A thick blue horizontal line underlines the text. On the right side, a blue graphic element consisting of two overlapping curved lines forms a stylized shape that resembles a cloud or a drop, partially overlapping the end of the word and the underline.

April 12th 2011

Lustre Community Releases

- Peter Jones
Support Manager
Whamcloud, Inc.

Agenda

- Rationale
- Lustre 1.8.x Releases
- Lustre 2.1
- Future 2.x Releases
- Summary

Rationale

- Late in 2010 Oracle changes strategy for Lustre
 - Lustre was placed into sustaining mode
 - All work suspended on 2.x
 - Initially unclear whether there would be further 1.8.x releases
- Lustre community needs stable 2.x release
 - All future Lustre feature development based on 2.x codeline
 - Lustre 2.0 went GA July 2010 but performance comparison to 1.8.x was deferred
 - Intention of 2.1 is to provide a stable 2.x release with performance that matches or exceeds 1.8.x in common usage patterns

Lustre 1.8.x Releases

- Most production Lustre sites running 1.8.x
 - 1.8.x has been available for some time and is quite stable
- Oracle has now clarified that it will continue to make Lustre 1.8.x releases
- No clear reason for the community to produce 1.8.x releases
 - Whamcloud is continuing to push fixes upstream
 - Whamcloud does have the capability to produce 1.8.x releases if needed at some future point

Lustre 2.1 – the Problem

- All Oracle work on 2.1 has ceased
 - There have been no checkins to master since the turn of the year
- How to fill the vacuum?
 - Who would do the work?
 - Where would the code be checked in?
 - Who would do the testing?

Lustre 2.1 – the Solution

- Some things are the same
 - Many of the same engineers performing the work
 - Community involvement in testing
- Some significant differences
 - Whamcloud had much of the necessary infrastructure in place to create Lustre releases and so volunteered to host 2.1 community release (git, JIRA, gerrit, jenkins, maloo etc)
 - 2.1 Release meetings open to anyone
 - Whamcloud contributor agreement means that **no single organization will ever hold Lustre copyright again**
 - All three community groups (EOFS, HPCFS and OpenSFS) are in support of this approach

Lustre 2.1 – the Scope

- RHEL6 Server and Client support
- Async journal commits by default
 - Added in 1.8.2; turned off by default
 - Already used in production at many sites (LLNL, ORNL, DDN sites)
- Ext4 by default
- 2.x performance to match/exceed 1.8.x
 - SMP Scaling\IO performance

Lustre 2.1 – the Logistics

- Adopted proven management model for Lustre releases
 - **Peter Jones** release manager, co-ordinating efforts
 - **Oleg Drokin** technical lead, acting as gatekeeper
- Broad community interest in release
 - 25 different organizations registered on mailing list (at least)
 - 11 different organizations represented at meetings
 - 3 different organizations submitted patches (LLNL, ORNL, Xyratex)
 - 4 different organizations offered to help with testing (Bull, Cray, LLNL, ORNL)

Lustre 2.1 – the Current Status

- Status is changing rapidly as we converge on an RC
 - You can view the current # of blockers using the “Lustre 2.1 blockers” filter found under Manage Filters – Popular in JIRA
- Get more directly involved
 - Mailing list <http://groups.google.com/group/lustre-21>
 - Weekly calls on Tuesday 9:30am PT 866-914-3976 534986#

Future 2.x Releases

- Lustre 2.1 is relatively simple
 - The scope of the release was already defined and most of the work was done
- It will be much harder to manage release content for future releases across multiple stakeholder groups
 - Need to find a workable long-term model
- Whamcloud will be producing future 2.x releases for its customers
 - Core Lustre code will be open to all and available to any other releases

Summary

- Oracle will continue to produce Lustre 1.8.x releases
- Lustre 2.1 due out this summer
- Lustre 2.x releases TBD



Thank You

- Peter Jones
Support Manager
Whamcloud, Inc.