



HYPERION RESEARCH

# Perspectives on HPC Storage and Interconnects


November 2021

[www.HyperionResearch.com](http://www.HyperionResearch.com)  
[www.hpcuserforum.com](http://www.hpcuserforum.com)

**Mark Nossokoff**

# HPC Storage Demand Continues

*Data-intensive workloads primary storage growth driver*



Source: Hyperion Research, 2021

- **Storage historically the highest growth HPC element**
- **Storage represents ~ 20% of HPC spending and growing**
- **Slight on-prem decline in 2025**
  - Declining exascale spending
  - Cyclical nature of the market
  - Accelerated adoption of cloud


Area (\$M)	2020	2021	2022	2023	2024	2025	CAGR 20-25
Server	\$13,523	\$14,550	\$16,947	\$18,565	\$19,947	\$19,901	8.0%
<b>Add-on Storage</b>	<b>\$5,423</b>	<b>\$5,931</b>	<b>\$6,981</b>	<b>\$7,835</b>	<b>\$8,479</b>	<b>\$8,454</b>	<b>9.3%</b>
Middleware	\$1,590	\$1,736	\$2,035	\$2,244	\$2,426	\$2,420	8.8%
Applications	\$4,600	\$4,913	\$5,626	\$6,060	\$6,396	\$6,353	6.7%
Service	\$2,146	\$2,254	\$2,531	\$2,674	\$2,767	\$2,738	5.0%
<b>Total Revenue</b>	<b>\$27,283</b>	<b>\$29,383</b>	<b>\$34,121</b>	<b>\$37,378</b>	<b>\$40,015</b>	<b>\$39,867</b>	<b>7.9%</b>

Source: Hyperion Research, 2021

# HPC Cloud Storage Forecast

*Cloud storage '20-25 CAGR ~1.9x on-prem storage CAGR*

- **\$1.3B of cloud storage spend in 2020**
- **Storage represented 1/3 of total cloud spend**
- **Cloud storage CAGR ~ 1.9x on-prem storage CAGR**
  - Cloud: 17.3%
  - On-prem: 9.3%

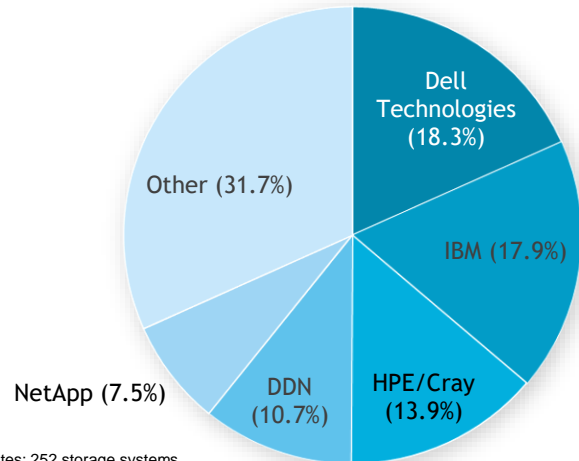


Source: Hyperion Research, 2021

# 2021 HPC Storage Vendor Preferences

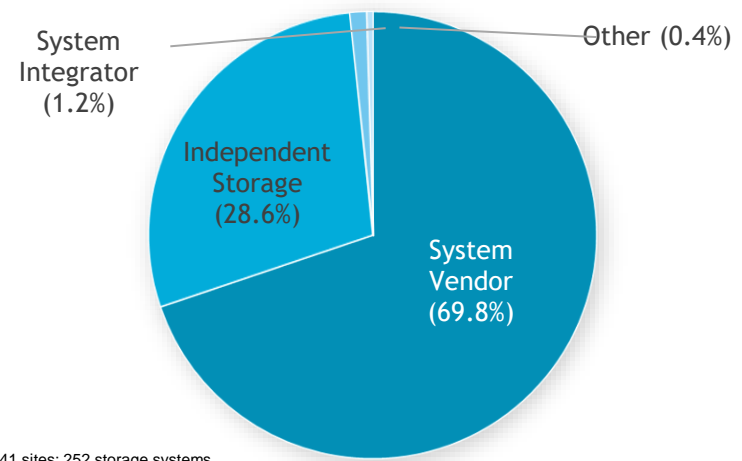
*Dell Technologies top preferred on-prem HPC storage vendor*

2021 On-Prem HPC Storage Vendor Preferences



n = 141 sites; 252 storage systems  
Source: Hyperion Research, 2021

2021 On-Prem HPC Storage Vendor Type



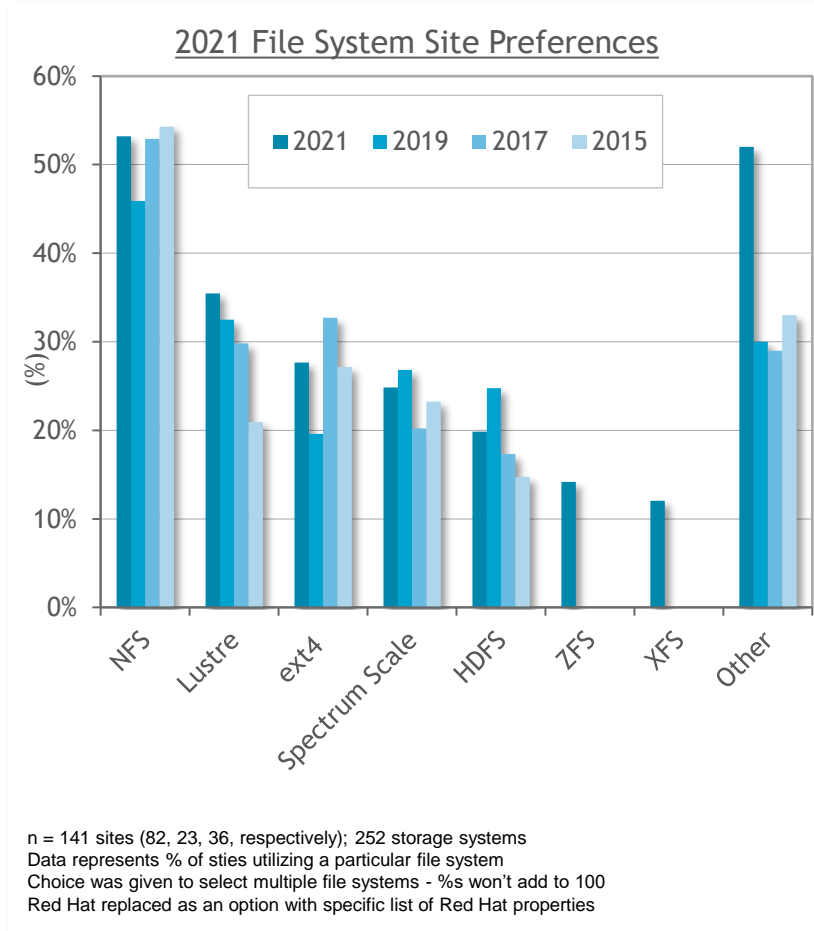
n = 141 sites; 252 storage systems  
Source: Hyperion Research, 2021

- **Dell Technologies** first overall; preferred by Industry
- **IBM** second preferred overall and in all sectors
- **HPE/Cray** preferred 3<sup>rd</sup> overall but tops in Government
- **DDN** 4<sup>th</sup> overall, preferred in Academia, and preferred independent
- **Top others: Seagate, Lenovo, Fujitsu, Huawei, ATOS**

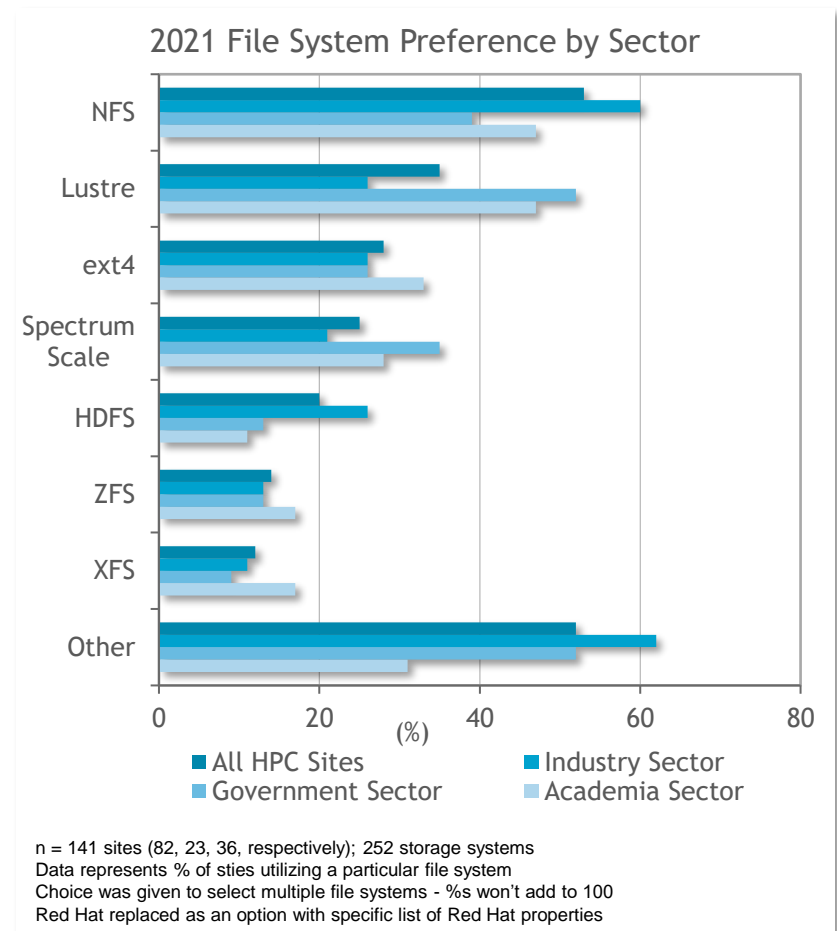
# Storage Ecosystem Technology

# File System Preferences

*Sharp increase in utilization of emerging file systems  
Industry sector leading exploration of “other” file systems*



Source: Hyperion Research, 2021



Source: Hyperion Research, 2021

# File System Market Observations

*Emerging file systems gaining traction*

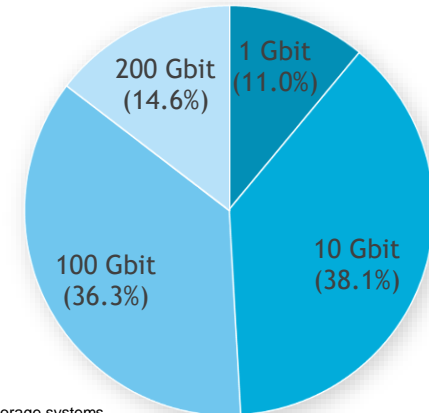
- **Sites average 1.8x file systems**
- **NFS and Lustre remain top two**
  - Both exhibited growth
  - Industry leans towards NFS
  - Government tends towards Lustre
- **ext4 rose to third most preferred**
- **Spectrum Scale dropped to fourth**
  - Sharp drop in adoption at government sites
  - Increased adoption within industry
- **Next five (10% - < 20%)**
  - HDFS, ZFS, XFS, BeeGFS, GFS
- **Next tier (2% - < 10%)**
  - CEPH, OneFS, WekaFS, Gluster, PVFS/OrangeFS, DAOS, PanFS

# System Interconnect Preferences

*InfiniBand edges ahead of Ethernet in aggregate adoption*

- **Initiated distinction between storage and system interconnects**
- **Ethernet**
  - Adopted at 44% of sites surveyed
  - 10 Gbit most widely adopted
- **InfiniBand**
  - Deployed at 46.1% of sites surveyed
  - EDR 100 Gbit most widely deployed
- **Omni-Path**
  - Some adoption across all sectors
  - 2<sup>nd</sup> largest deployed in academic sites (13.9%)

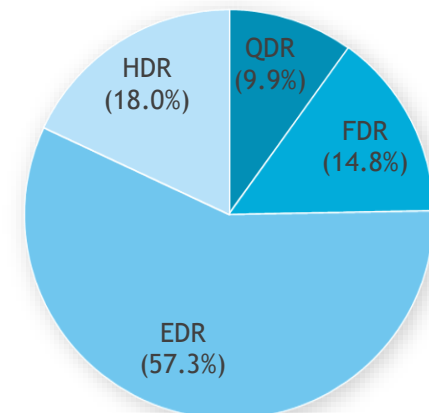
2021 Ethernet Adoption



44%  
Total

n = 141 sites; 252 storage systems  
Source: Hyperion Research, 2021

2021 InfiniBand Adoption

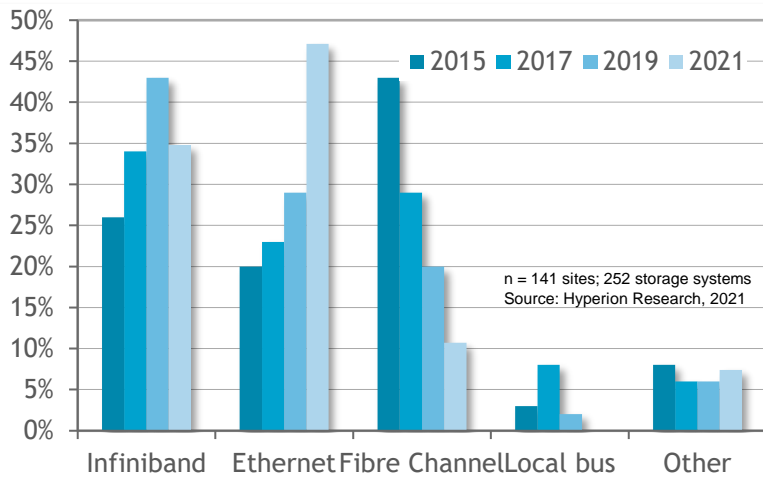


46.1%  
Total

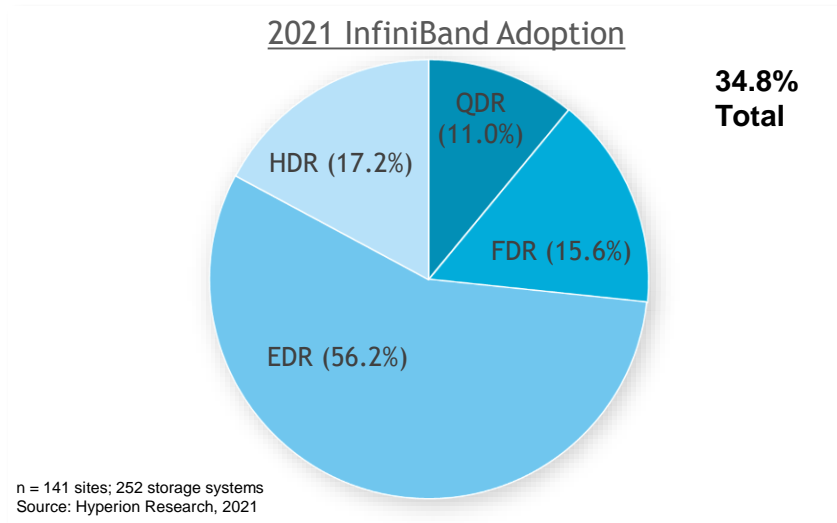

n = 141 sites; 252 storage systems  
Source: Hyperion Research, 2021

# Storage Interconnect Preferences

*Ethernet emerged as preferred storage interconnect*



- **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**



# 2H21 Storage Industry Review

# 2H21 HPC Storage and Interconnect Activity

## *Broad areas of investment and innovation*

- **Pawsey 130PB system refresh**
  - Saving data for longer periods of time
  - More researchers globally accessing the data
  - In addition to Setonix announcement in 1H21
- **IBM Spectrum Fusion**
  - Red Hat bearing fruit
  - Containerized IT for increased adoption of HPC-enabled AI
- **Lustre cloud enhancements**
  - AWS
  - Azure
  - Google Cloud
- **DNA Storage financing**
  - CATALOG
  - Molecular Assemblies
  - DNA Script
- **Interconnect / Networking announcements**
  - Rockport Networks – switchless network
  - Accelerating CXL adoption plans

# 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

**THANK YOU!**

**Want to continue the conversation?**



[info@hyperionres.com](mailto:info@hyperionres.com)

[mnoosokoff@hyperionres.com](mailto:mnoosokoff@hyperionres.com)