Xyratex Update

Michael K. Connolly
Partner and Alliances Development

xyratex
A Seagate Company

Is Now……

The Continued Power of Xyratex

Global Solutions Provider of High Quality

Data Storage Hardware, Software and Services

Broad data storage solutions expertise

Storage Media, Storage Platforms, Clustered File Systems

370 storage related patents

Solving complex technical data storage requirements

ClusterStor™ Scale-Out Data Storage Solutions

OneStor™ OEM Storage & Capital Equipment

xyratex
A Seagate Company

What Does this Mean?

- To Xyratex:
  - Why Did Seagate Purchase Xyratex?
  - Xyratex Brand and Structure
  - Significant Growth Opportunities
  - More Resources to Support and Serve Customers
  - Complimentary and Expanded Products and Solutions
  - Access to Emerging and Innovative Technologies

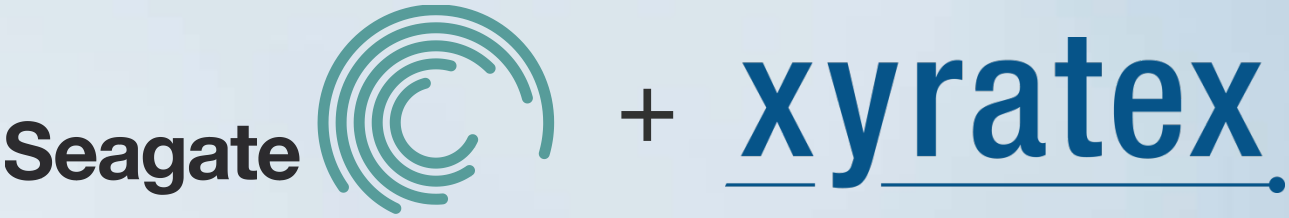

What Does this Mean?

- To the Community:
  - Continued Commitment to Support and Solution Development around Lustre®
  - Xyratex Renewed OpenSFS Promoter Level Membership
  - Seagate is Supportive of Community Participation
  - Focus on Collaboration to Yield the Highest Performance, Stability, Reliability, and Robustness of Lustre
  - Upcoming and Exciting Community Plans

xyratex
A Seagate Company

Some Events Completed and Planned in 2014

March | April | May | June

Rice Oil & Gas
ISUM
HPC AC
HPC OSL TES

Ft. Meade
HPC for Wall St.
GEOINT 2014
LUG 2014
Bio-IT 2014

CUG 2014

EAGE 2014
HPC-AC at ISC
ISC 2014
HPCS–Canada

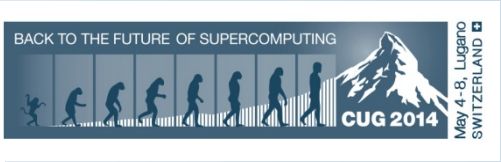

EAGE

ISUM 2014

geoint SYMPOSIUM TENTH ANNUAL

ISC INTERNATIONAL SUPERCOMPUTING CONFERENCE

HPC ADVISORY COUNCIL NETWORK OF EXPERTISE

HPCS 26 Years of Supercomputing

HPC OSL TES 2014
High Performance Computing and Open Source & Linux

xyratex
A Seagate Company

Growth and Accomplishments

- From 2013 through 2014:
  - Great Year for ClusterStor and Lustre
    - Awarded: Vendor to Watch!
    - Grew revenue significantly with many new customers in multiple verticals
    - Hired many new employees in various support and dev disciplines
    - Participated at vast number of community and industry events
    - Xyratex contributed Lustre Improvements in version 2.5:
      - 4 MB I/Os, Network Request Scheduler, Metadata performance + over 240 bugs fixed
  - New ClusterStor Solutions and Features
    - Launched ClusterStor 1500 and 9000 versions
      - New CS OS software release 1.4 with new multiple features
    - Planned integration of Lustre 2.5 within 2014
    - On the Horizon….?

xyratex
A Seagate Company

The New ClusterStor™ 9000

Engineered Solution for HPC and Big Data

- Now 50% faster than previous generation
  - Up to 9 GB/s per 5U-84 Scalable Storage Unit
  - ClusterStor™ OS version 1.4 with New Features
  - Leverages latest Lustre® capabilities

- Engineered solution optimized for
  - End-to-End Speed
  - Enterprise Bullet Proof Reliability
  - Clear Industry Leading Efficiency by Wide Margin

xyratex
A Seagate Company

ClusterStor OS Version 1.4

- New in Version 1.4:
  - GridRAID (formerly called PDRAID)
    - *available in mid 2014*
  - New and Improved Monitoring Dashboard
    - High level view into the entire storage
    - Node Status
    - File System Throughput
    - Inventory
    - Top System Statistics
  - The SSU+n feature, where the maximum value for n=3, whereby up to 3 Expansion Storage Units (ESUs) can be added to each SSU
  - NIS GUI Support - added GUI support for configuring NIS as an option for Lustre users.

Perspective on Traditional RAID

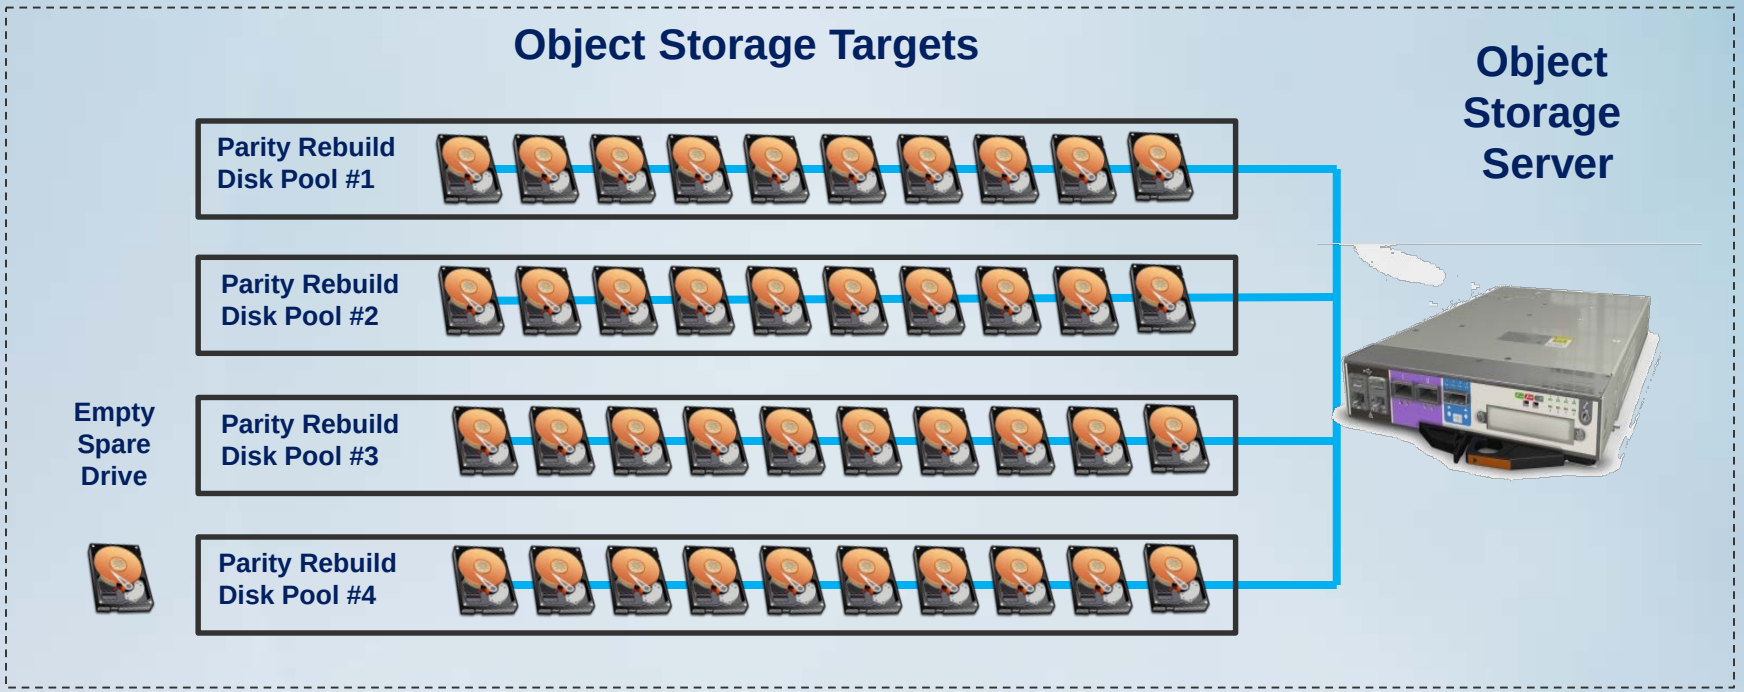

- Traditional RAID rebuild time increasingly long for high capacity drives
  - Faster data protection and recovery technology required

ClusterStor GridRAID

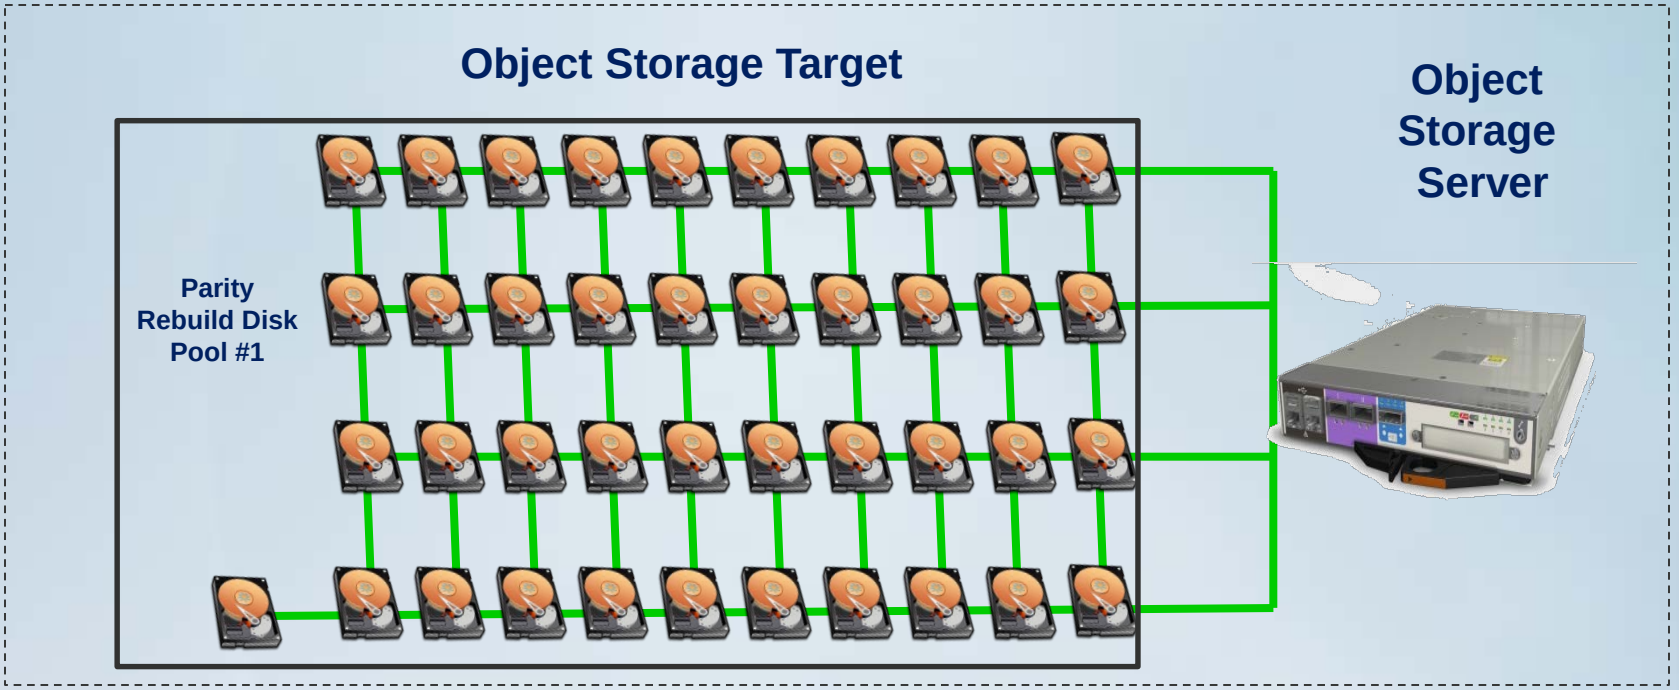

- ClusterStor GridRAID restores full data redundancy up to 400% faster than traditional RAID6 implementations
- Mandatory to effectively implement high capacity drives and solutions
- Consolidates 4 to1 reduction in Object Storage Targets
- Seamless ClusterStor RAID management

xyratex
A Seagate Company

How does this Translate to Lustre?

- Traditional RAID 6 with 50MB/s rebuild rate of a 4TB Drive takes approximately 22 Hours to complete Rebuild
  - Translates: A OST performance will degrade to 45% of peak performance under heavy I/O load during rebuild for 22 Hours
    - Note 4 OSTs per OSS with RAID 6
- GridRaid using the same 50MB/s, has a reconstruction time of the same 4TB Drive of 5.5 Hours to complete recovery
  - Translates: A OST during reconstruction rate can have 0 impact or the same 45% impact performance, depending on what drives are used for the data stripe, for 5.5 Hours under heavy I/O load during recovery
    - Note 1 OSTs per OSS with GridRaid

Benefits of GridRAID for Lustre

- Improved MTTR
  - Repair >4X faster than RAID 6 @ 50MB/s per drive
- Distributed Recovery Workload
  - Less disruption to client IO
  - Overcomes the single drive bandwidth bottlenecks
- Rebalancing up to 4x Faster than Rebuild
  - Lower system impact
- Reduced Lustre OST Configuration
  - Reduces OST count by a factor of 4
- 4 Distributed Spare Volumes per SSU
  - RAID 6 provides 2 Global Hot Spares per SSU

In Summary

The Future Looks Bright for XYRATEX and LUSTRE

Thank You

Michael_Connolly@xyratex.com