



Lustre Beyond HPC

Presented to the Lustre* User Group
Beijing– October 2013

Brent Gorda
General Manager
High Performance Data Division, Intel Corporation



Agenda

- From Whamcloud to Intel
- Today's Storage Challenges
- HPC and Big Data
- Conclusions, and opportunities



From Whamcloud* to Intel

Whamcloud Started July 16, 2010

- Brent Gorda – CEO
- Eric Barton – CTO

Acquired by Intel in July 2012

- Now Intel's High Performance Data Division
- **Same** world class team...
- Passionately focused on moving Lustre* software forward
- Equipped with more resources and opportunities...



Lustre* @ Intel: Today

Intel continues Whamcloud's pledge:


- We fully support OpenSFS and EOFS
- With backing of OpenSFS/EOFS: All Lustre releases
- Fully committed to open, collaborative development
- Delivered in real-time to the worldwide community
- We remain focused on breakthrough innovations
- Create features for HPC and big data workloads
- Intel encourages you to join and be active with Lustre



Lustre* @ Intel – Moving Lustre Forward


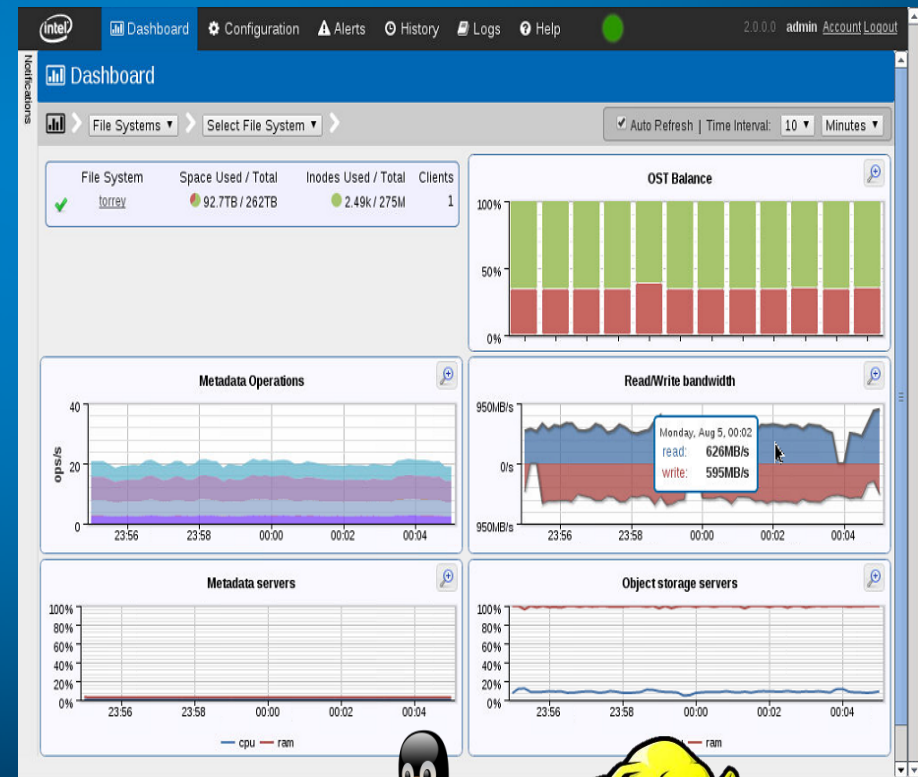
Creating an Enterprise Grade Ecosystem

- Partners are looking for a business-quality solution
- Intel Enterprise Edition for Lustre* Software
- Adds features such as Hadoop, AWS, IML, 24x7 support



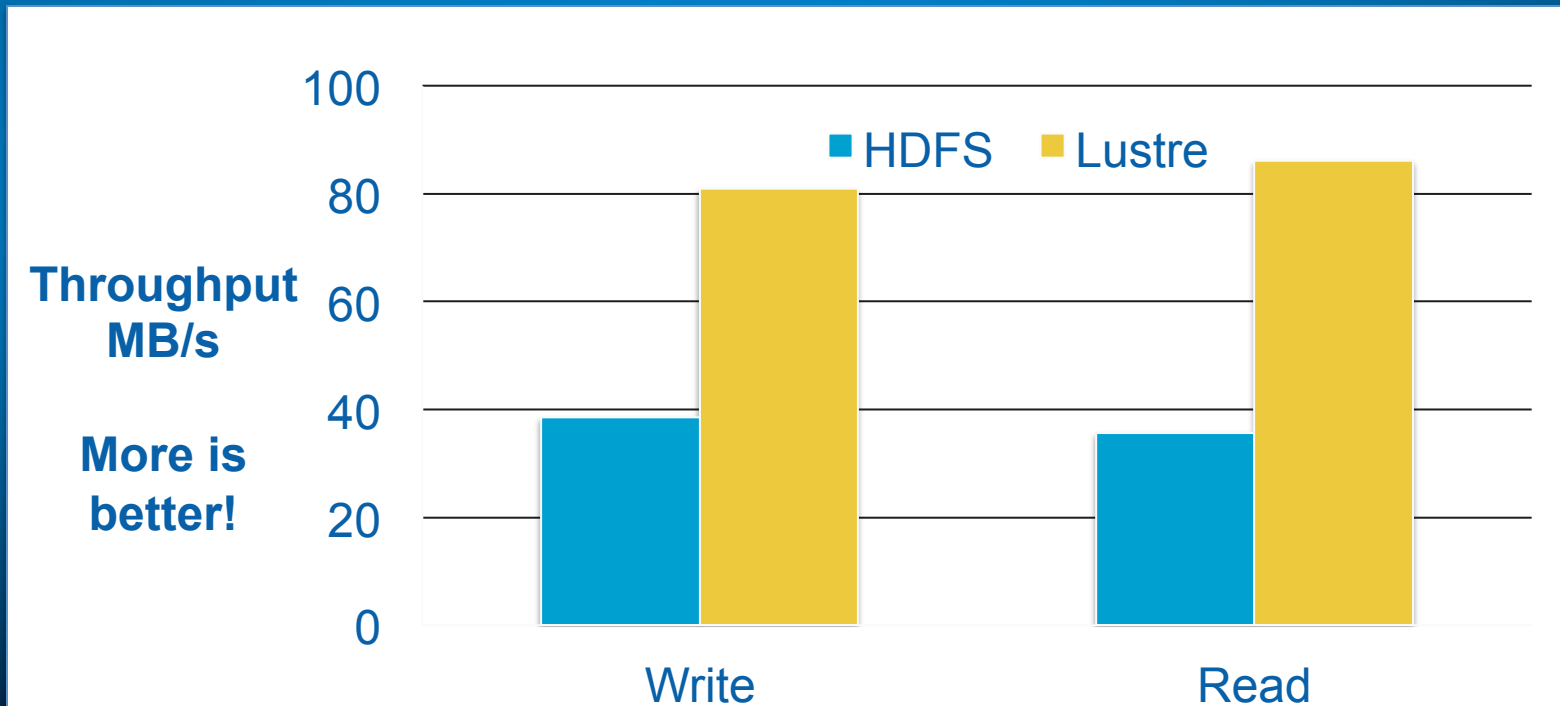
Storage for Data Intensive Applications Powered by Intel® Enterprise Edition for Lustre* software

- Production quality Lustre file system backed by Intel
- Intel® Manager for Lustre* UI
- Intel® Lustre* Adaptor for Apache® Hadoop
- 24x7 Enterprise support
- Professional services and training
- Available from OEM, integrator and reseller partners

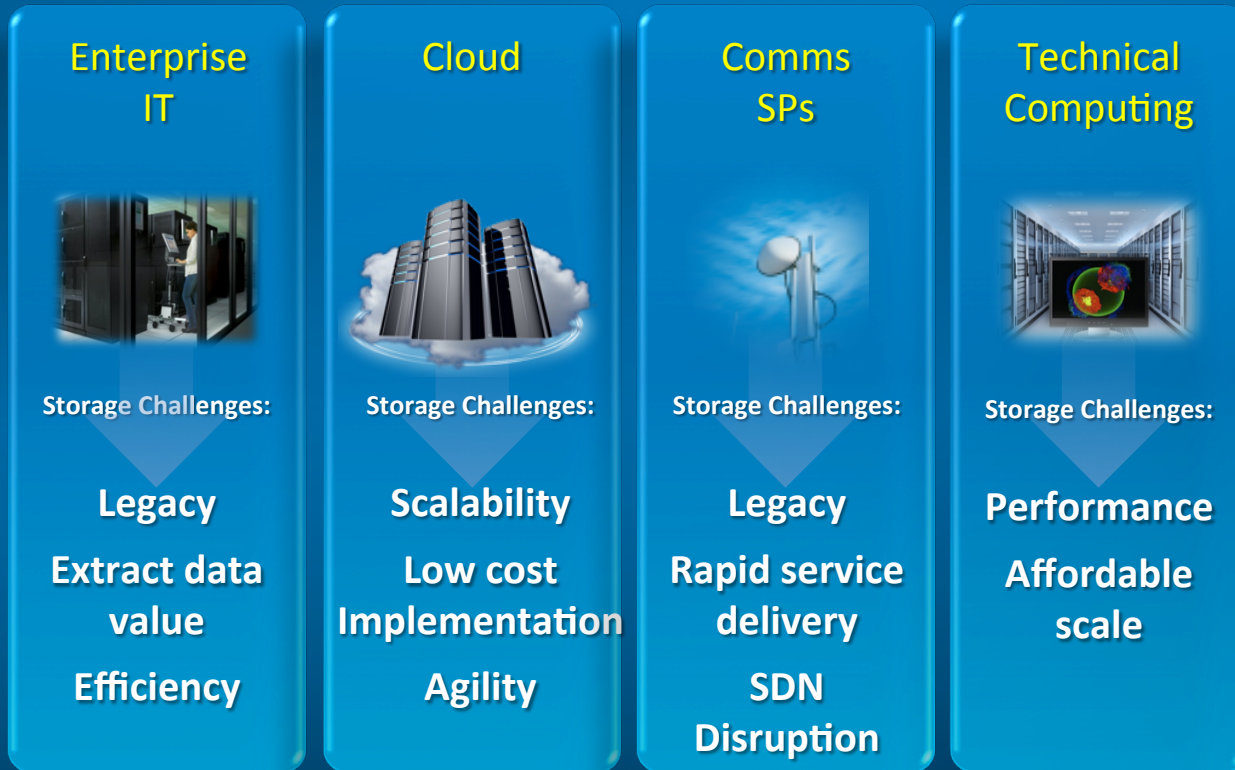


TestDFSIO Benchmark

- Tests the raw performance of a file system
- Write and read very large files (35G each) in parallel
- One mapper per file. Single reducer to collect stats.
- Embarrassingly parallel, does not test shuffle & sort



Storage Challenges Differ By Market



'Big Data' Workloads Are Reshaping Storage

80%

Of data is
unstructured

15X

Growth rate for
unstructured data

67%

Of HPC users run "Big
Data" applications

Parallel Storage for HPC and 'Big Data'

- Compute, Fabric & Storage Must Scale and be balanced
- Without adequate storage performance, compute performance goes unused
- Scalable storage software is critical
- Lustre is key



The Big Data Paradigm Change

Codes Based on Analytic Models

Well Known HPC Workloads

HPC Today

- ✓ Compute Focused
- ✓ Minimizes Data Movement
- ✓ I/O predominantly For Checkpoints
- ✓ Datasets are ~Petabytes
- ✓ Data is sampled or generated

Codes Based on Data Driven Models

New Big Data Workload

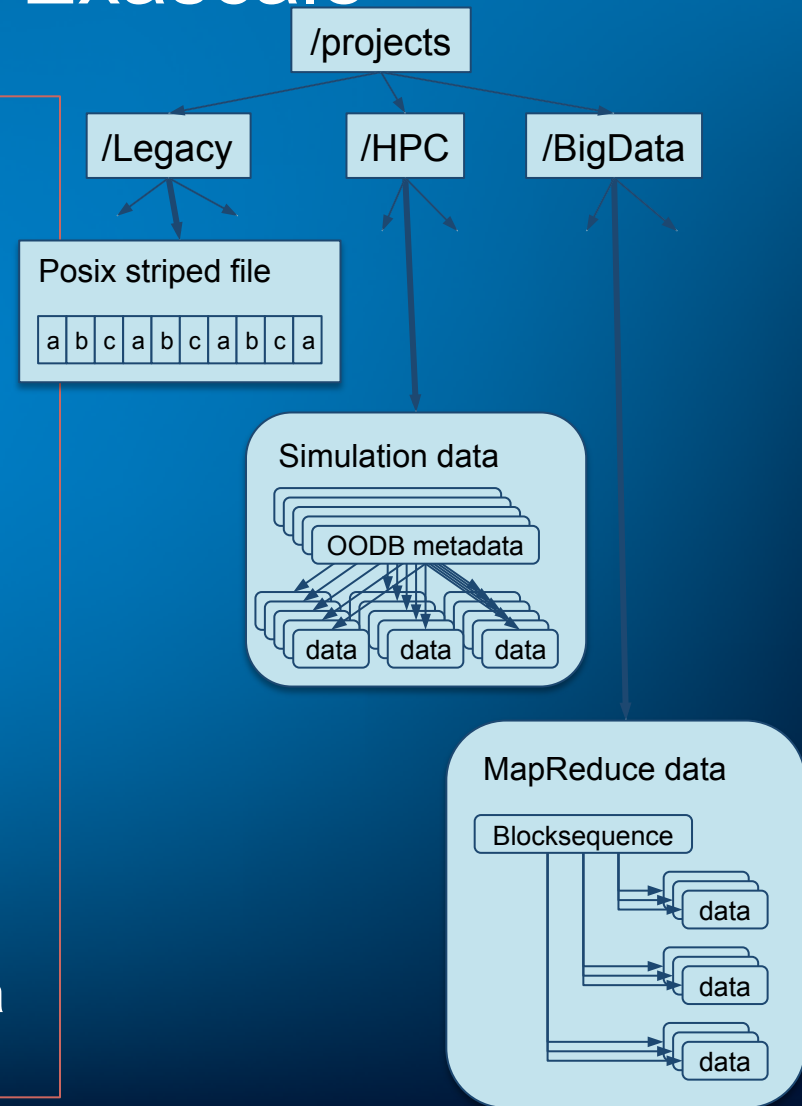
HPC Tomorrow

- ✓ I/O Focused
- ✓ Lots of Data Movement
- ✓ I/O predominantly for storing & retrieving data
- ✓ Datasets are ~100's of Petabytes
- ✓ All data is needed all the time

System Design Points Will Change!

Moving Lustre* Forward - Exascale


- Integrated I/O Stack
 - Epoch transaction model
 - Non-blocking scalable object I/O
- HDF5
 - High level application object I/O model
 - I/O forwarding
- I/O Dispatcher
 - Burst Buffer management
 - Impedance match application I/O performance to storage system capabilities
- DAOS
 - Conventional namespace for administration, security & accounting
 - DAOS container files for transactional, scalable, object IO
- Detailed technical sessions on agenda



Conclusions and Opportunities

- **Lustre* was purpose built for performance at scale**
 - The technology is mature, stable, fast and popular
 - Offers a stable, performance based, multi-vendor solution
- **Community participation with OpenSFS is vital**
 - As a community, we grow together
- **Scalable storage software is critical for maximum application performance**
 - Data is growing – the need / opportunity is there
 - Storage is not the problem. Data is the problem

⇒ **Intel is committed to Lustre – for HPC + Big Data**
Intel is moving Lustre forward



Software

Datacenter Software Division
High Performance Data Organization



Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel, Look Inside and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.
Copyright ©2013 Intel Corporation.