

Special Report

2022 HPC End Users Perspectives on Trends and Forecast in HPC Storage and Interconnects - Key Findings

Mark Nossokoff, Jaclyn Ludema, and Earl Joseph
February 2023

HYPERION RESEARCH OPINION

Key findings from a recent Hyperion Research study indicate that HPC storage solutions, and associated storage and system interconnects, continue to be critical for HPC infrastructure to deliver optimal capabilities and provide the fastest time to results for the systems' users. Data-intensive workloads driven by new AI/ML/DL workloads, increasing scale of traditional HPC modelling and simulation, emerging edge computing, and emerging composable systems are placing greater demands and requirements on HPC storage systems. Insights into the critical factors driving these and other trends are detailed in the 2022 iteration of Hyperion Research's annual MCS end users' study, *2022 HPC Multi-Client Study: Trends and Forecasts in HPC Storage and Interconnects*. Key findings from the report are summarized in this document.

Hyperion Research conducts an annual end-user Multi-Client Study (MCS) to measure and track key trends across the spectrum of the HPC market. The latest iteration of the MCS encompassed 181 HPC end-user sites with 3,830 HPC systems. Reports produced as a result of the study include:

- AI and HPDA Usage and Future Technology Trends
- Vertical/Application Workload Areas and Technical Computing System Software and Middleware
- Use of Public/External Clouds for HPC Workloads, Trends, and Drivers
- Processors, Coprocessors/Accelerators, and HPC Budgets
- Trends and Forecasts in HPC Storage and Interconnects

Overall, the HPC market is expected to grow by about 6% to 7% over the next five years and reach \$40 billion by 2026. HPC storage will be the highest growth areas of the HPC market, driven by the capacity demands of AI, machine learning, deep learning, and other analytics methods as well as increasing use of iterative modeling and simulation (M&S) methods. Feeding this demand in storage is innovation and collaboration within all elements of the HPC storage ecosystem (e.g., file systems, storage systems, interconnect, and networking solutions).

SELECT KEY FINDINGS

Select key findings in this report include dynamics relative to storage capacities, file system adoption, storage system vendor preferences, interconnect technology adoption and vendor preferences, and storage media adoption.

Top Industry Sites Emerged with Demanding HPC Storage Requirements

Respondents reported an average of 148 PB of add-on storage per HPC site, nearly double the average among last year's sample. This was primarily due to increasing storage capacities among industry sites surveyed. Average industry site storage capacity is over 200 PB and average site single-file bandwidth (almost 3.6 TB/s), driven by seven sites indicating capacity of more than 1 EB.

Average Site Storage Capacity Across All Sectors Doubled

Average storage capacities across all surveyed sites approximately doubled from 74 PB to 148 PB. Excluding the sites indicating more than 1 EB of capacity, the average capacity across all sectors was 47.1 PB.

File System Utilization is Relatively Stable

There was little change relative to file system utilization from the prior iteration of this study. NFS remains the dominant scale-out file-based file system, while Lustre and Spectrum Scale (formerly GPFS) remain the top parallel file systems. That said, new entrants are emerging to challenge the status-quo, taking advantage of green-field opportunities to leverage heterogeneous systems supporting both HPC and AI applications, as well as hybrid systems that require seamless and performant operation whether on-premises or in the cloud

System Interconnect Architecture Expected to Shift with Sites' Next HPC Procurements

System interconnect architectures are expected to shift from a preference of independent node-node and node-storage networks to a preference of converged networks with respondents' next HPC procurement.

Ethernet and InfiniBand Leading Interconnect Technologies for Different System Architectures

100 Gbit Ethernet (and Ethernet overall) is the preferred technology for independent networks (both node-node and node-storage), while 100 Gbit InfiniBand (and InfiniBand overall) is the preferred interconnect for converged networks.

Tape Persists at On-Premises HPC Sites

Tape continues to be utilized at roughly half of the sites surveyed, with average tape capacity at those sites being approximately 2.5x greater than average disk (both HDD and SSD) capacity across all of the sites.

Sites Are Bullish with Storage Budget Expectations

49% of the sites surveyed expect to increase their storage budgets by more than 5% over the following year. This is up compared with 36% of respondents who answered similarly in last year's study.

FUTURE OUTLOOK

Storage and interconnects are projected to continue to increase in relevance from both business and architectural perspectives. In 2022, the HPC storage segment was the second largest on-premises HPC segment with almost \$6 billion being spent, representing 20% of the overall HPC market. Viewed another way, for every dollar spent on HPC technical servers in 2022, an additional \$0.43 was spent on storage. This segment is expected to remain the 2nd largest segment and is projected to grow at the highest rate at 8.6% over the forecast horizon.

Interest in HPC storage and interconnects continues to grow across several important dimensions, including vendors' R&D investments and share of budgets being spent by users. Without vendors' continued innovation in these key areas of the HPC and AI solution ecosystem, the scale of users' research and time to achieve their results will be muted as a result of unbalanced and sub-optimal system performance.

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 2023 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.
