Cray Storage Update LUG 2018

Nathan Schumann, Sr. Product Manager

Safe Harbor Statement

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts. These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

Cray and ClusterStor

Alignment **Commitment** **Integration**

Cray ClusterStor Platforms

ClusterStor L300

- Lustre 2.7
- 12 Gbit SAS enclosures
- Broadwell based ESMs
- EDR IB
- OPA
- 100 GbE
- 2x 40 GbE
- 6/8/10 TB HDDs

Large sequential I/O workloads

ClusterStor L300N

- NXD I/O Manager
- 3.2 TB SSDs
- Advanced MMU

Many applications with mixed I/O patterns

ClusterStor L300N: Any Workload, Any Time

Goal: Support Widest Application I/O Variety Within Budget

Compute Cluster Applications

Sensor Data Capture & Processing
Data mining - Simulations
Modelling - Software development
Visualization of complex data -
Rapid mathematical calculations

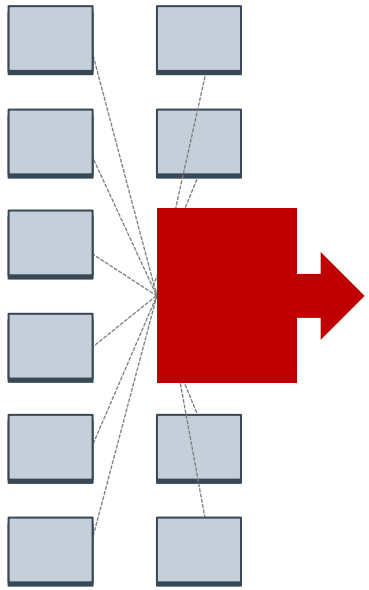

Challenge: Lost Productivity Due To Poor I/O Predictability

Mixed I/O Patterns=

Unpredictable Application Performance

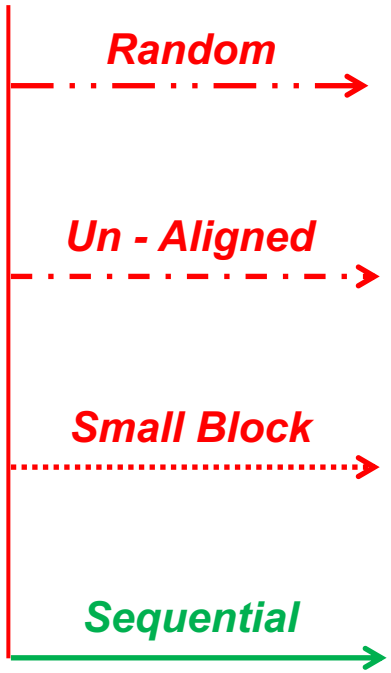

Solution: Smart I/O Management Provides Consistent Performance

ClusterStor NXD =

Transparent Redirection of I/O to Appropriate SSD or HDD Medium

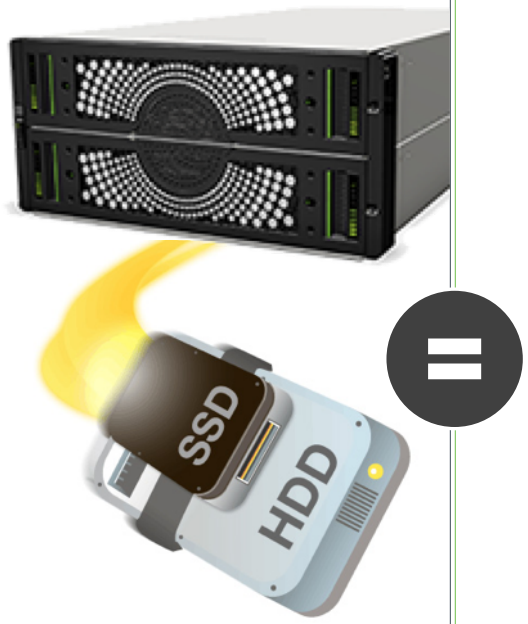

Result: I/O Optimized For Each Application I/O Profile

ClusterStor NXD =

Predictable and Fastest Compute Application Time-to-Solution

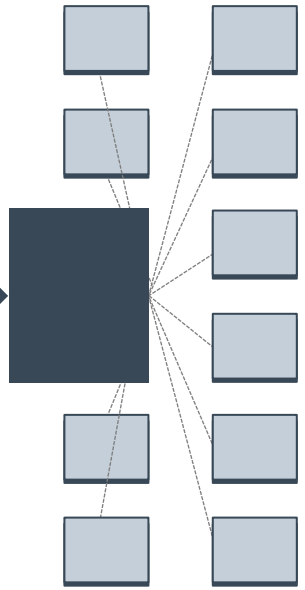

COMPUTE | STORE | ANALYZE

Three Horizons of ClusterStor

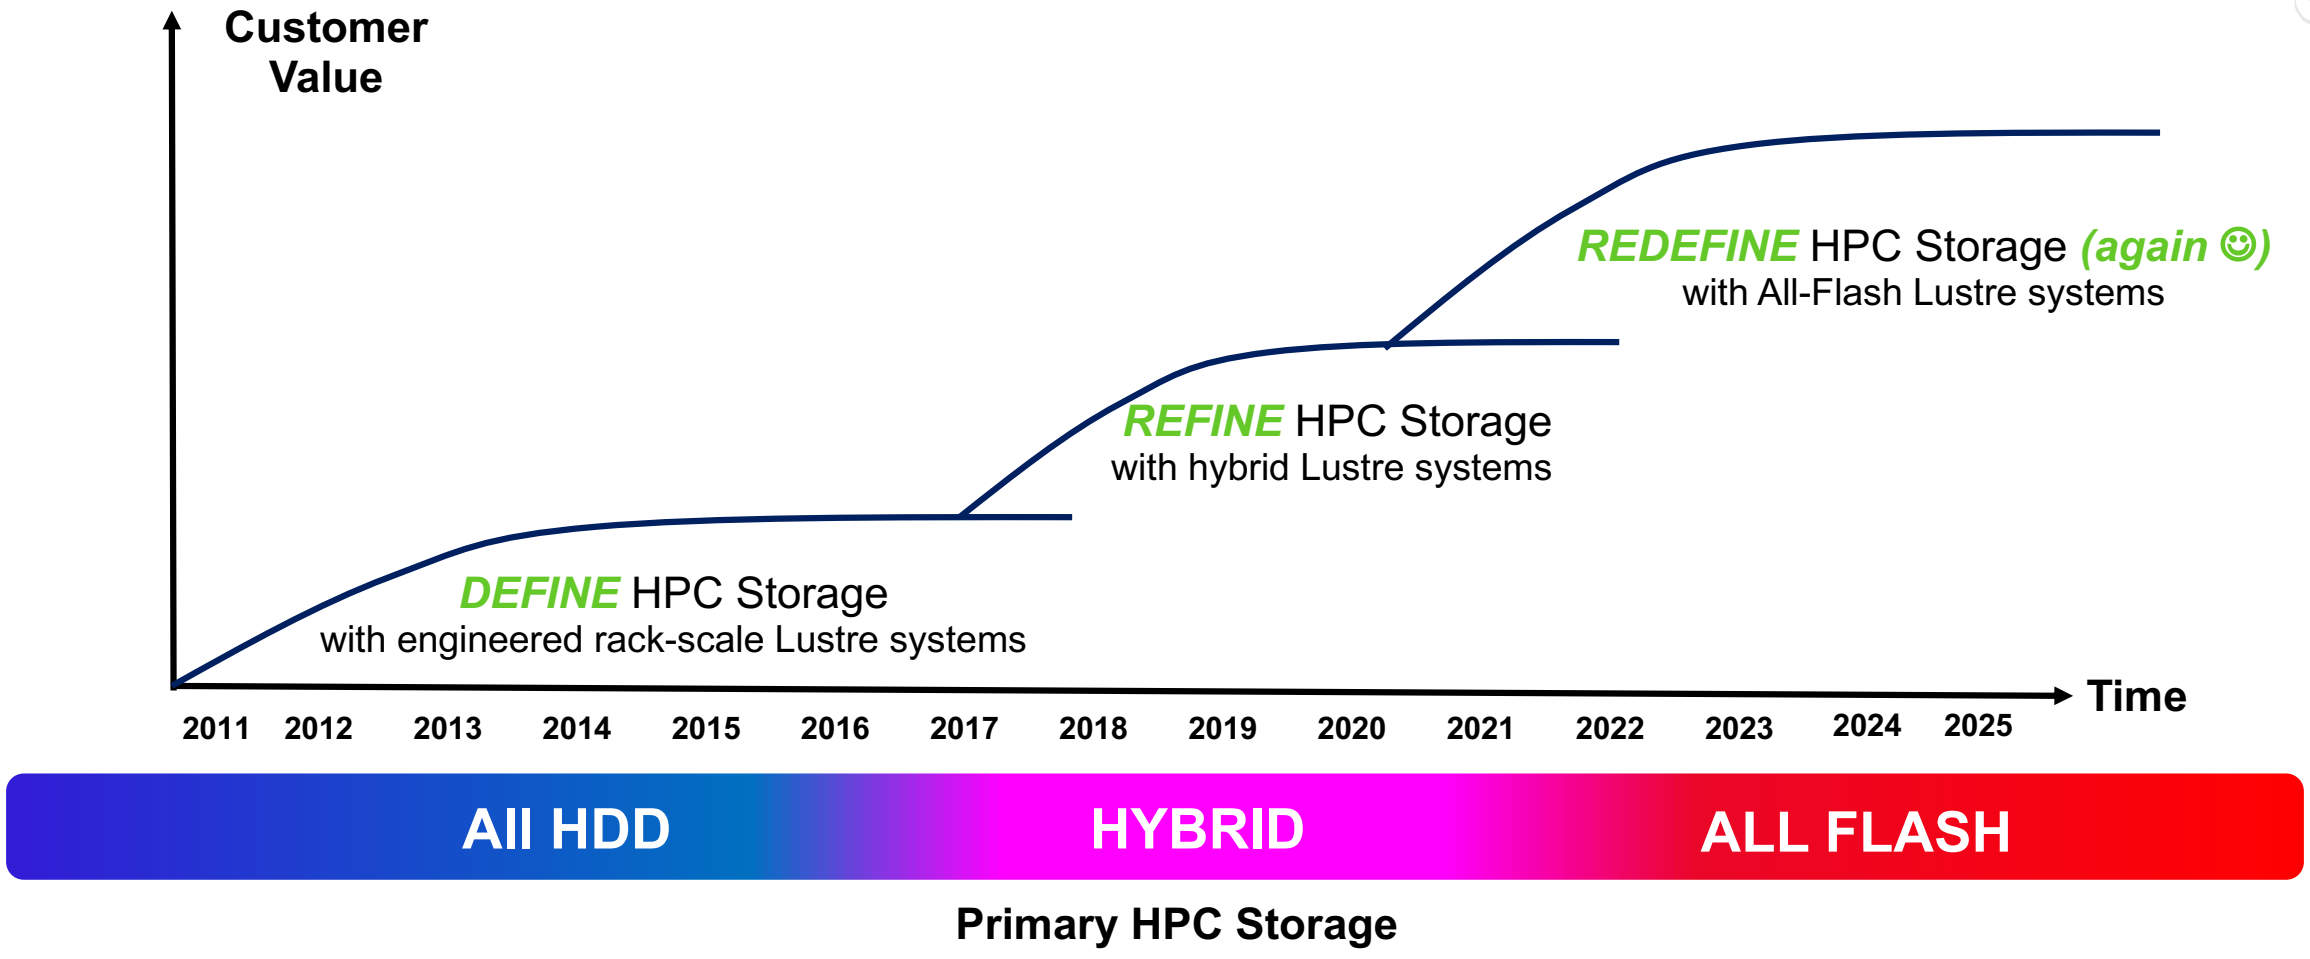

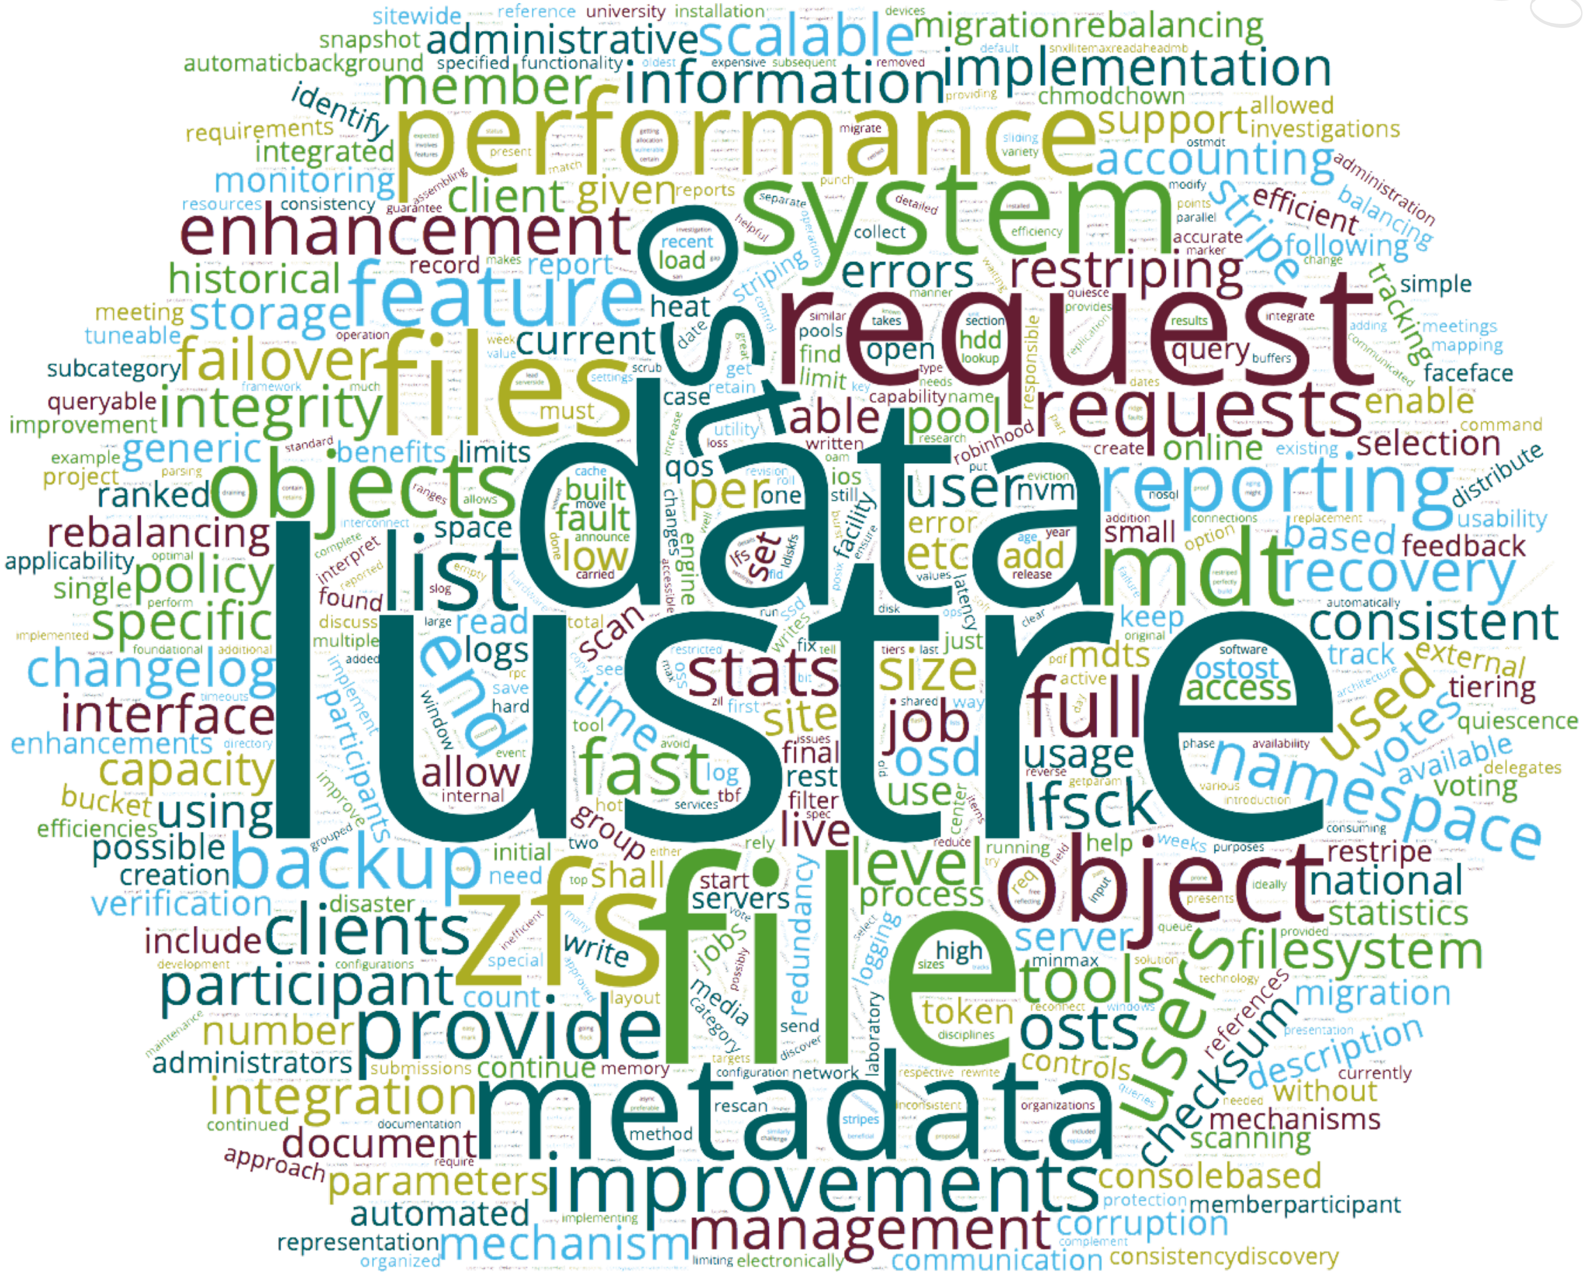

Lustre Improvements for Flash

Reduce client code overhead	Reduce network RPCs	Reduce MDS overhead
Small IO improvements • Tiny writes [2.11]* • Fast reads [2.11]*	• Lock Ahead [2.11]* • Immediate short IO [2.11]* • Small overwrites [2.11]* • Data on MDT [2.11] • Size on MDT*	• DNE.2 [2.8] (reduce MDS bottlenecks) • Progressive file layouts [2.10] (more compact layouts)

Other improvements:

- OST/MDT server tunings for SSD
- Multi-channel PTLRPC (QoS)

ClusterStor Lustre Release Adoption

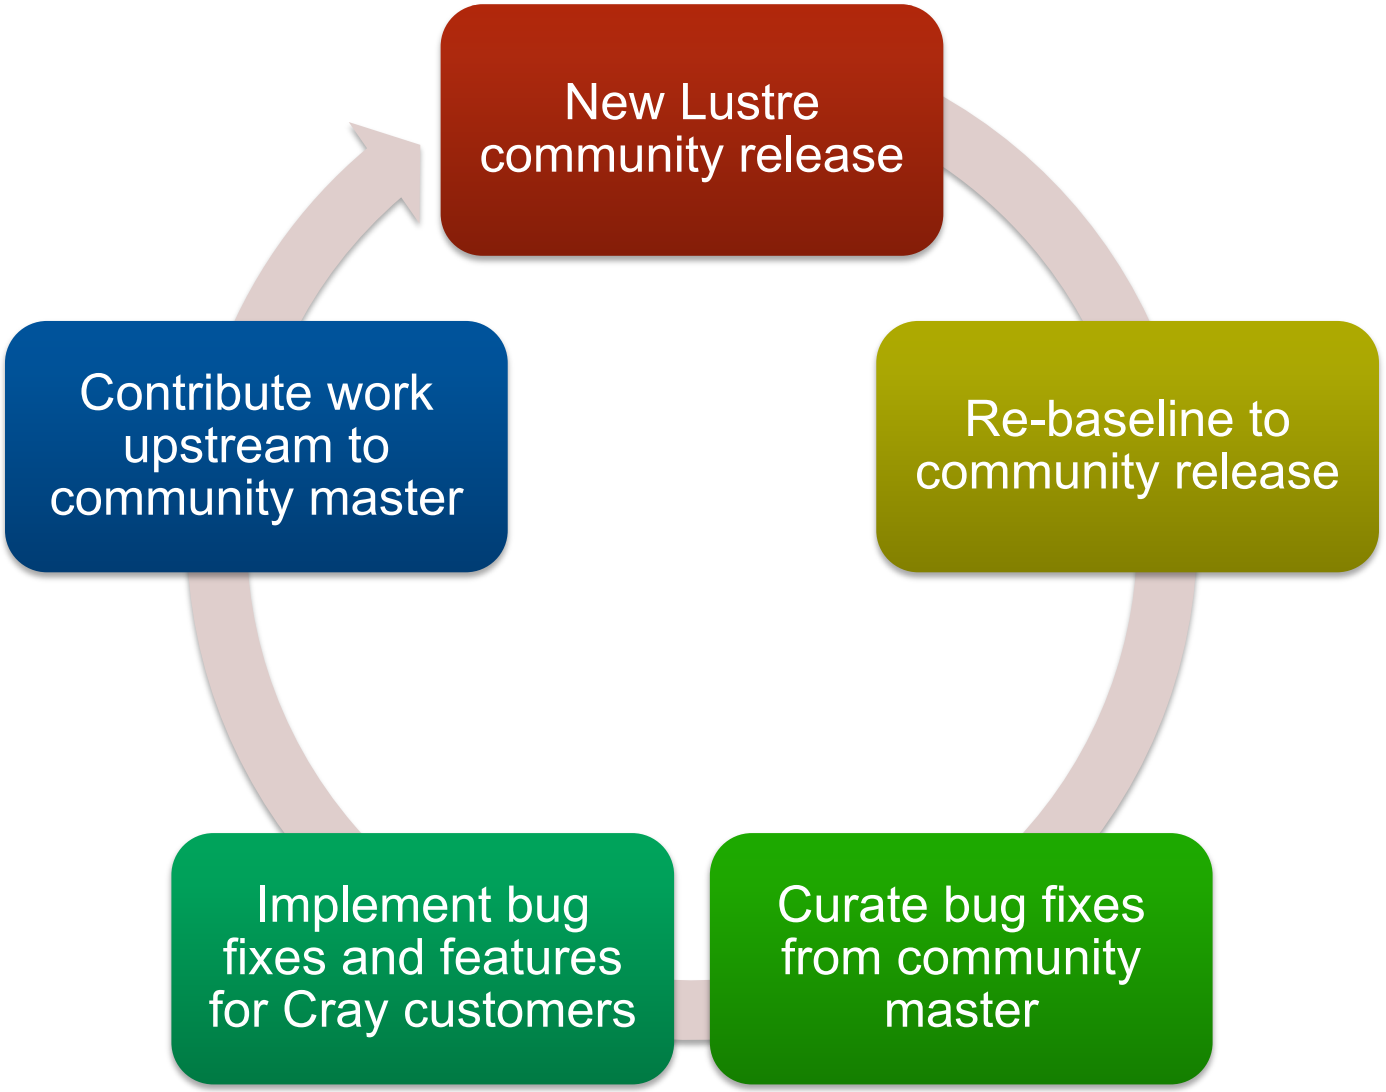

Lustre Features Coming to ClusterStor

Feature Name	Lustre Release	ClusterStor Release
Lock ahead	2.11	2.x
Multiple MD modify RPCs per client	IEEL 3.0, 2.8	3.0
DNE phase IIb (async commit and recovery)	2.8	3.next
UID/GID mapping	2.9	3.next
Subdirectory mounts	2.9	3.next
Large bulk I/O	2.9	3.next
Multirail LNet	2.10	3.next
Progressive file layouts	2.10	3.next
Project quotas	2.10	3.next
Data on MDT (DoM)	2.11	3.next
File-level redundancy (FLR) - delayed resync	2.11	3.next

VIEW for ClusterStor™

Job Performance Details and Scoring

Jobs
Searchable fields for the most important data to the job step level

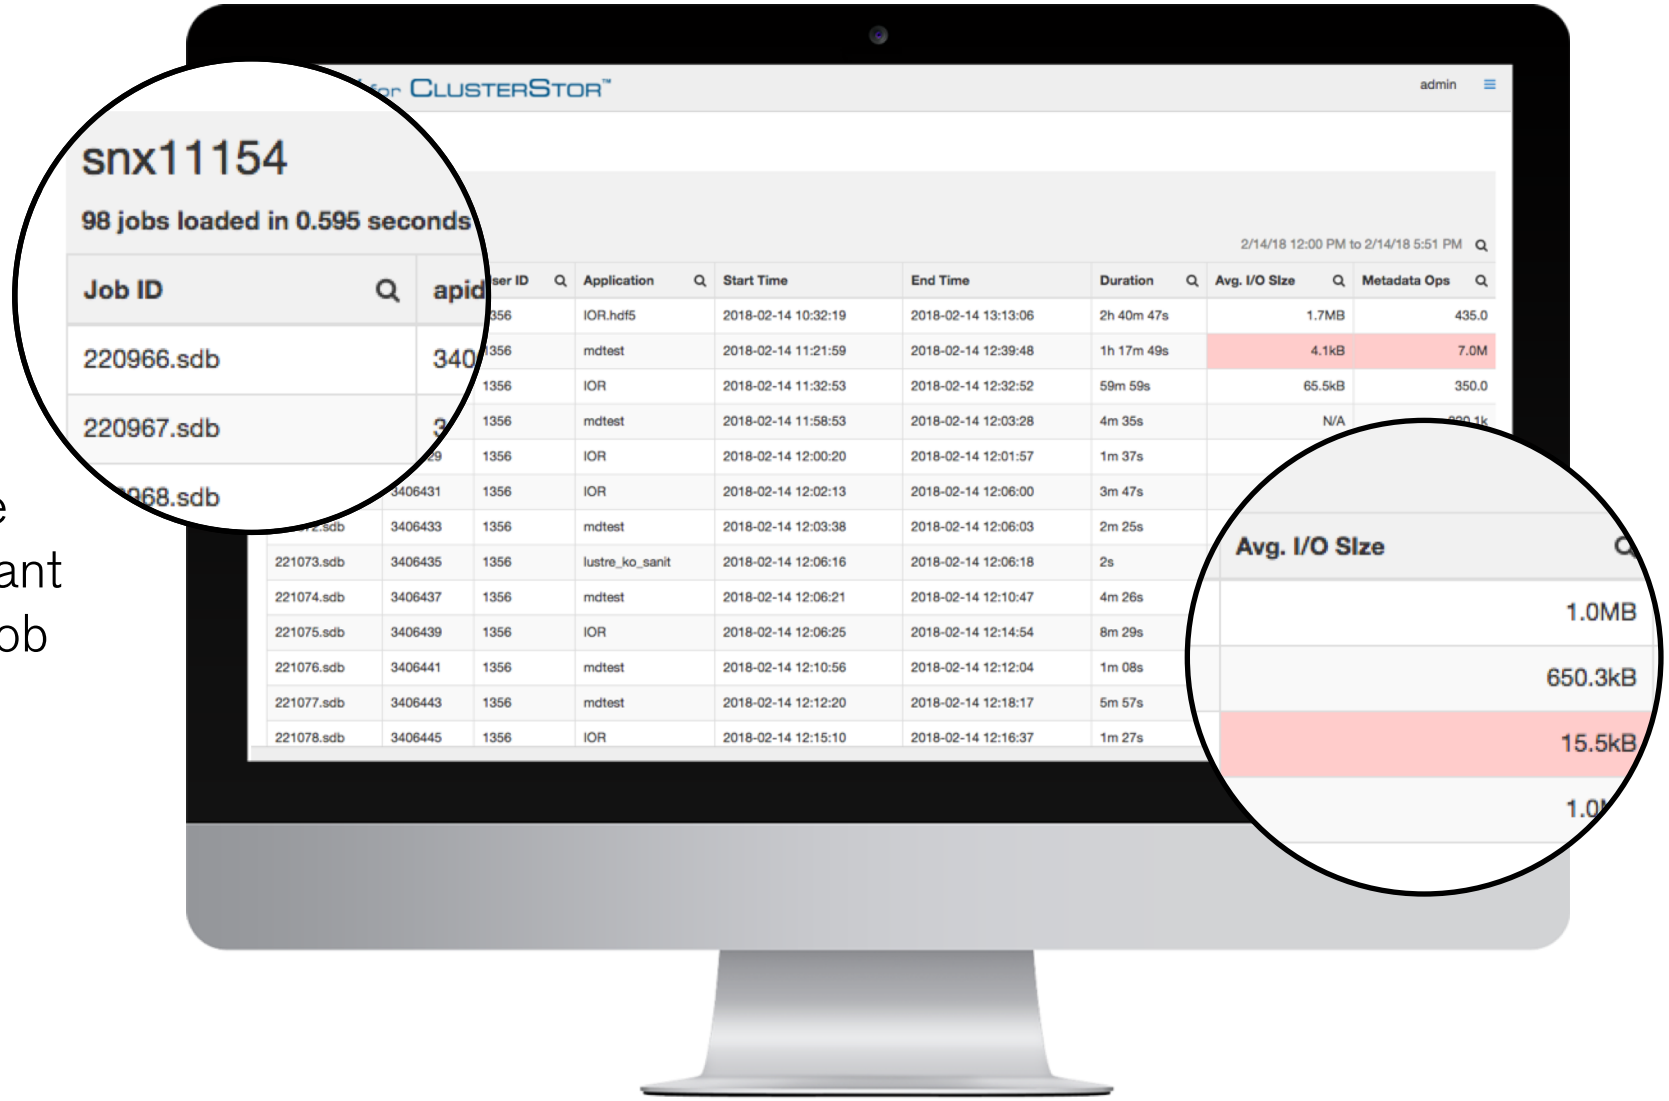

Scoring
Know which jobs might be causing issues with visual cues

Performance Visualization and Comparison

Visualize
Performance graphs over the life of the job for write, read, and metadata operations

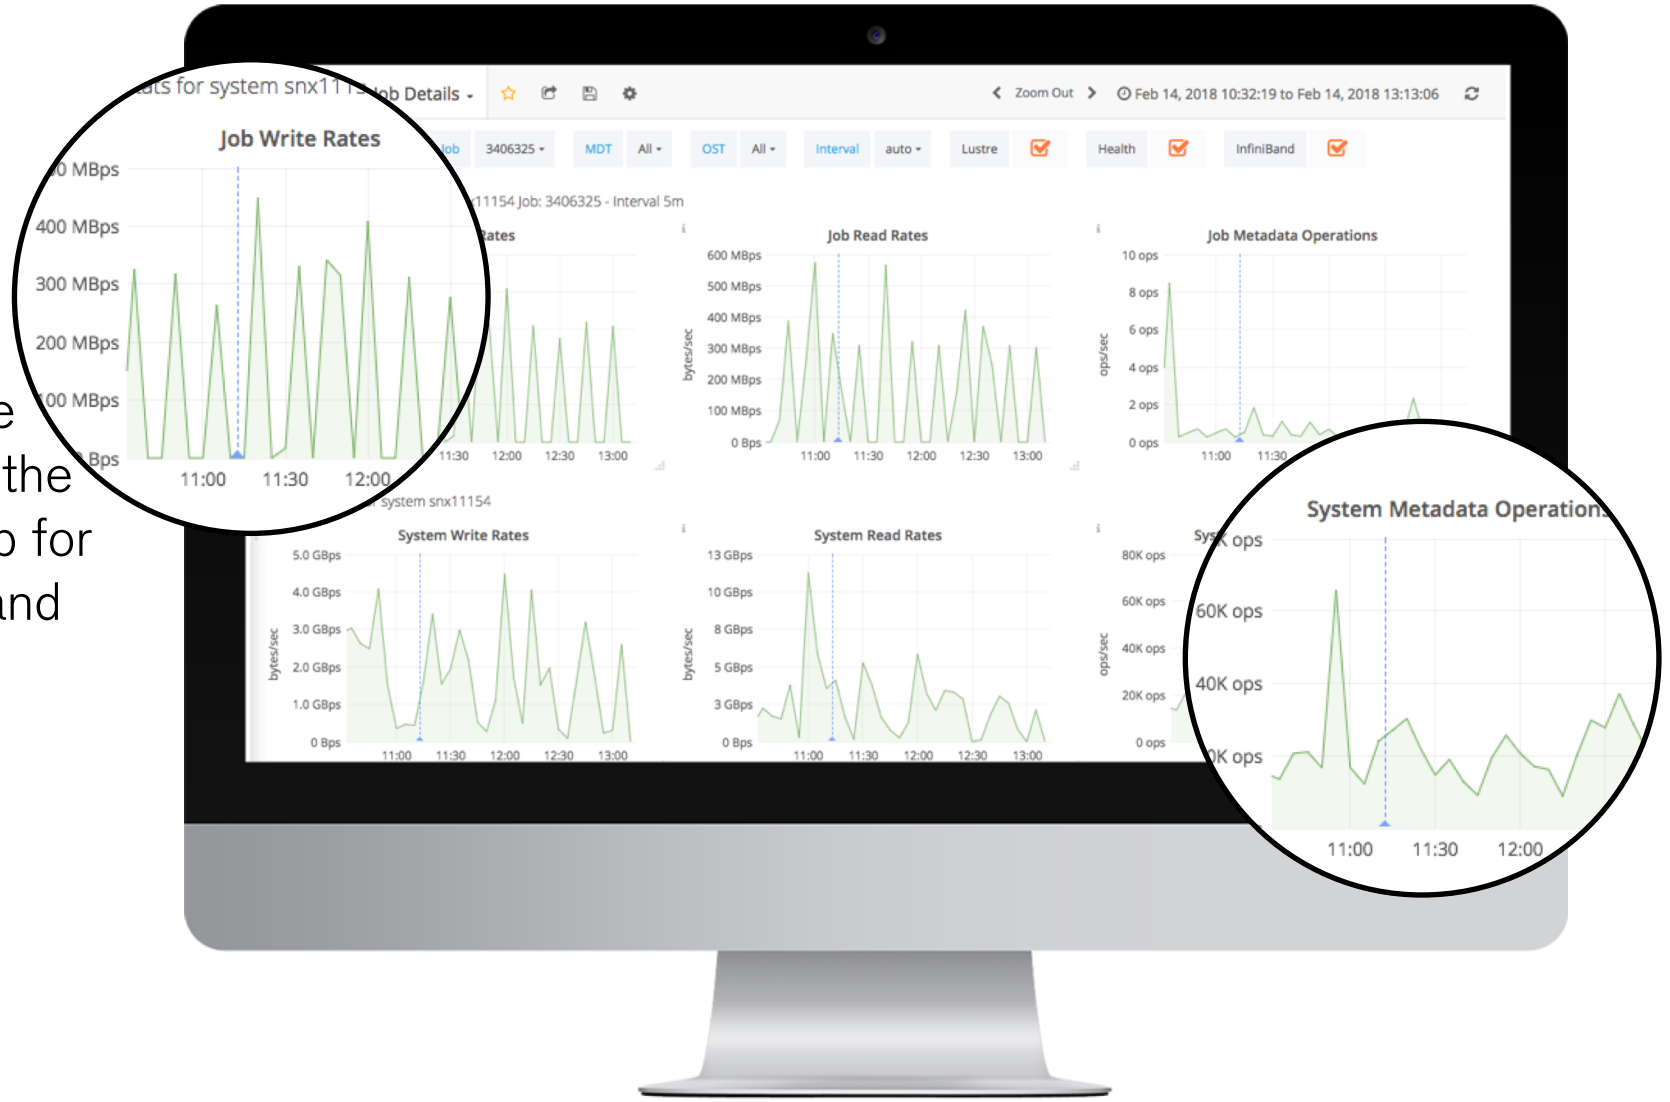

Compare
Compare this job to the rest of the system at a glance

Questions?