



PATH FORWARD FOR LUSTRE COMMUNITY

HPCFS AT ORLANDO LUG 2011

BILL BOAS



SFW AND LUSTRE

- ▶ **Focus on Engineering Services, I/O Performance, Integration**
 - Over 50 assignments in last 8 years
 - Integration of OFED for Range of Storage and File System Products
 - Custom System Integration, “Green Beret” services
- ▶ **Delivering Products and Solutions in HPC, Virtualization, Private Clouds, Enterprise Data Centers with 10/40 GbE, InfiniBand**
- ▶ **Customers are Leaders in their Markets – Research, Defense, Finance, Service Provider**
- ▶ **Focus is Highly Efficient Computing, Low Latency Networks and Storage, Lustre**
- ▶ **Pro-Buono work - Officers of OpenFabrics Alliance, InfiniBand Trade Association, High Performance Cluster File System Foundation**
- ▶ **Lustre and Related Experience**
 - Managed Lustre Program at LLNL 2002-6, Consultants to NRL JCTD 2006-9
 - Engineering Srvcs for CFS, Cray, DDN, LBL, LSI, Microsoft, MLNX, NGA, NRL, Oracle, Sun
 - Deployment Srvcs at BOM Australia, Hitachi, LANL, NRL, Raytheon, Sandia, Sun

The Lustre Market, Potential Resources exist

- ▶ Customers/users and vendors need to believe there are a long term, sufficient technical resources, funded and equipped to sustain and evolve Lustre so there is a viable market
- ▶ Today the market includes the fabric/network, OSS, disk controllers and actual storage (disk+flash+HSM hardware)
- ▶ Assuming 500 user sites of Lustre spending \$2M per year on this hardware is \$1Billion
- ▶ A 1% share of that for on-going Lustre development yields \$10M/year within their budgets that could be available if customers asked their hardware vendors to allocate funds
- ▶ Some of the h/w vendors: Bull, Cray, Dell, DDN, Fujitsu, HP, Hitachi, LSI, MLNX, NEC, Oracle, QLogic, RAID, SGI, Terascale, Xiotech, Xyratex – there are many others.....
- ▶ Note that no software or support revenue is counted here
- ▶ Can customers start expressing this case to their h/w vendors?
- ▶ No consideration of what customers are spending themselves

LUSTRE-SOURCE.ORG WE CAN ALL AGREE ON

- ▶ Single set of Working Groups collaborating internationally to make technical decisions on direction, master feature content and tree
- ▶ Technical Respect the basis for gatekeepers, release and bug management
- ▶ Encouraging participation from all vendors and all users at any level
- ▶ Community Repository at neutral site, Core Tree, Bugs DB
- ▶ Build on 2.1 work being driven by Whamcloud, sites and vendors can specialize from the core and contribute back select code
- ▶ Contributors keep copyright in code, binds the code together
- ▶ Collaborative testing amongst all in community
- ▶ Commitment to GPL v2 license, ensures long term
- ▶ Community agrees not to assert any patent claim in the code
- ▶ A single client implementation that interoperates across all networks and OSSs

Working Items We Can All Agree On

- ▶ What is in this release
- ▶ How is a release candidate chosen, who gatekeeps and who release manages
- ▶ What's in the next release
- ▶ How the core tree is set up
- ▶ How the core tree is released and advances
- ▶ How sites and vendors use the core tree plus to meet their needs
- ▶ What is the bug fixing and patching process
- ▶ How do we specify stability and performance criteria
- ▶ Building a single repository in neutral territory
- ▶ Building a collaborative testing process
- ▶ At the margin, vendors and sites have differing needs
- ▶ Lustre needs “marketing and promotion” to prospects and customers

If we have these substantial areas of agreement

Can we project these to the market place

And

Implement them

Is there a Lustre Market?

- ▶ Yes, lets consider who are the customers, suppliers and size
- ▶ Customers are those that use Lustre, how many are there 3-500
- ▶ The market has Suppliers of Hardware, Software, Support Services and Consulting including OEMs, System Integrators and resellers
- ▶ Suppliers are Bull, Cray, Dawning, DDN, Dell, Fujitsu, HP, IBM, LSI, NEC, Seagate, SGI, SFW, Terascala,
- ▶ How many petabytes of storage space for Lustre do customers deploy per year? 100-200 ?
- ▶ What is the street price, incl. OSS, Controllers., network, etc. for a petabyte? \$1MM
- ▶ Leads to an estimate of \$1-200MM per year for storage hardware
- ▶ Customers have average of ½ FTE, FTE costs \$200k with burden
- ▶ Estimate customers are spending \$3-50MM self support

Leaders in High Performance, Low Latency Data Motion

“Fabric Computing that Works”

Trusted Strategic Advisors

Reliable Implementation Partners