

Special Study

Snapshot of HPC File System Landscape

Mark Nossokoff and Melissa Riddle
January 2022

HYPERION RESEARCH OPINION

Results from the most recent Hyperion Research Multi-Client Study (MCS) show continued evolution within HPC file system adoption. The *2021 HPC Multi-Client Study: Trends and Forecasts in HPC Storage and Interconnects* report details the storage and interconnect-related responses and analysis associated with the 252 HPC storage systems deployed at 141 global HPC respondent sites representing 2,006 HPC systems. The survey sample had a total of 338 file systems. This Special Study examines file system utilization.

File systems bridge the gap between the applications consuming the data and the physical devices where the data resides, while also being tasked with managing an overall storage solution that is responding to new requirements including evolving pricing models, use cases, and data access patterns. Key takeaways from this year's study include:

- Technical requirements for traditional HPC storage use cases (e.g., home directory, scratch, archive) continue to evolve and fragment, leading to specialization by storage system vendors.
- Data at scale stemming from increasing cloud utilization, large data volumes created on-prem and at the edge, and data locality is causing users to reconsider appropriate application of file, block, and object storage.
- Complex data access patterns including increasingly mixed bandwidth and throughput I/O profiles driven by AI, machine learning, and deep learning workloads are exposing the limits of the current incumbent file systems and driving exciting innovation within the HPC storage community.
- Increased adoption of HPC-enabled AI techniques by traditional enterprise IT datacenters to run their data analytics workloads is creating demand for improved ease of use and integration of data services not typically required by traditional HPC datacenters.

THE BROADER HPC MARKET

The Add-On Storage segment of the broader on-prem HPC market continues to display the highest forecast growth rate of all the segments with a 5-year CAGR of 7.9%. This may well be a conservative estimate as the data demands of HPDA and AI workloads shows no sign of slowing in the foreseeable future and may exhibit even higher growth than currently forecasted. Table 1 summarizes the forecast revenues of the various HPC segments.

TABLE 1

On-Prem HPC Revenues by the Broader HPC Market Areas (\$ Millions)

Revenue (\$M)	2020	2021	2022	2023	2024	2025	CAGR '20-'25
HPC Servers	\$13,523	\$14,550	\$16,947	\$18,565	\$19,947	\$19,901	8.00%
Add-On Storage	\$5,423	\$5,931	\$6,981	\$7,835	\$8,479	\$8,454	9.30%
Middleware	\$1,590	\$1,736	\$2,035	\$2,244	\$2,426	\$2,420	8.80%
Applications	\$4,600	\$4,913	\$5,626	\$6,060	\$6,396	\$6,353	6.70%
Service	\$2,146	\$2,254	\$2,531	\$2,674	\$2,767	\$2,738	5.00%
Total Revenue	\$27,283	\$29,383	\$34,121	\$37,378	\$40,015	\$39,867	7.90%

Source: Hyperion Research 2021

FILE SYSTEM LANDSCAPE

In the most recent survey, 141 supercomputer sites representing 2,006 HPC systems were asked which file systems were used on their systems. As respondents were given the choice to select multiple file systems, the results reflected 338 file systems deployed yielding an average of 2.4 file systems utilized per site. This is up from an average of 1.8 from the prior MCS. As a result of the increased number of file systems deployed across the survey sites, increases in adoption can be viewed as an amplified positive and declines in adoption can be viewed as an amplified negative. The table below shows the percentage of respondents' sites utilizing each file system. Respondents were able to select multiple responses.

NFS remains the predominant file system deployed and has grown to being adopted at 53.2% of the sites surveyed, up from 45.9% the previous year. The industry and academic sectors both increased deployments of NFS while government adoption of NFS declined. NFS's strength in this survey is largely due to intentional over-representation of the commercial industry sector and the smaller scale of their HPC systems, which don't require a large-scale parallel file system.

The second most widely deployed file system is Lustre. Its adoption has grown to reach 35.5% from 32.5%. Industry adoption of Lustre was responsible for the majority of the growth, increasing from

18.9% to 25.6% of sites deploying it. Government adoption showed modest growth while academia remained flat.

Prior to 2020, Red Hat had been listed as a single "file system" for survey respondents to choose from. In this latest survey, several different Red Hat file systems were individually listed, including ext4, XFS, Gluster, and CEPH. ext4 emerged as the leading Red Hat file system property and is singly compared with the prior Red Hat choice for historical purposes. It is likely ext4 has been deployed both as a stand-alone file system and in conjunction with a higher-level parallel file system.

GPFS/Spectrum Scale saw a slight decline, from 26.8% to 24.8%. The decline was driven by a steep decline in adoption in the government sector (50% to 34.8%) while the industry sector exhibited a marginal growth in adoption (17.1% to 20.7%). The increase in industry adoption reflects the trend of industry's traditional enterprise IT adopting of HPC-enabled AI infrastructures for their rising number of data-intensive applications. GPFS/Spectrum is regarded as having more robust traditional enterprise IT-required data management features (replications, snapshots, etc.) than the more HPC-focused parallel file systems, such as Lustre.

HDFS adoption has declined to 19.9% from the previous 24.7% adoption it enjoyed. Initially deployed largely for ingesting data for big data workloads, HDFS is falling out of favor in lieu of alternative methods that have evolved to ingest modern HPDA and AI datasets.

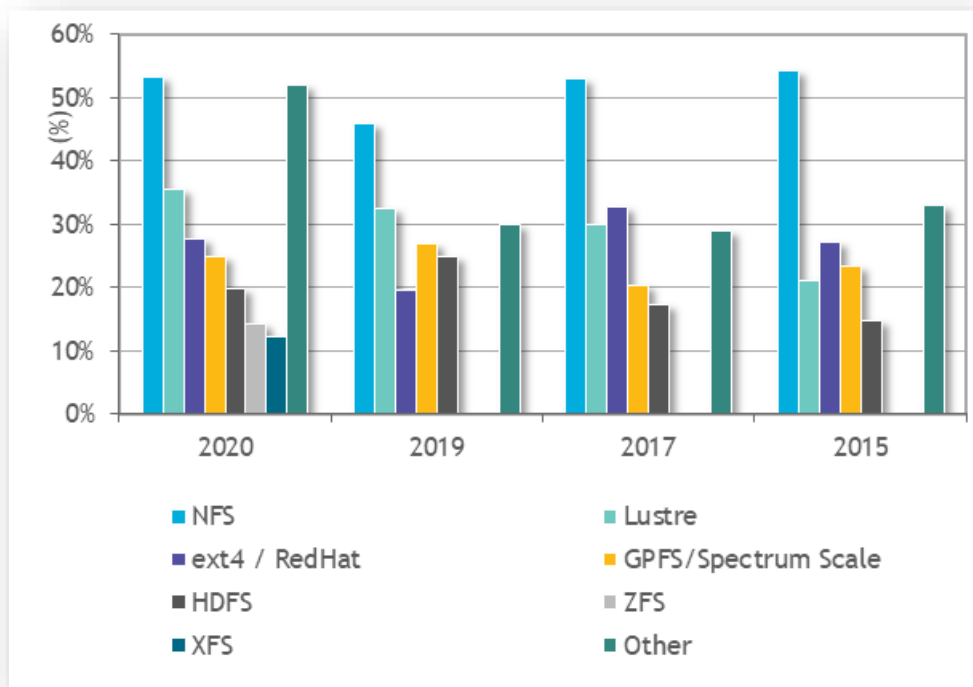
While not meriting explicit representation in Figure 2, it should be noted CEPH is the only other previously individually tracked file system that exhibited growth in adoption, going from 4.6% to 8.5%. Each sector reflected growth in CEPH utilization. Growth was due in large part to the increasing amount of unstructured data being generated and stored, one of the key use cases of CEPH.

Although not new to HPC veterans, several file systems now being individually tracked showed relatively meaningful adoption rates. BeeGFS topped these files systems, being deployed at 9.9% of the 141 surveyed sites. Of the 14 sites utilizing BeeGFS, 9 were in industry.

Figure 1 shows the historical evolution of the HPC file system landscape while Figure 2 reflects the most recent results to the question: *What type(s) of file system(s) do you use on your HPC technical servers or clusters?* Respondents were able to select multiple responses.

FIGURE 1

File System Adoption 2015 -2020



Note: Data represents percentage of sites utilizing a particular file system.

Note: Respondents were given the choice to select multiple file systems - percentages will not sum to 100%

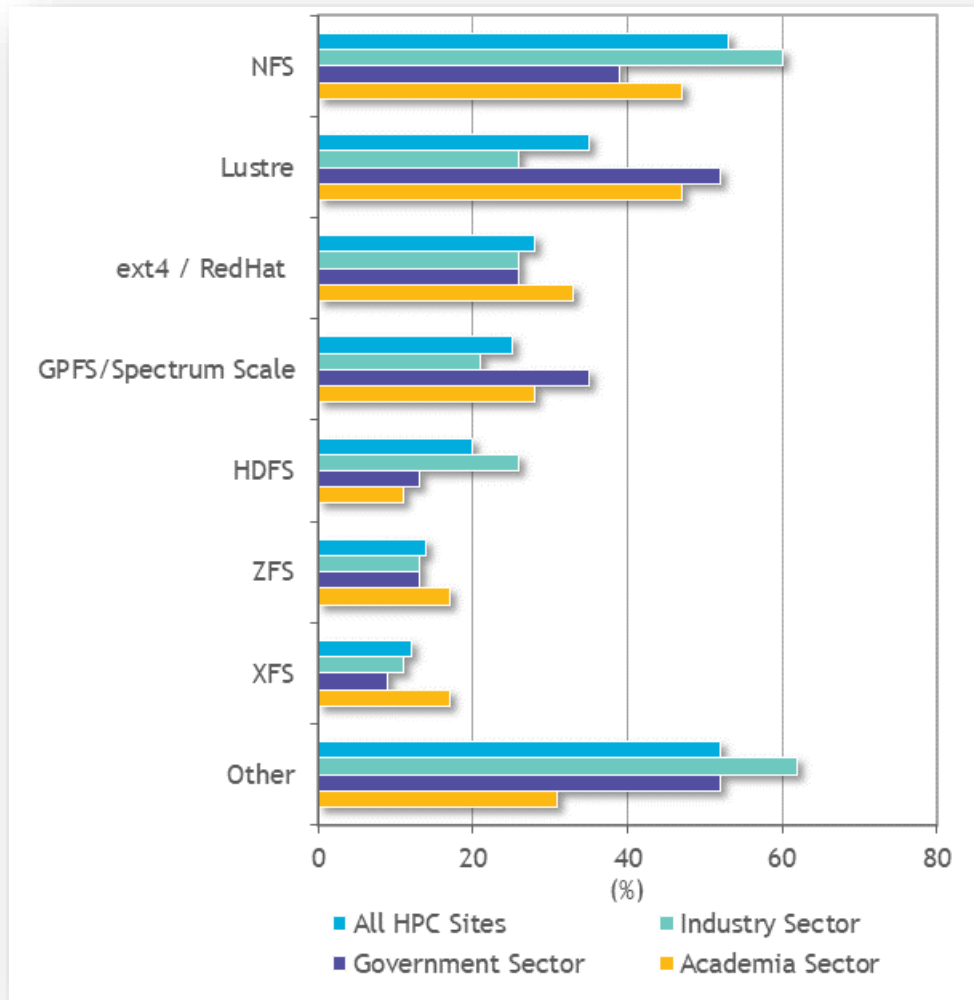
Note: Prior to 2020, Red Hat was listed as a single file system choice. In 2020, several file systems were provided as choices and ext4 was uniquely identified at this adoption rate.

Note: Prior to 2020, ZFS and XFS were captured in the Other category.

Source: Hyperion Research, 2021

FIGURE 2

2020 MCS HPC File System Usage



Note: Data represents percentage of sites utilizing a particular file system.

Note: Respondents were given the choice to select multiple file systems - percentages will not sum to 100%

Note: Prior to 2020, Red Hat was listed as a single file system choice. In 2020, several file systems were provided as choices and ext4 was uniquely identified at this adoption rate.

Source: Hyperion Research, 2021

FUTURE OUTLOOK

The future outlook for HPC file systems remains much like it did from the prior survey. Changes are expected to continue throughout the HPC file system landscape in the coming years as events are accelerating across all fronts: business, technology, workload shifts, and competitive upstarts. Business and technology should warrant the most significant changes as cloud adoption for storage is expected to rise, coupled with the emergence of utility and service-level based business models, regardless of whether the system is deployed on-prem or in the cloud.

Users should understand the increasingly diverse sets of workloads and new demands being placed on their file systems. Care should be taken to understand whether having different file systems to address different workloads (AI training, AI inference, simulation), different data types (structured vs. unstructured; file vs. block vs. object), and different access patterns (user directories, scratch, archive) are required or whether a single file system will be robust enough to address the needs of all the different scenarios. This is especially true for the increasing deployment of active archives across the HPC landscape. These systems tend to be optimized for object data and deliver extremely large capacities with moderate throughput (GB/s) requirements and minimal-to-no latency requirements.

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). We provide 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.