

Special Report

Perspectives on HPC Storage and Interconnects in the Second Half of 2021

Mark Nossokoff and Bob Sorensen
January 2022

HYPERION RESEARCH OPINION

Storage and interconnects continue to be important elements of HPC system architecture and are expected to take on even greater significance with increasing demanding and diverse requirements driven by both traditional compute-intensive HPC mod/sim workloads and data-intensive AI workloads. The significance is reflected in recent market data and near-term forecasts, technology adoption and utilization trends, industry announcements in the second half of 2021, and future storage and related technology research direction of Hyperion Research.

Storage market insights discussed in this paper includes:

- Broader HPC on-premises market forecasted to be in line with traditional near-term growth but with a slight decline from 2024 to 2025.
- HPC cloud storage projected to approach \$9 billion in 2025 with a CAGR almost twice that of on-premises storage.
- Dell Technologies and DDN continue to be the leading on-premises systems' storage provider and independent storage provider, respectively.

Technology highlights summarized are:

- NFS and Lustre remain the primary file systems utilized, with other maturing and evolving file systems rising in adoption, including BeeGFS, CEPH, OneFS, WekaFS, DAOS, and PanFS.
- Ethernet surpassed InfiniBand as the preferred system interconnect, with Omni-path appearing third overall and as second within academia.
- Interconnect adoption is split largely between ethernet and Infiniband for both system networks and storage networks, with 10-Gbit ethernet being the most broadly adopted ethernet rate.

A list of notable industry activity in 2H21 covered includes:

- Refresh of the Pawsey Supercomputer Center machine to include a 130PB object storage upgrade
- Release of IBM Spectrum Fusion to address increased container workloads also reflects the fruits of Red Hat acquisition
- Lustre in the cloud announcements reflecting increased parallel file system investments across AWS, Microsoft Azure, and Google Cloud Platform
- DNA storage financing totaling \$190M between CATALOG and DNA Script
- Interconnect innovations addressing latency at both internode (Rockport Networks) and intranode (accelerating CXL adoption plans)


Note: This page is intentionally blank.

MARKET DATA

Similar to prior forecasts, the most recent Hyperion Research market forecast detailed in Figure 1 shows storage remaining the 2nd largest HPC area and the highest growth area, culminating in almost \$8.5B in projected storage revenue in 2025. New to the forecast is a slight decline in the out years of the forecast period. This reflects the waning of the initial wave of exascale-class spending, the overall cyclical nature of the market (including fluctuating budgets of major governmental and regional datacenters) and anticipated accelerated adoption of the cloud for storage deployments

FIGURE 1

On-premises HPC Market



Storage growth in the use of clouds for HPC workloads is even stronger than on-premises, as shown in Figure 2. Cloud storage spending is expected to grow almost twice as fast as on-prem storage over the forecast period, to almost \$3 billion in 2025. Recent research indicates roughly 1/3 of HPC cloud spending goes towards storage and storage-related items. There is a potential upside to user spending in the cloud for storage should attention focus on architectural challenges relative to data locality and where data should best reside (on-premises vs. cloud vs. edge). The cloud revenue reflects what users are spending to run their workloads in the cloud, as opposed to what the CSPs are spending to deploy HPC resources for users.

FIGURE 2

Storage in the Cloud





Figure 3 shares HPC sites' storage vendor preferences deployed across their infrastructure. This data and subsequent technology ecosystem data are results from Hyperion Research's most recent annual site survey that looked at 141 HPC sites, representing a total of 252 storage systems. The study emphasized the Industry sector with 82 respondents from Industry respondents, 23 from Government, and 36 from Academia.

Dell Technologies remains the most preferred storage vendor at the survey sites and most preferred by Industry. IBM was the 2nd most preferred overall and across each sector. HPE/Cray was third overall and most preferred in Government. DDN continues as the most preferred independent storage vendor, most preferred in Academia and 4th overall.

FIGURE 3

Storage in the Cloud



TECHNOLOGY TRENDS

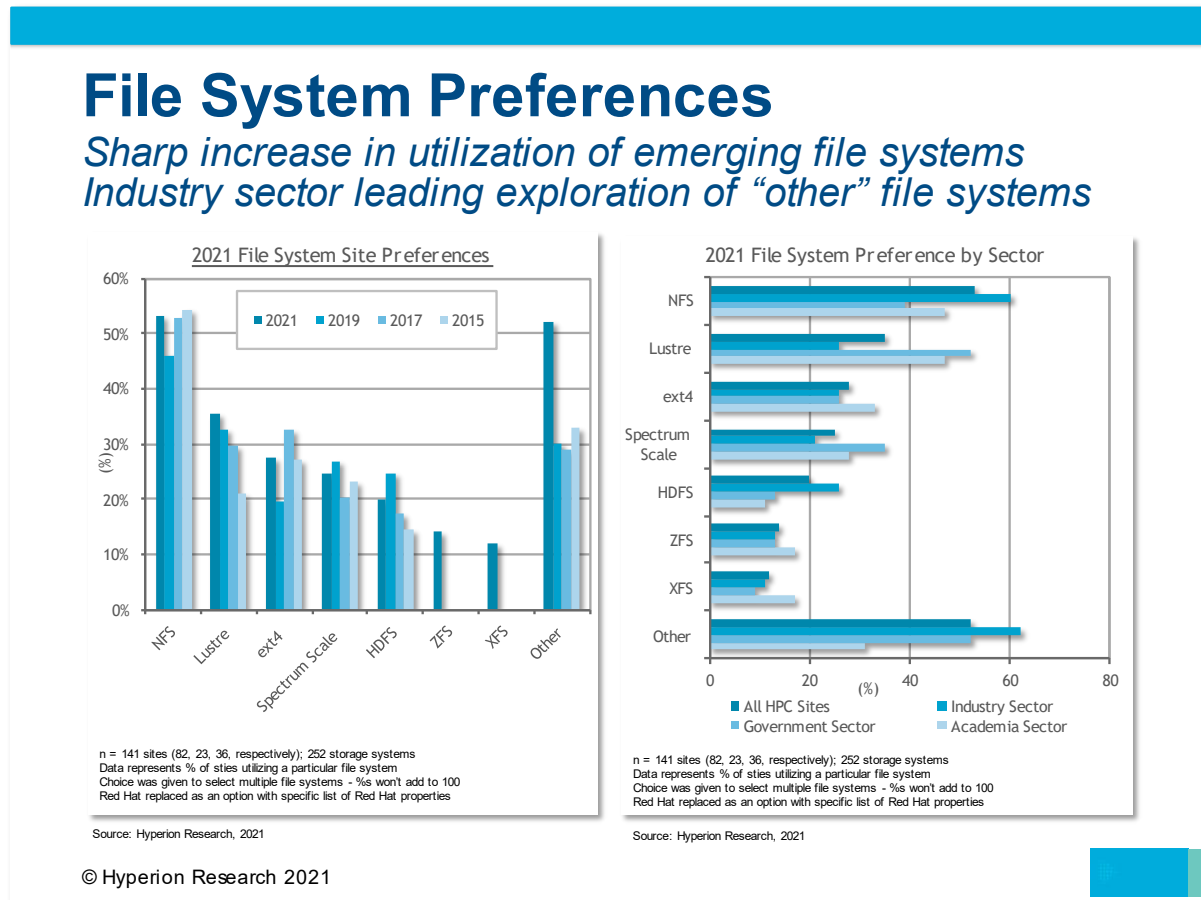
There is no shortage of technology development and innovation occurring within the HPC storage interconnect ecosystem. File systems and interconnects are profiled here.

File Systems

Figure 4 summarizes the file system results of the most recent Hyperion Research site survey. Users from 141 sites representing 252 storage systems were asked to share what file systems are deployed across all their systems. There was an average of 1.8x file systems deployed at each site.

FIGURE 4

File System Preferences



NFS and Lustre remain the most deployed file systems, with both experiencing growth over the past year. Lustre is heavily favored by the Government sector while the Industry sector favors NFS.

Ext4 emerged as the third most deployed while Spectrum Scale dropped to fourth and showed a slight year-over-year decline, particularly at government sites. That said, there was a noticeable increase in adoption of Spectrum Scale within the Industry sector. This reflects increasing adoption of HPC-enabled AI techniques within the industrial sector to support the growing use of AI. This sector also often requires typical enterprise IT data services (including snapshots and replications) which are areas of strength for Spectrum Scale.

File systems in the "Other" category were a combination of file systems growing in adoption since the prior iteration of the survey published in 2020 (e.g., BeeGFS, CEPH, OneFS, WekaFS, DAOS, PanFS) and those with declining adoption (e.g., GFS, Gluster, PVFS/OrangeFS).

Interconnects

One of two interconnect architectural approaches are typically employed within an HPC system:

- Separate, independent interconnects for interprocessor communication (IPC) and storage I/O
- Single, converged system interconnect for both IPC and storage traffic

Acknowledging these different architectural approaches, a distinction was made in the most recent Hyperion Research survey between system interconnects and storage interconnects.

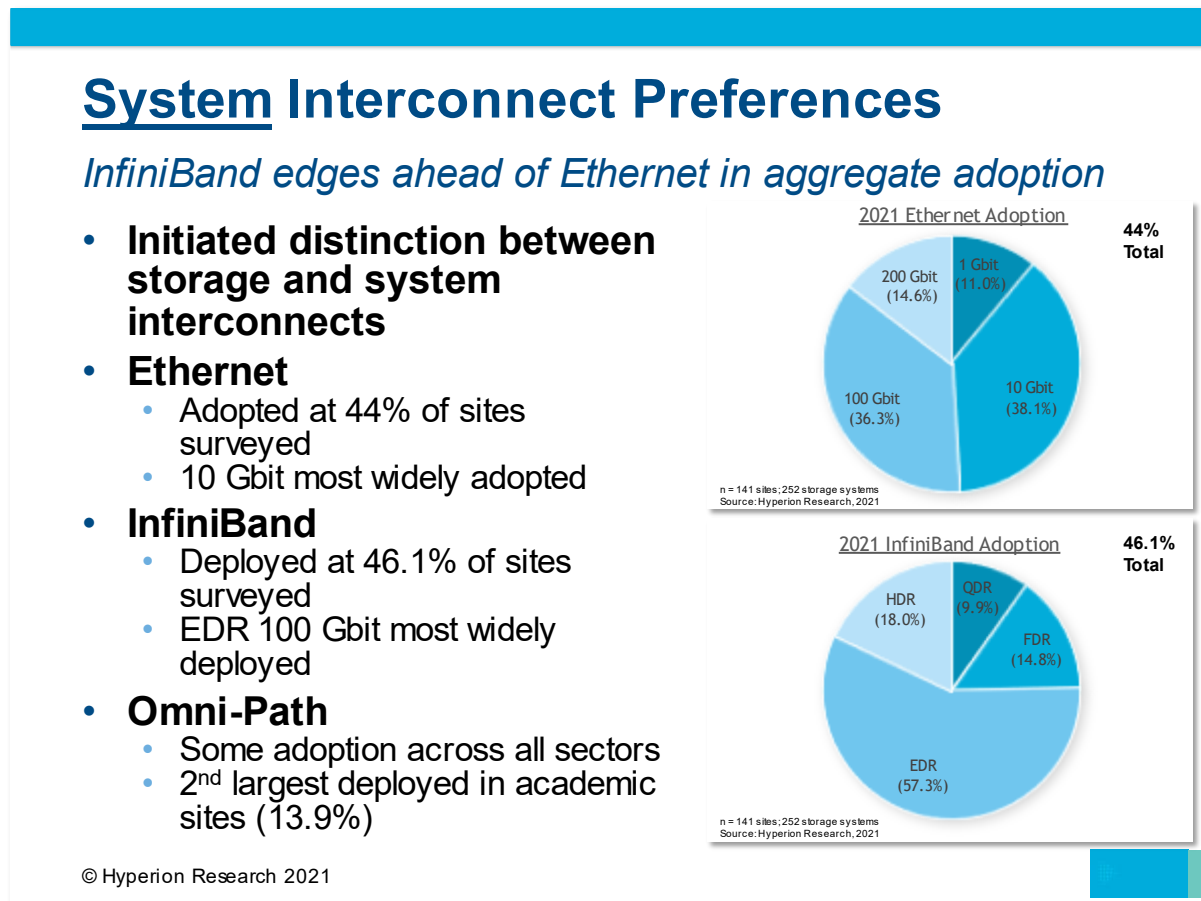
On the system interconnect side, summarized in Figure 5, respondents indicated a slightly higher adoption of InfiniBand over Ethernet (46% - 44%), with over half of InfiniBand users deploying EDR IB.

For those that indicated deploying ethernet, 10 Gbit was the predominant speed. This suggests that a number of HPC sites don't require the highest performing interconnect solutions and are content with cost-effective, "good enough" system performance. It also reflects the weighting of Industry respondents for this site survey.

OmniPath emerged as the third largest system interconnect and second largest at academic sites.

FIGURE 5

System Interconnects




Shown in Figure 6, ethernet showed a large increase in storage interconnect adoption and emerged as the preferred storage interconnect, being deployed at 47% of the sites, with 10 Gbit being the predominant line rate, similar to ethernet adoption for the system interconnect. InfiniBand was second and showed a small decline.

There was a material difference in preference between sectors, with Industry leaning toward ethernet and academia leaning towards IB. Government was evenly split between the two.

Fibre Channel continues to decline as a preferred storage interconnect for HPC systems.

FIGURE 6

Storage Interconnects



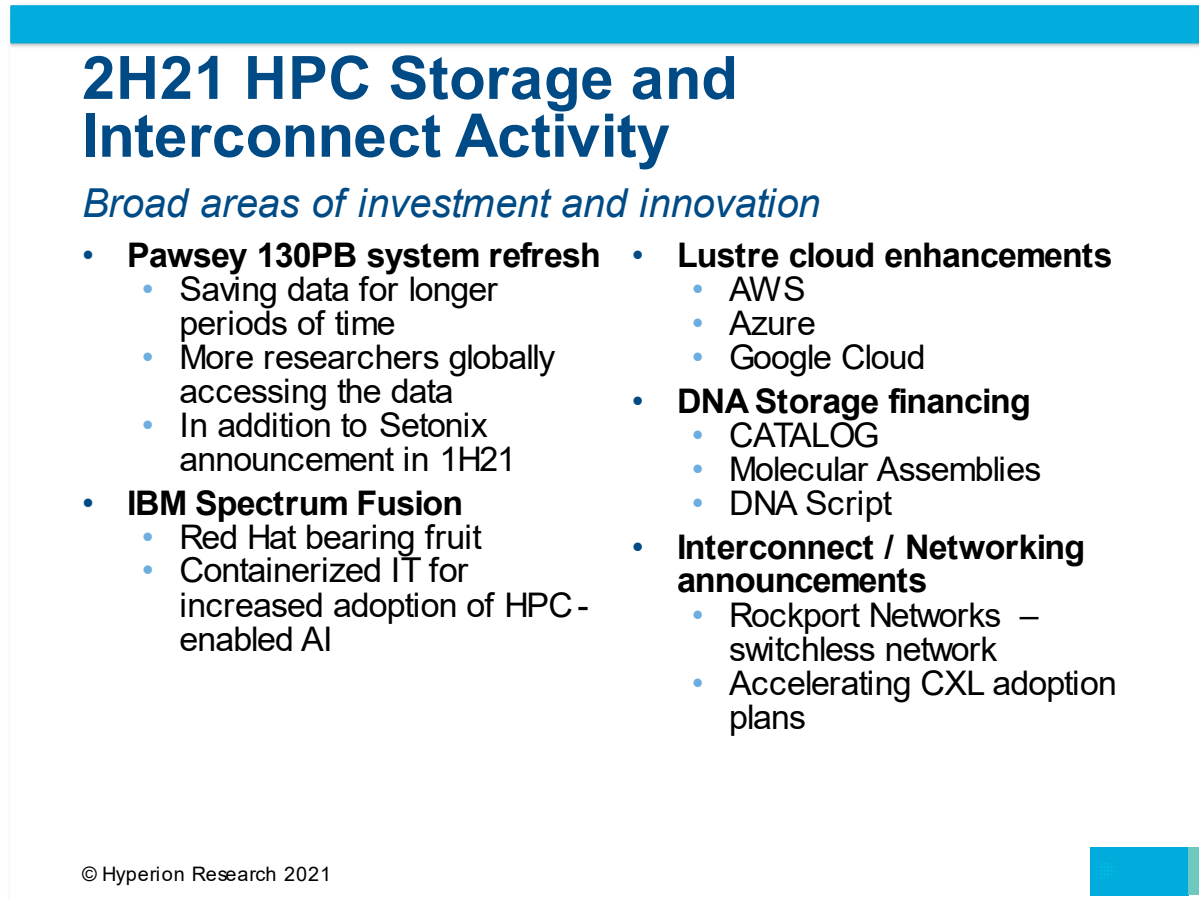
- **Initiated distinction between storage and system interconnects**
- **Sector preference differences**
 - Industry leans ethernet
 - Government split between ethernet and InfiniBand
 - Academia leans InfiniBand
- **Omni-Path large part of other**

INDUSTRY ACTIVITY IN THE SECOND HALF OF 2021

The items summarized in Figure 7 cover a broad range of representative investments and innovations occurring across the spectrum of HPC storage and interconnect ecosystem. The level and breadth of activities cited are typical of the wide-ranging events that occur within the HPC community across any six month period.

FIGURE 7

2H21 HPC Storage and Interconnect Activity



Pawsey Supercomputer Center announced a sizeable storage upgrade of 130PB of object-storage that will improve data availability, data transfer speed, and overall storage capacity. This is incremental to the previously announced and currently being deployed Setonix system.

While many users are evaluating how to best leverage containers for their HPC workloads, IBM began shipping Spectrum Fusion container-native storage solution for OpenShift. As AI, hybrid cloud, and HPC converge, container-base solutions are designed to streamline migration of those workloads to and between clouds.

Three of the top CSPs indicated enhancements to their Lustre-based offerings, indicating continued investment to provide the necessary scale and support for running complex HPC workloads in the cloud.

On the future storage technology front, there are several announcements of substantial new funding rounds in the DNA Storage community. Targeted at long-term archiving applications, DNA-based storage technology represents continued interest and investment in new, novel storage solutions.

Concluding with several interconnect activities, at the external, inter-node network, Rockport Networks came out of stealth, taking a unique switchless approach for HPC system interconnects aimed to dramatically improve HPC workload completion performance. And at the internal, intra-node, CPU/GPU/memory interconnect hierarchy, there appeared to be an acceleration of announcements from vendors looking to address latency via adoption of CXL.

FUTURE OUTLOOK

2022 portends to be as dynamic, if not more so, than 2021. Areas Hyperion Research intends to pay particular attention are identified in Figure 8.

FIGURE 8

Future Research Direction

Future Research Direction

Broad range of topics across diverse storage ecosystem

- **Grow on-prem census data**
 - Use cases
 - Temporal, durable
 - File, block, object
 - Scratch, user, home directory, project, campaign, archive
 - Capacity
 - Internal to servers and compute nodes
- **Expand cloud storage coverage**
 - Cloud storage landscape model
 - Cloud workload usage and requirements
- **Broaden general storage and data management areas**
 - Storage consumption models
 - Data and workflow management
 - Migration of HPC workloads to the cloud
 - Impact of edge computing
- **Develop broader technology coverage**
 - Storage for Containers
 - POSIX compliancy
 - Interconnects
 - Additional memory topics

© Hyperion Research 2021

There's no shortage of areas to explore. Storage and interconnect research agenda items include (in addition to current market coverage):

- Developing a deeper picture of how storage is deployed and utilized on-premises
- Refining the cloud storage landscape model and how the cloud can best be utilized to optimize HPC storage solutions
- Broadening overall data management approaches, including utility- and as-a-service consumption models and workflow management from edge to cloud to on-premises, inclusive of edge computing and computational storage
- Growing coverage of emerging HPC and AI related technology areas including containers, POSIX compliancy, deeper intra-node and compute-level interconnects and additional memory topics such as Big Memory and in-and-near memory computing

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology, and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue

St. Paul, MN 55102

USA

612.812.5798

www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2022 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.