

CRAY

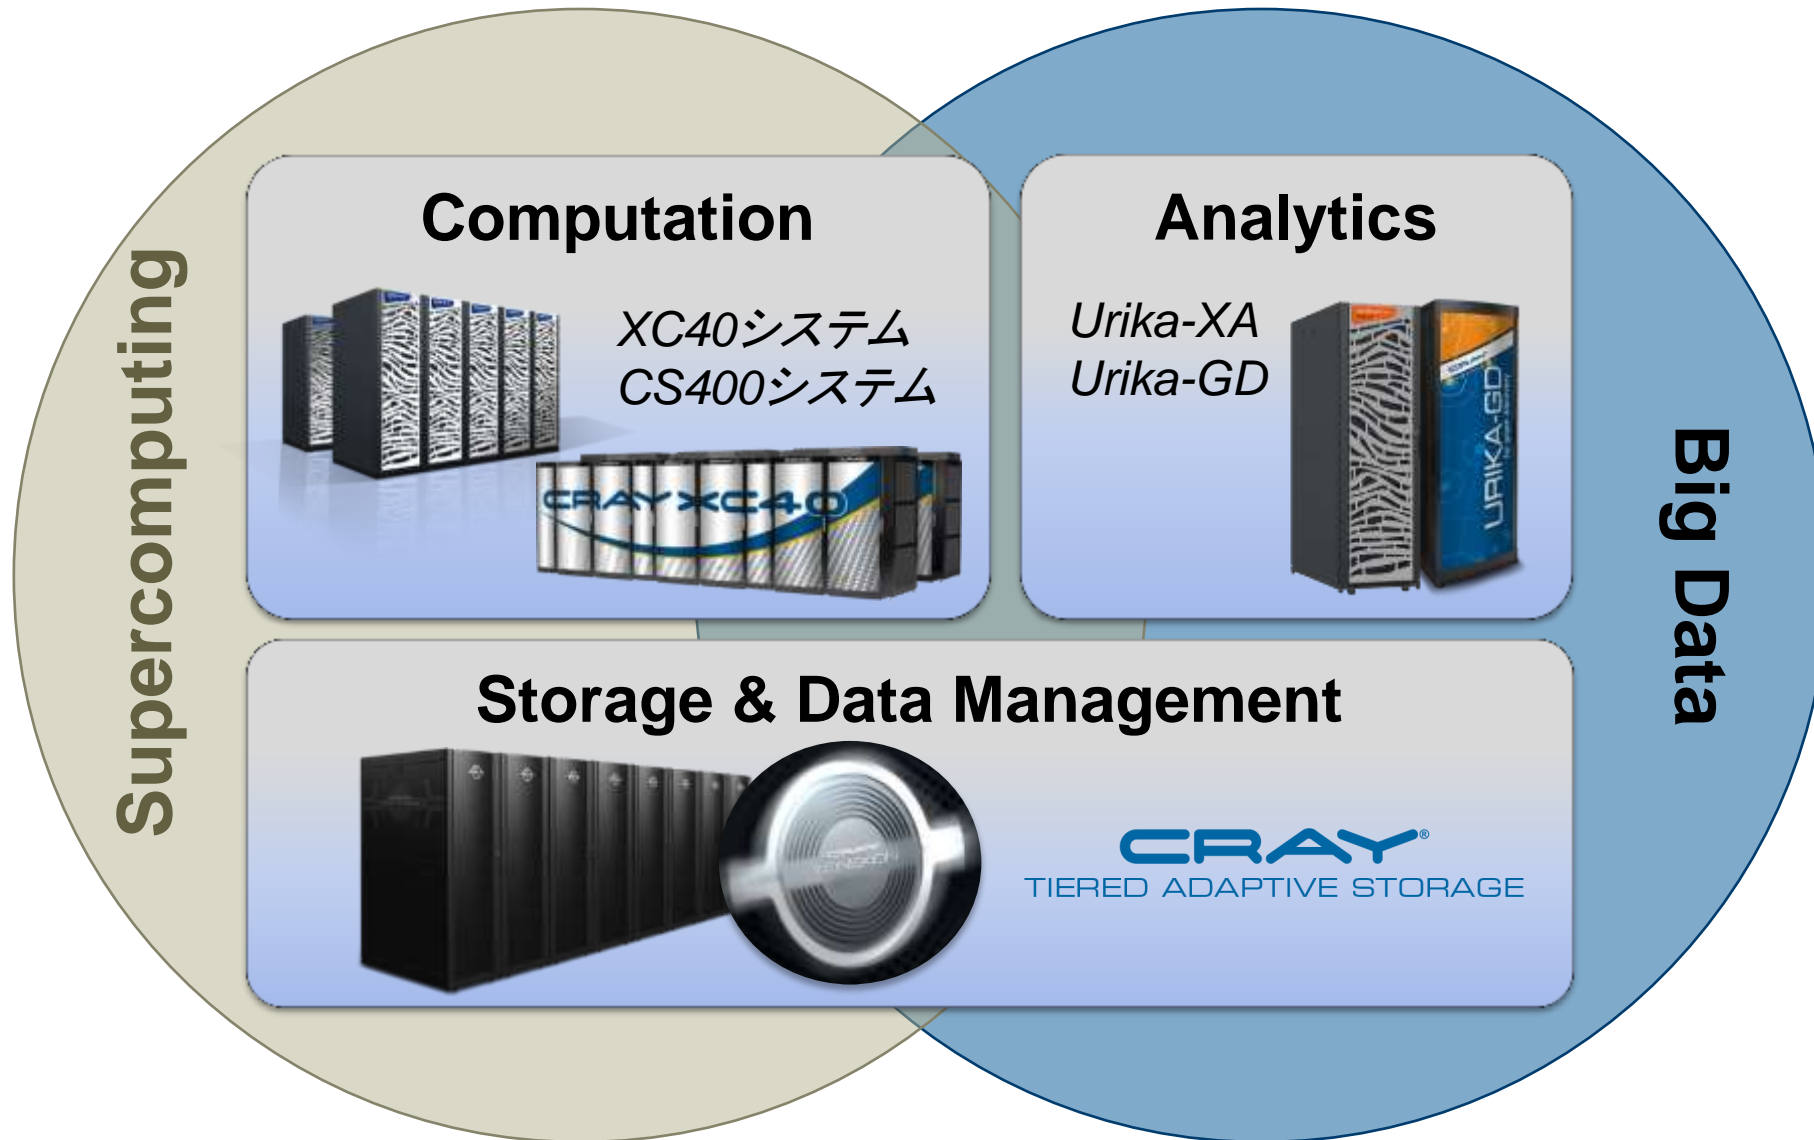
第15回PCクラスタシンポジウム
ストレージ製品とデータ・ソリューションのご紹介

クレイ・ジャパン・インク

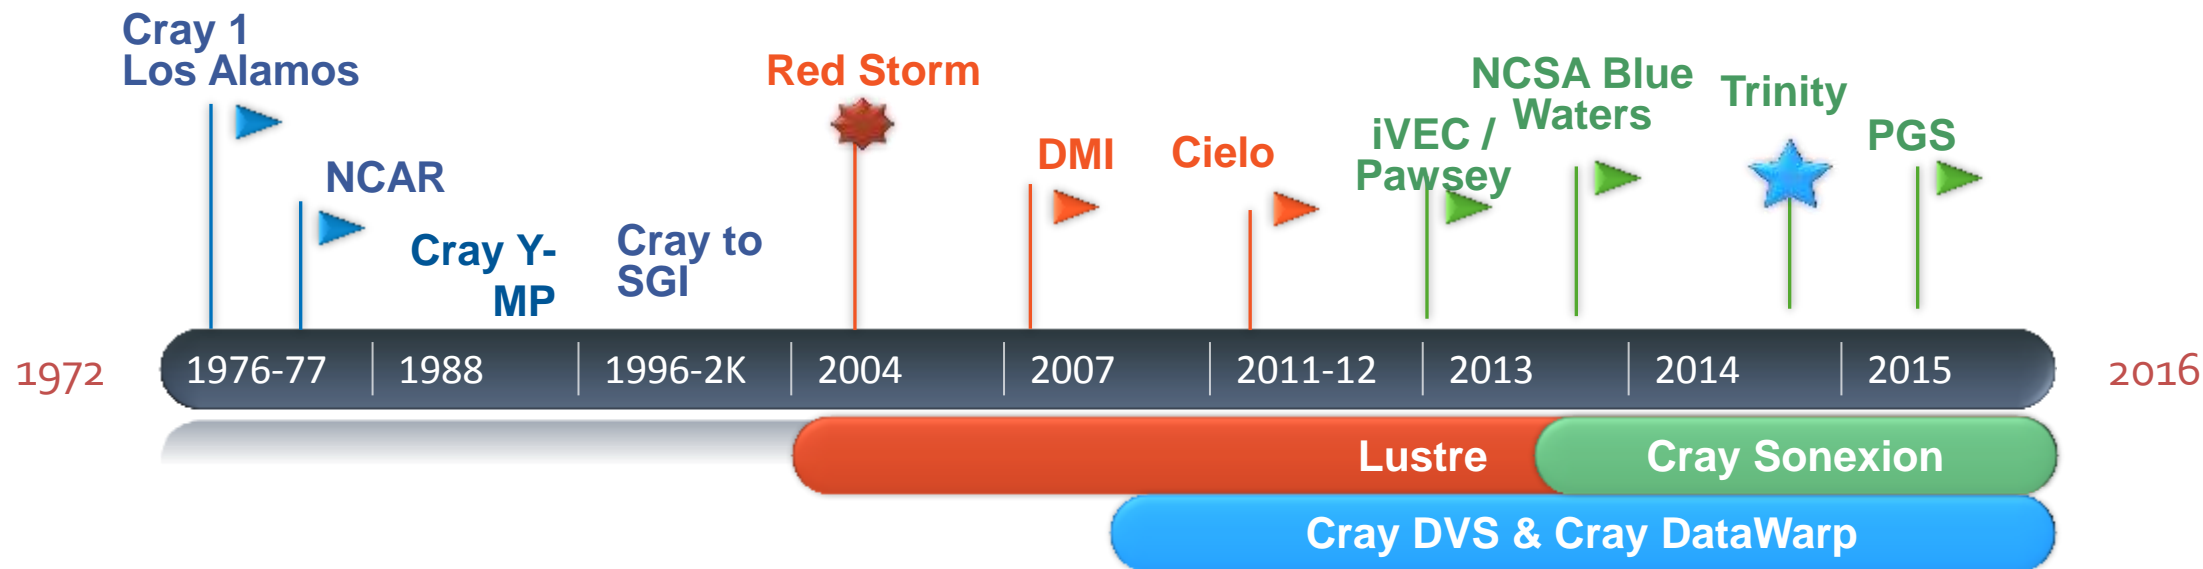
Safe Harbor Statement

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.

クレイ・アダプティブ・スーパーコンピューティング



Innovation



*The Cray DataWarp technology in the Trinity system will provide **the first multi-petabyte multi-terabyte/sec IO burst handling capability, ever.***

--Gary Grider, HPC Division Leader, LANL

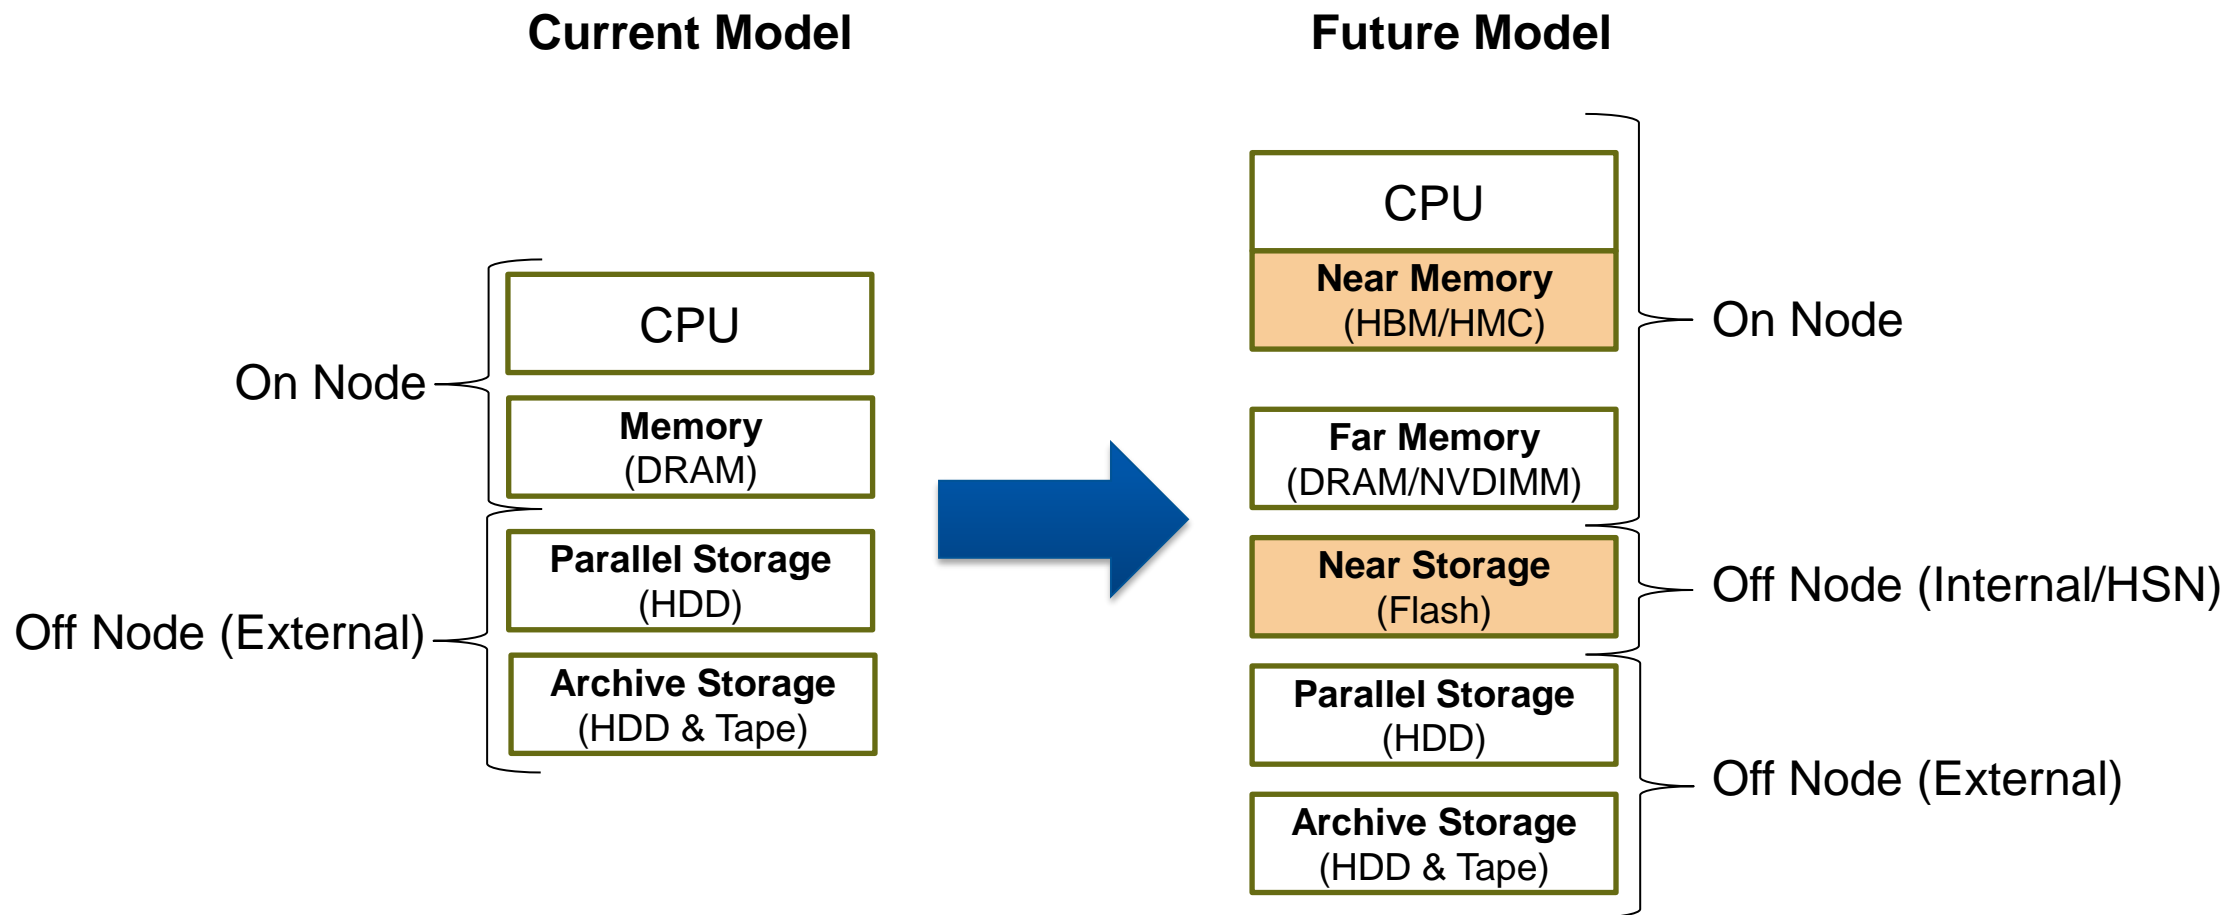




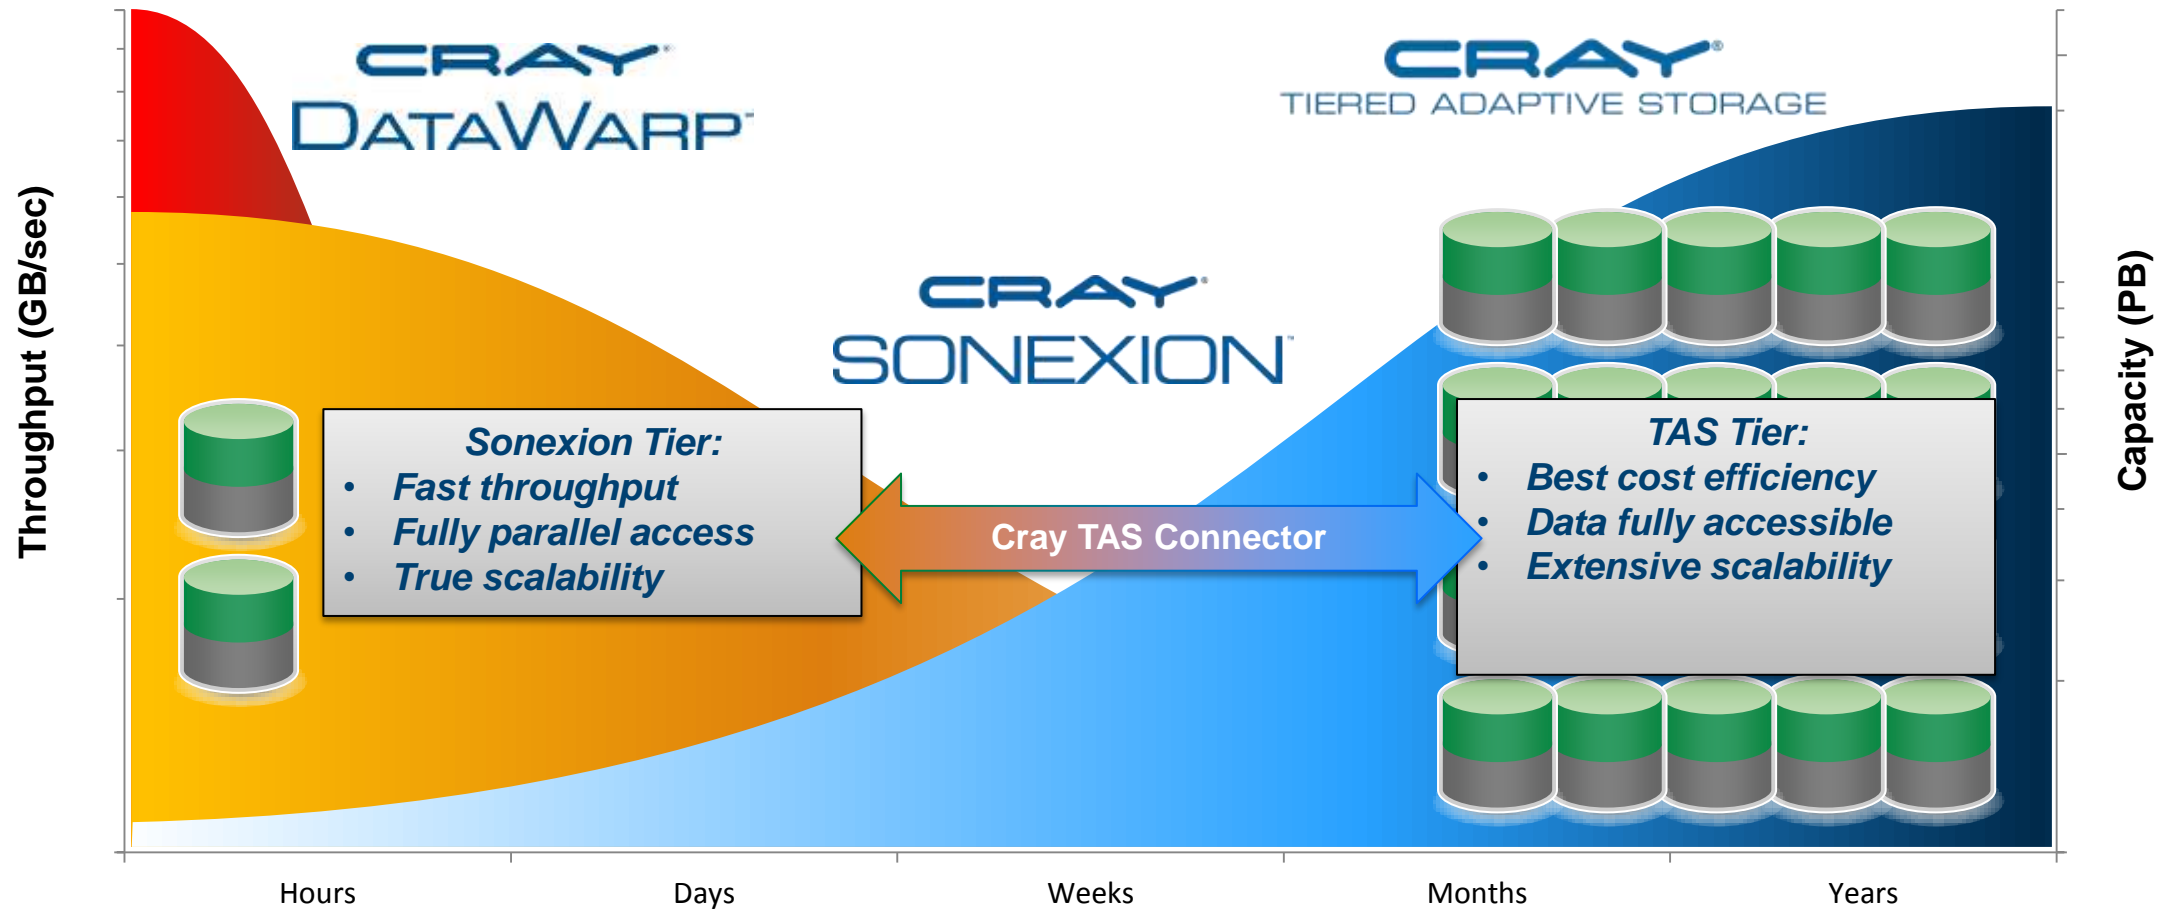
Storage Trends

Convergence of Supercomputing and Big Data


Computing Memory and Storage Trends



Cray Storage Solutions - Span Data Lifecycle



Cray Storage Offerings



CRAY
DATAWARP™



I/O Acceleration

CRAY
SONEXION™



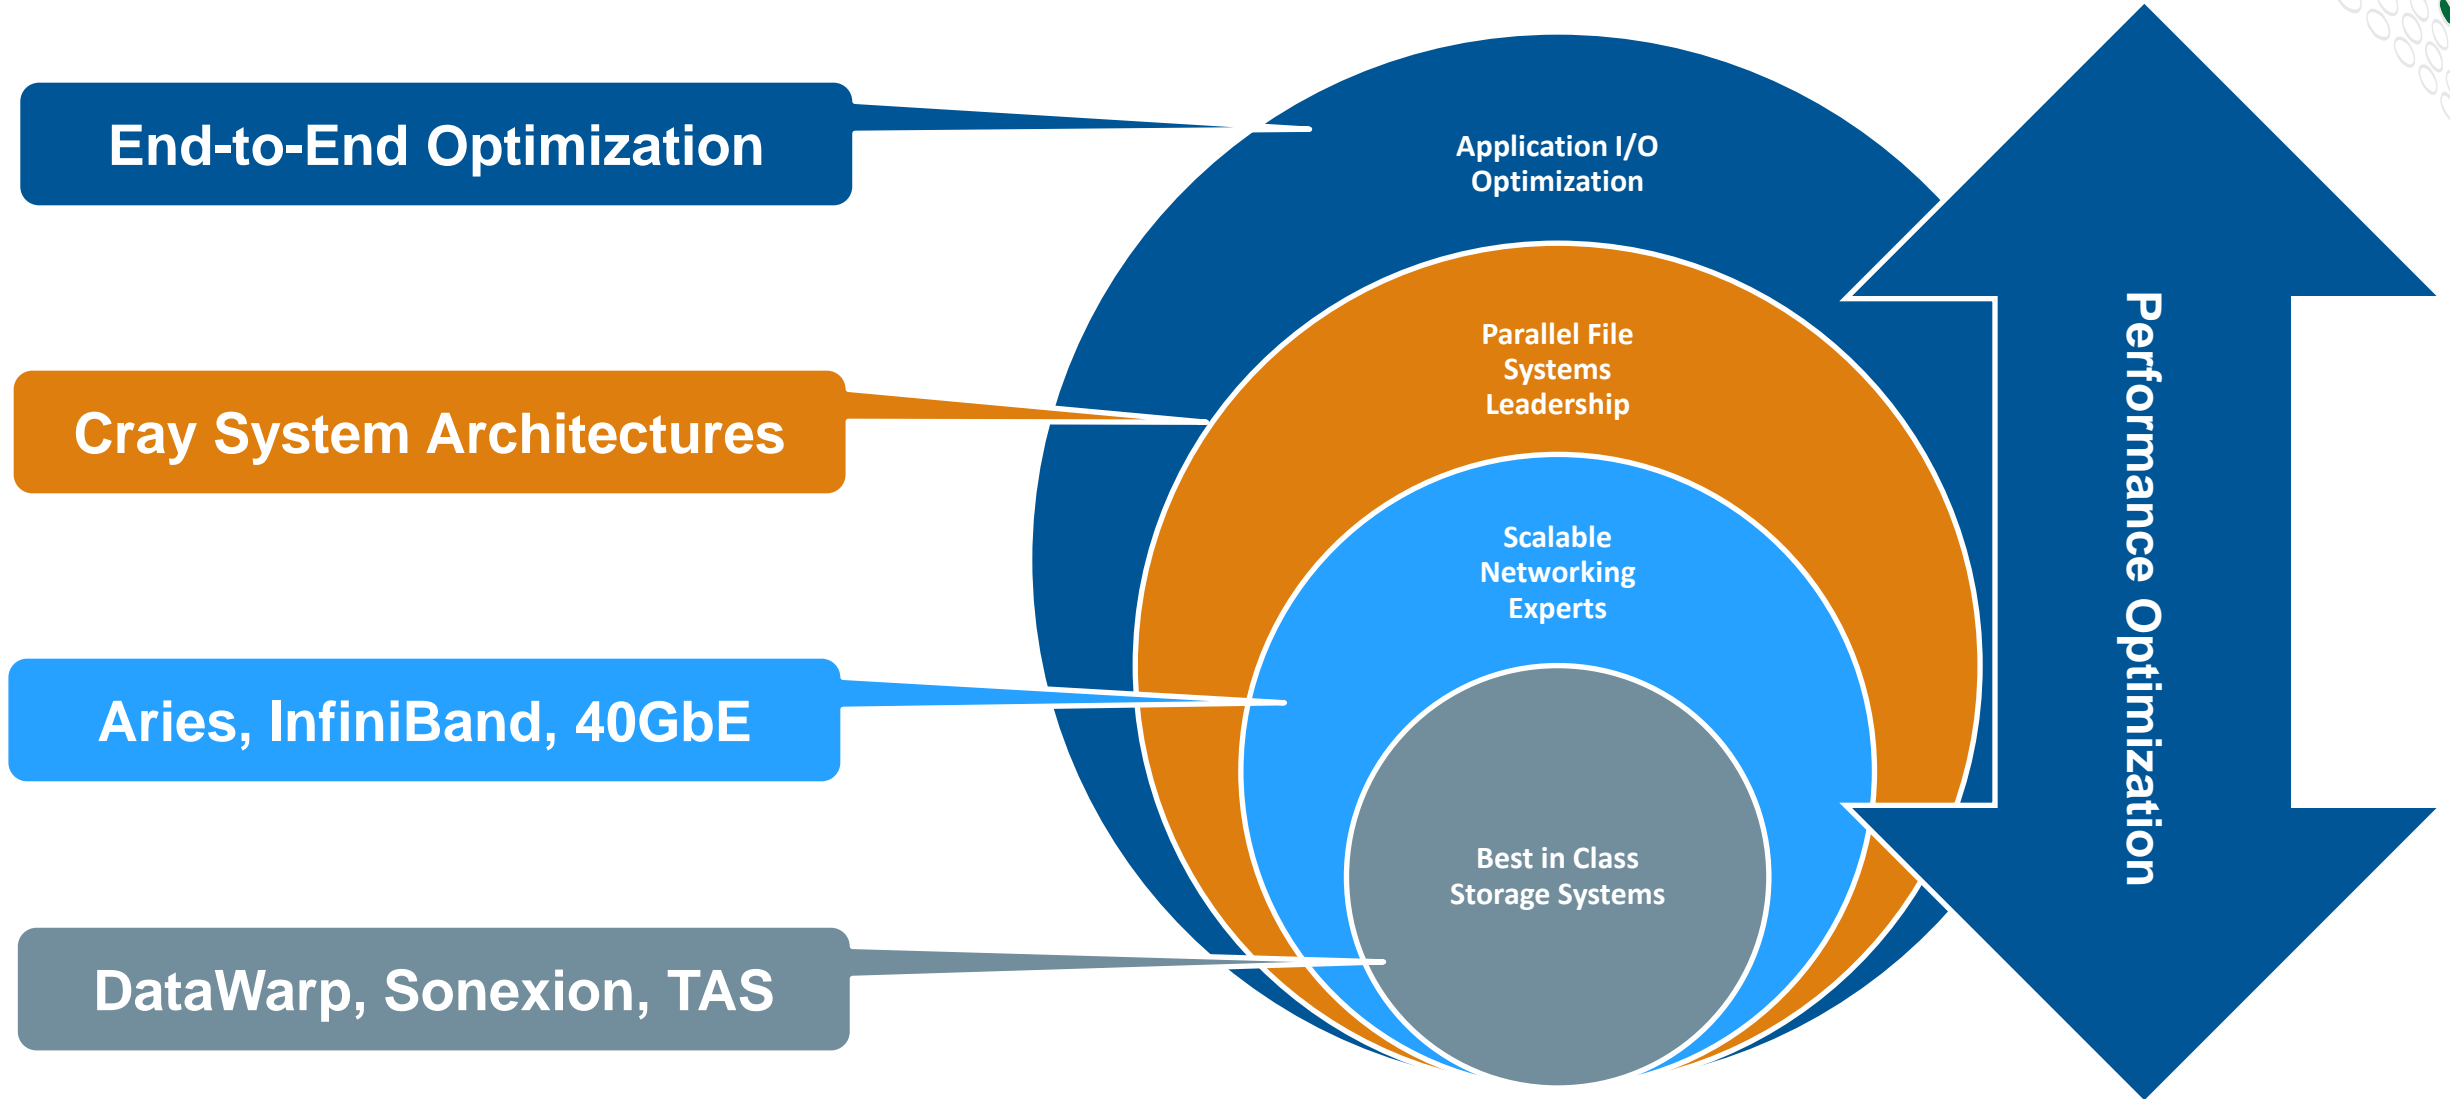

High Performance Storage

CRAY
TIERED ADAPTIVE STORAGE



Efficient Long-term Storage

Summary – Optimized End-to-End Storage Solutions



High Performance Flash Storage & I/O Acceleration

CRAY®
DATAWARP™

Flash Storage IO Acceleration System for Cray XC40



Performance

- Scales from 70 thousand to 40 million IOPS
- Accommodate wide range of workloads



Efficiencies

- 5x the performance of disk - same cost
- Offloads I/O intensive workloads




CRAY®
DATAWARP™

Cohesive

- Flexible usage models
- Automated workflows

Bridges the performance gap between compute and storage

DataWarp Benefits



- **DataWarp absorbs spikes**
 - Reduces size of underlying PFS
 - Provision PFS for sustained performance instead of peak
 - Accommodates range of applications
- **Improves efficiency**
 - Flash tier: 3-5x the performance of disk at the same cost
 - Machine efficiency improved
 - Better use of compute resources

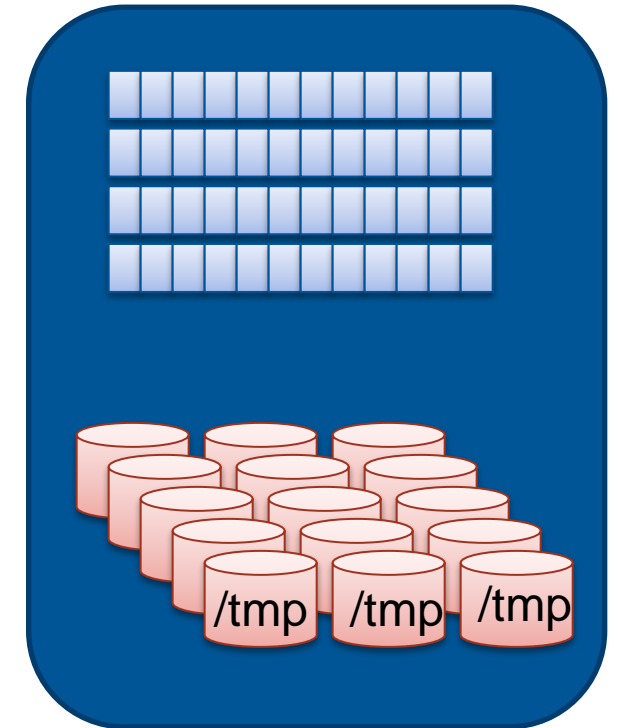
Use Case: Local Storage on Demand

Per Node Scratch

- Each compute node in a job is assigned a private part of the allocated SSD space
- Much faster than “faking it” with a parallel file system

Per Node Swap Space

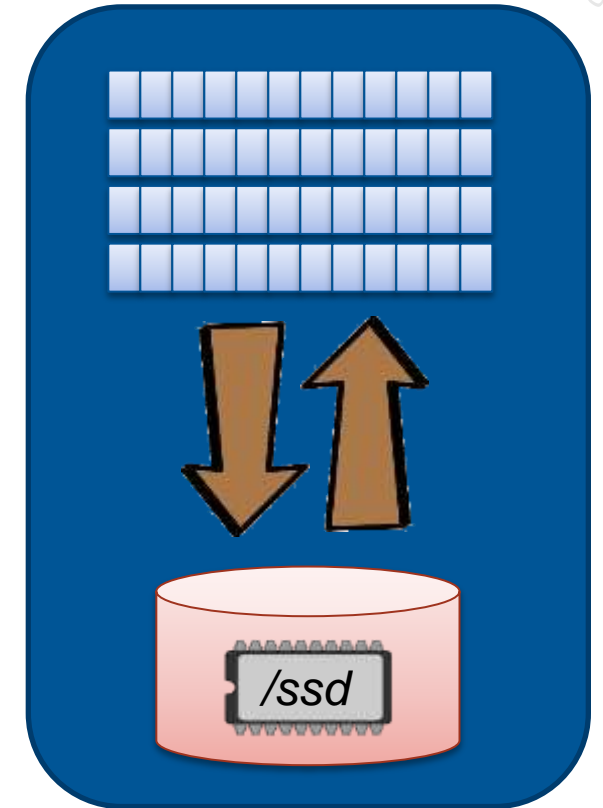
- Dynamic compute node swap space



CRAY®
DATAWARP™

Shared Fast Scratch

- High Bandwidth access to shared files
- Files can be striped across multiple DataWarp Nodes
- Space can be temporary for the job, or be marked as persistent to work between jobs




CRAY
DATAWARP™

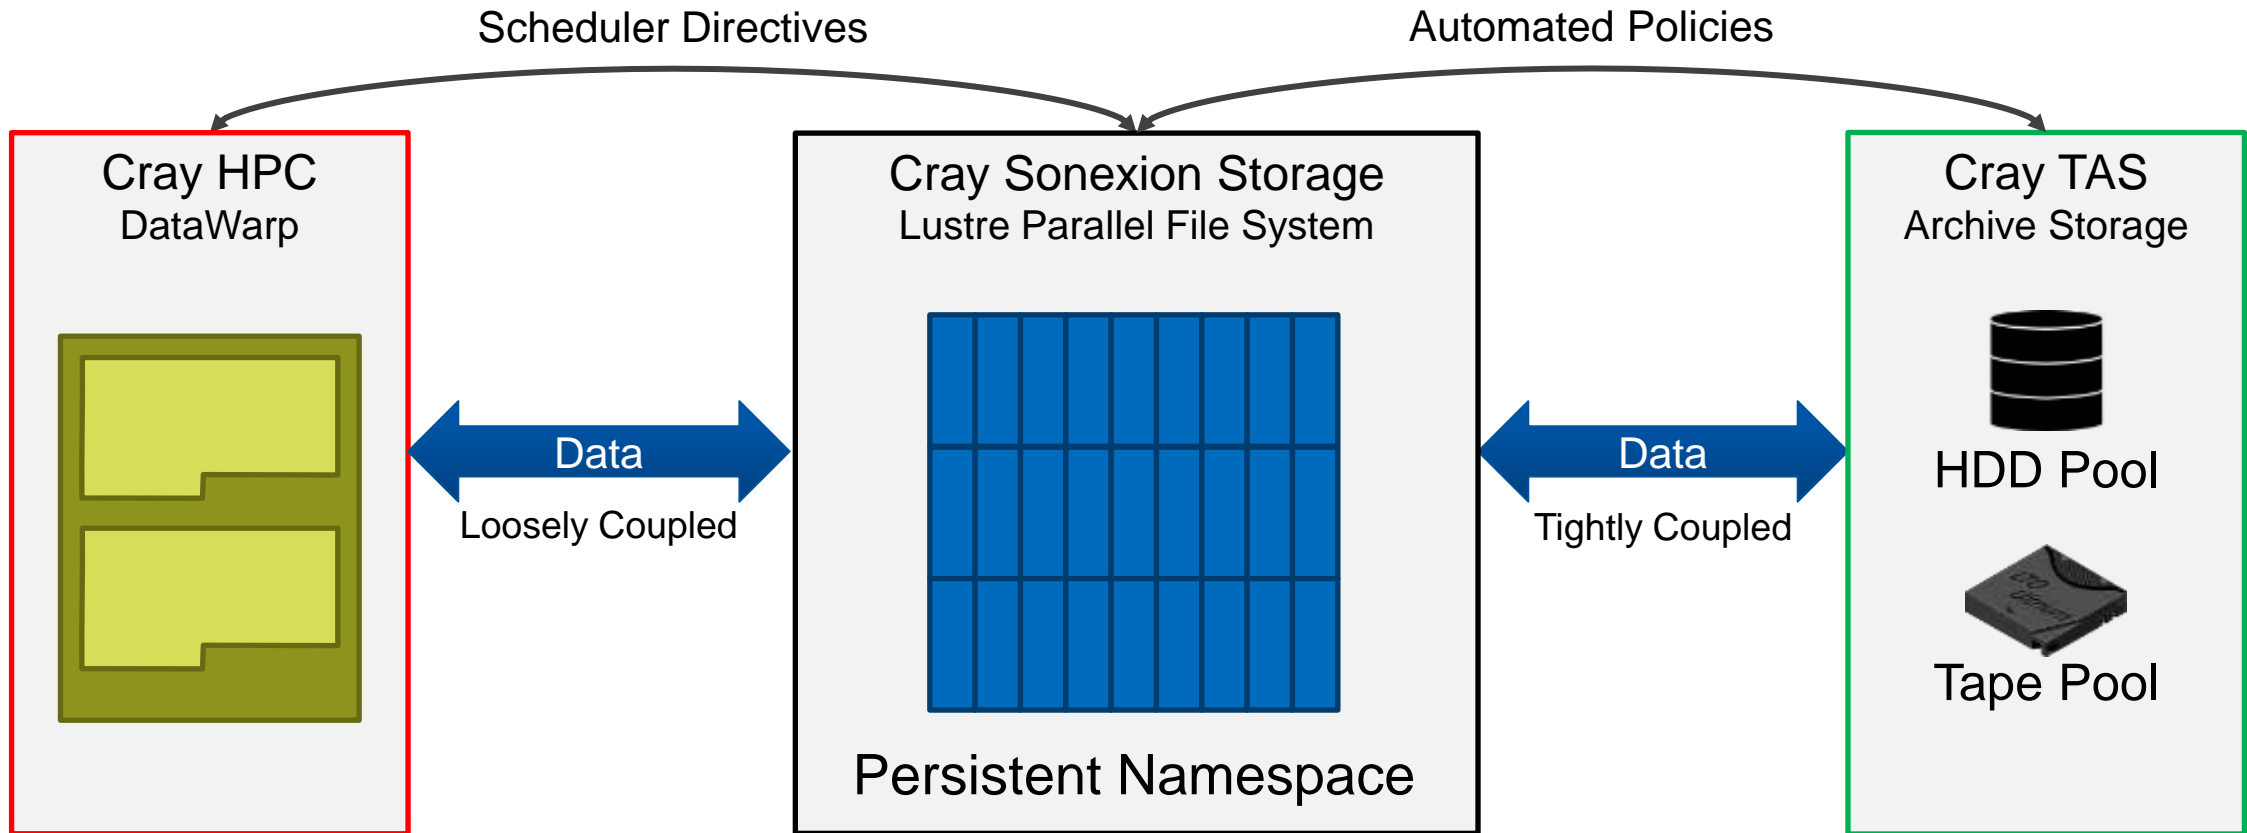
Use Case: Checkpoint / Restart

Fast Checkpoint / Restart

- User asks for enough SSD to cover the number of concurrently resident checkpoints
- High Bandwidth checkpoints are written to SSDs
- Followed by an asynchronous explicit or transparent copy out to rotating storage



Automated Coordination Across Storage Tiers



High Performance Primary Storage

CRAY®
SONEXION®

Efficient, proven, scale-out Lustre system for any Linux HPC environment



**Innovate
Faster**

- Deploy faster
- Achieve results faster

**Scale
Efficiently**

- Sustained performance over 1 TB/s
- 25% fewer components
- Over 3 PBs usable in one 42U rack

Be Confident

- Proven Cray architectures
- Single point of support

Optimize performance and capacity at maximum density

Efficiency Comparison – 1TB/s Example



Deployment Comparison	Modular Storage with External Servers	Cray Sonexion (NCSA Blue Waters)	Monolithic Storage with External Servers
Capacity	22 PBytes	22 PBytes	32 PBytes
Bandwidth	1 TByte/sec.	1 TByte/sec.	1 TByte/sec.
LNET routers	942	482	440
Storage units	472	180	360
Hard drives	28,320	15,120	20,160
External servers	942	0	294
Director IB switches	6	0	2
IB cables	5,468	482	1512
Racks	94	30	40
Cost	\$\$\$	\$	\$

Simplify Management



CRAY SONEXION SYSTEM MANAGER

Dashboard Home Center Performance Log Review Support Health Configure

Master Status

File System Throughput

Inventory

Type	Severity	Status
Racks	2	Installed
Services	1	Installed
Storage Hardware	4	OK
Storage Media	103	Problems Found
Power & Cooling	30	OK

Top System Statistics

Item	Value
File System Peak Read	340.28 MB/s
Current Read	0 B/s
Peak Write	334.01 MB/s
Current Write	0 B/s
Metadata Current Operations	0 Ops/s
Storage Number of OSTs in use	24
Number of Disks in Use	1,054
Capacity in use	38.48 TB
Capacity available	1.59 PB
Power Current cluster usage	11.73 KW

CRAY SONEXION SYSTEM MANAGER

Dashboard Home Center Performance Log Review Support Health Configure

46 UP 0 DOWN 0 UNREACHABLE 0 SPENDING 0 (4) TOTAL

Current Network Status

Service Status Details For Service Group 'Inventory'

Host	Service	Status	Last Check	Duration	Attempt	Status Information
101104001	Archie and Disk Status	OK	2015-08-04 03:30:41	47s 16.46s (7s)	\$0	There is an issue with one or more services. T1091 Not checked
101104014	Archie and Disk Status	OK	2015-08-04 03:27:07	47s 28.22s (1s)	\$0	There is an issue with one or more services. T1091 Not checked
1011154015	Archie and Disk Status	OK	2015-08-04 03:35:01	47s 28.36s (1s)	\$0	There is an issue with one or more services. T1091 Not checked

Calendar Mail People Tasks

Tiered Data Management

CRAY®

TIERED ADAPTIVE STORAGE

Cray Tiered Adaptive Storage (TAS)

CRAY®

Protect and access your high-value data when you need it, for as long as you need it



Protect

- Preserve assets efficiently
- Transparently migrate data

Access

- Continuously accessible
- Flexible access models

Sustain

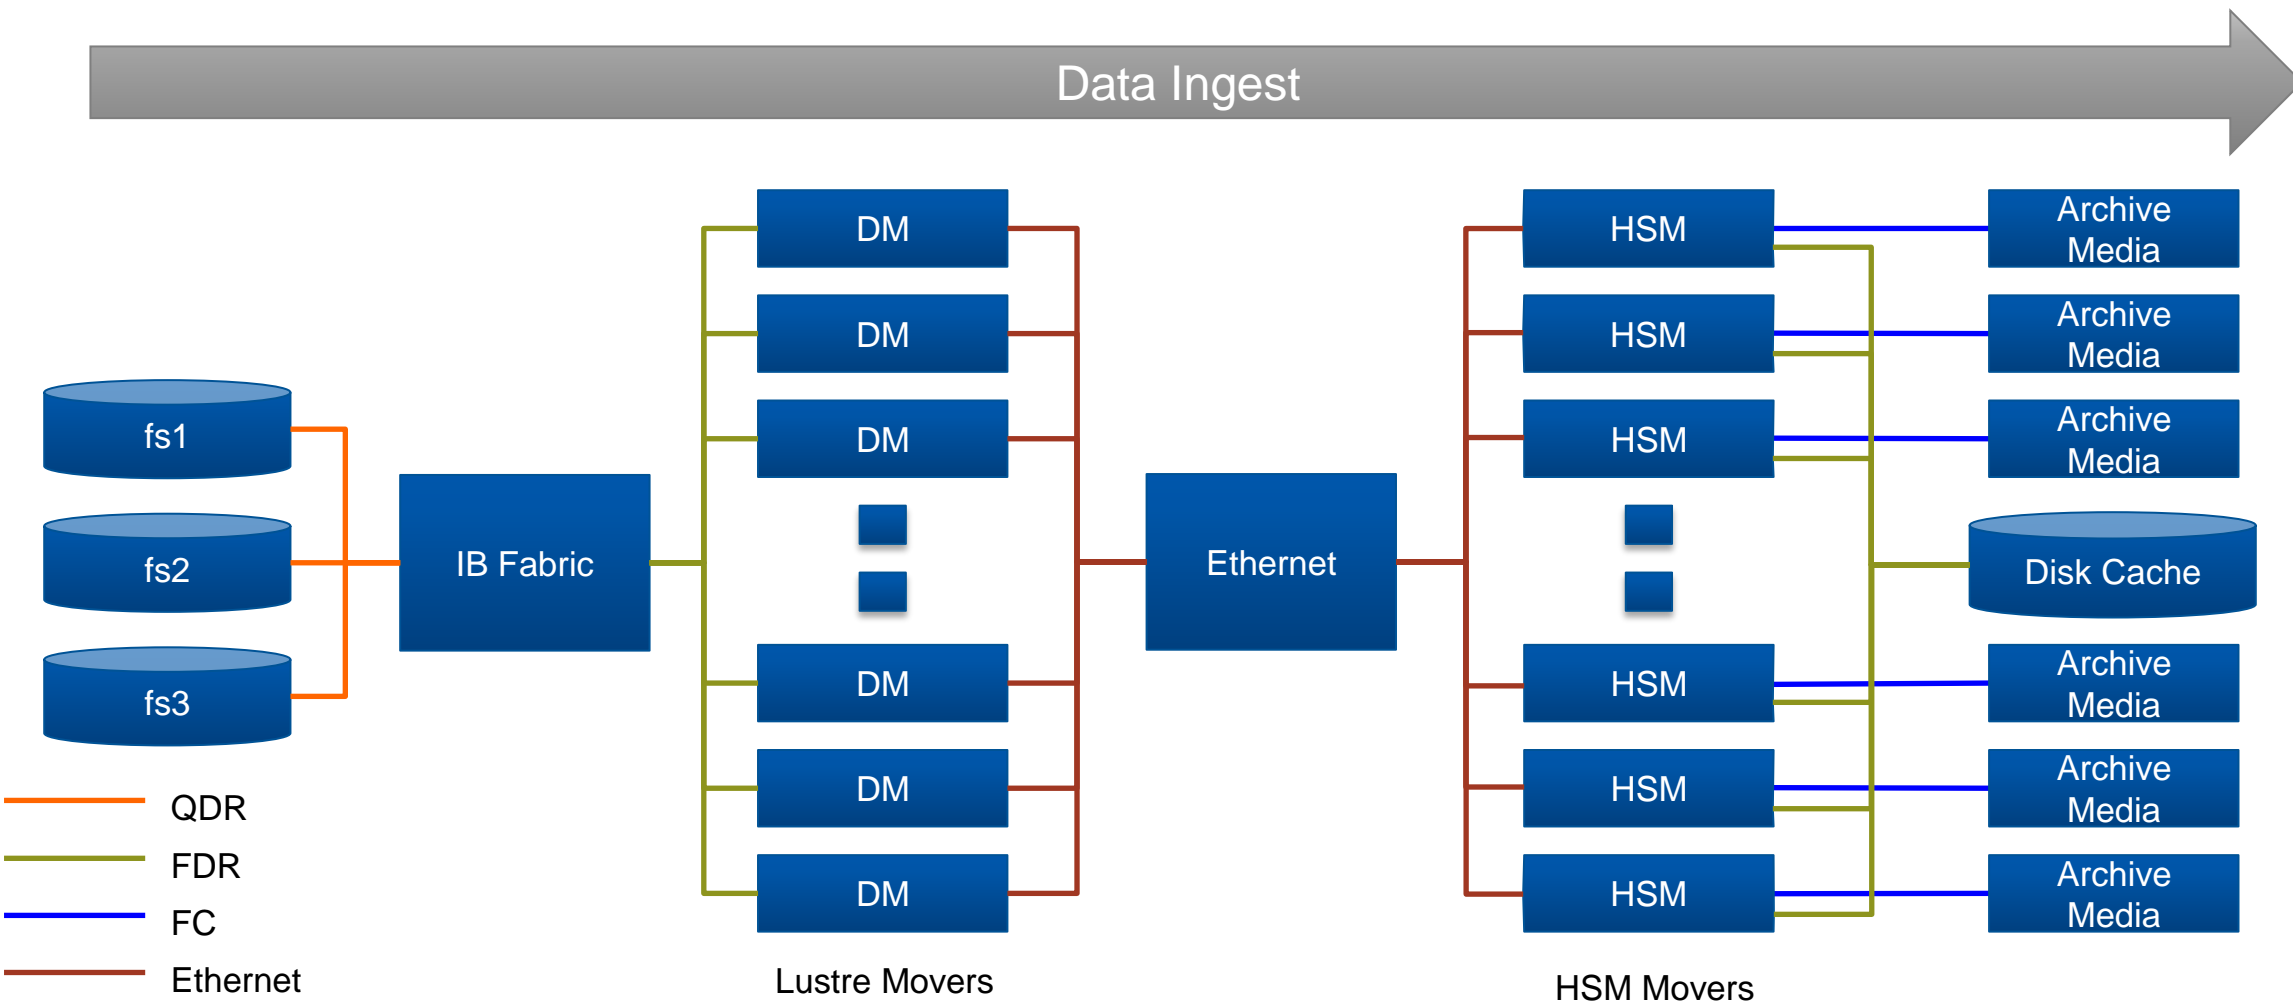

- Break vendor lock-in
- Multigenerational data preservation

Powered By



Efficient storage management with tiers

Traditional HSM for HPC – Complex



Simplified Data Management for Big Data and HPC



Common Access Protocols:
Lustre, NFS, SMB, HTTP, FTP

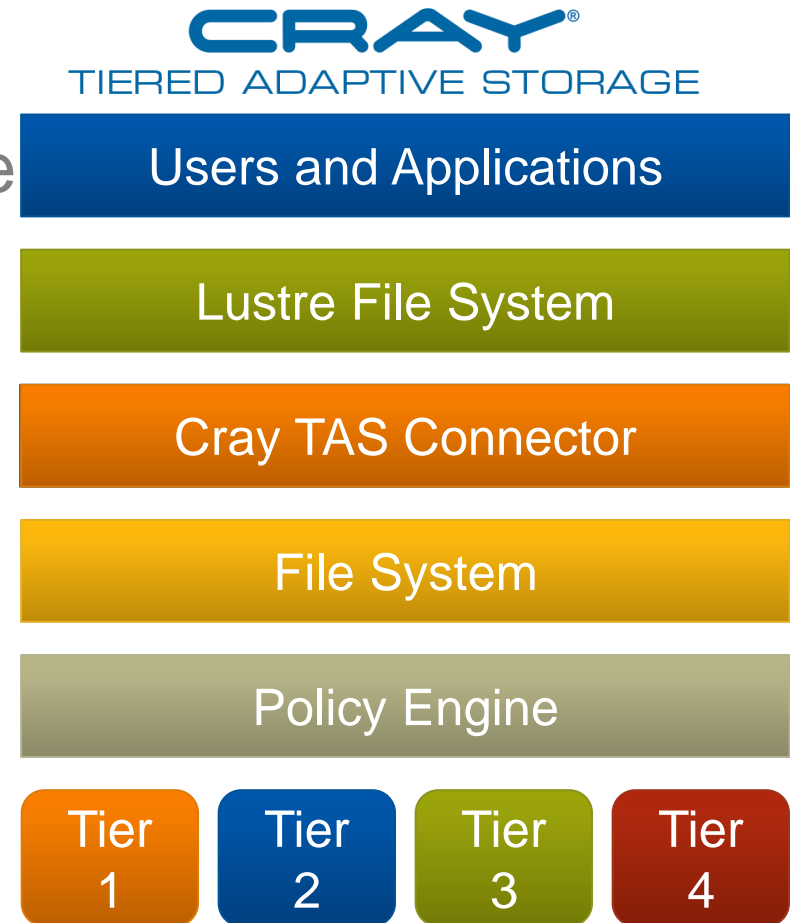


TIERED ADAPTIVE STORAGE

- QDR
- FDR
- FC
- Ethernet

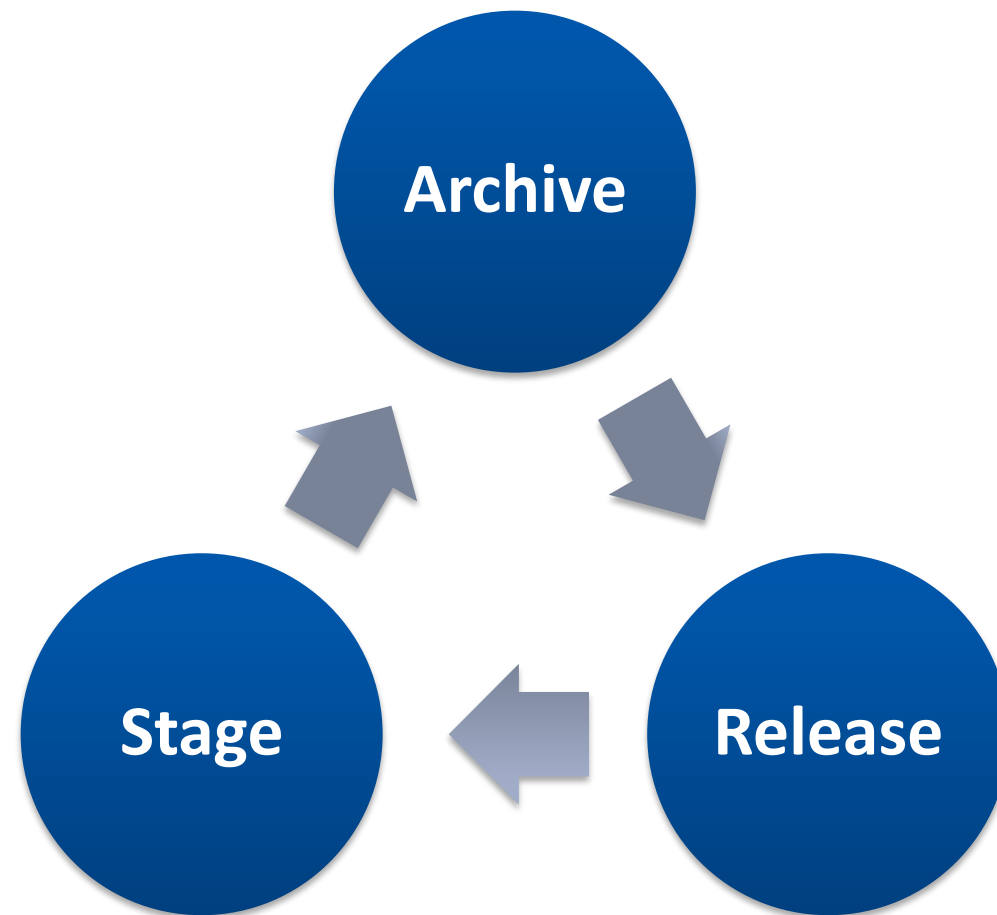
Protect Data over its Lifespan

- **Transparent tiering for users**
 - Data always accessible regardless of tier
 - File system appears infinitely large
 - Files always visible from the file system
- **Automated data management**
 - Policy-based data management
 - 24x7 data management
 - Multiple copies and disaster recovery
- **Works with any Lustre 2.5**
 - Cray TAS Connector for Lustre HSM



Policy-based Data Movement

- Familiar Actions & Policies
- Transparently *Archive* from disk cache to archive media
- Manage disk space or *Release* archived files from disk
- Automatically *Stage* released files back when accessed



Sustain Long-term Repositories



- **Open data format**

- Based on POSIX TAR
- Data is accessible without TAS
- No vendor lock-in

the default setting is

Open

- **Data protected at scale**

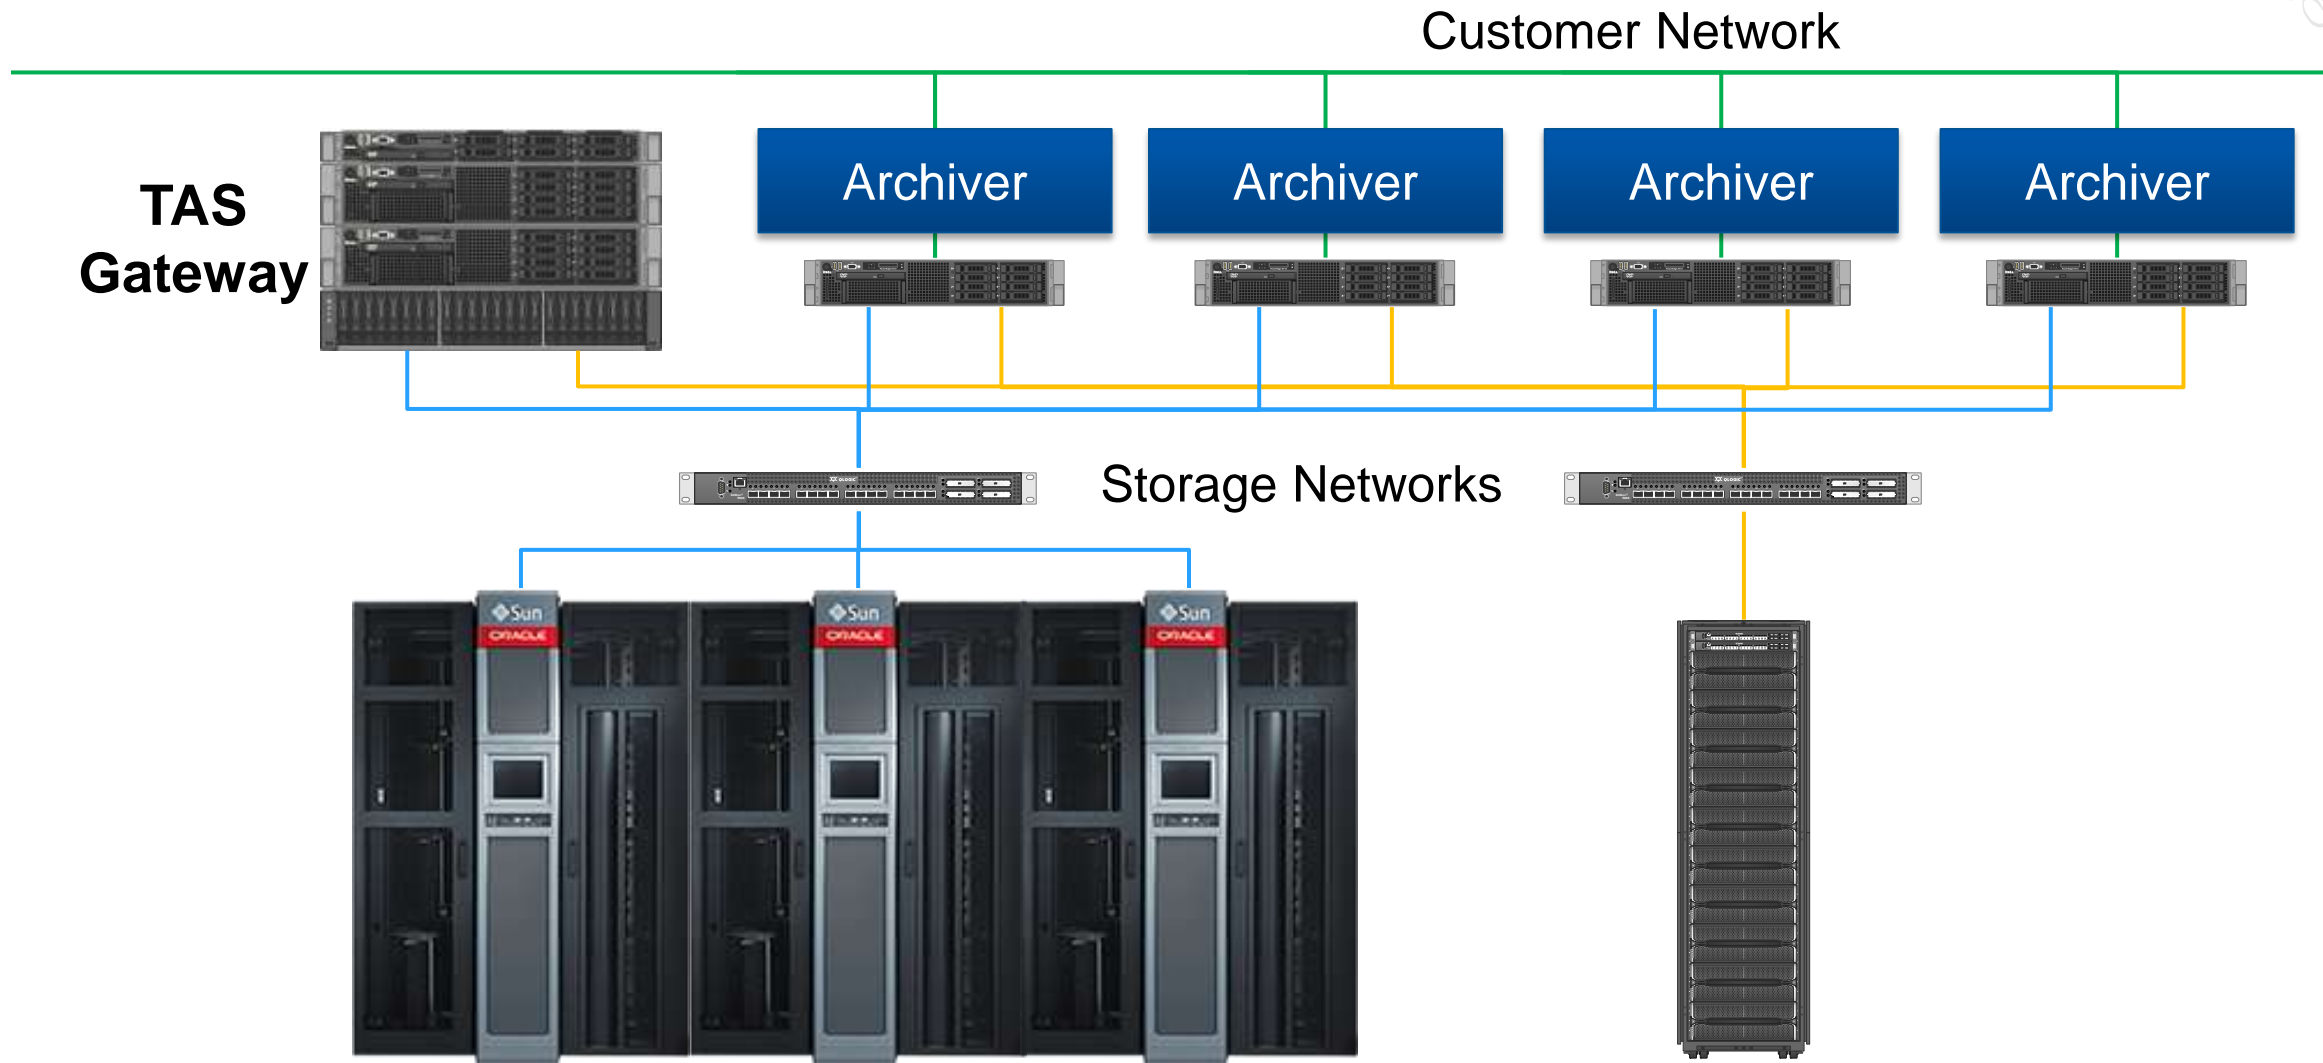

- Support for 100's of PB of managed data
- Integration with Lustre HSM

- **Future-ready technology migration**

- Support for multigenerational data management
- Migrate with technology



Parallel Archiving with Cray TAS



Cray TAS – Tiered Data Management Summary



Protect and access your high-value data when you need it, for as long as you need it

- **Store and Protect**

- Up to 5 copies—across any media
- Transparently migrate data across tiers, from ingest to archive

- **Access Data Forever**

- All data always accessible to apps and users
- Your choice of file protocol

- **Stay Open – Sustainable Infrastructure**

- Open formats break vendor lock-in
- Data transparently migrates across generations of storage infrastructure, disk and tape

CRAY[®]
TIERED ADAPTIVE STORAGE



- お問い合わせ先:
クレイ・ジャパン・インク
プリセールスチーム
磯野智之 isono@cray.com

**The future is
seldom the same
as the past**

Seymour Cray
June 4, 1995

