White Paper

HPC Storage: The Unsung Hero of HPC Solutions

Sponsored by: Dell Technologies and Intel Corporation

Mark Nossokoff and Alex Norton
October 2020

HYPERION RESEARCH OPINION

The pace of change within the HPC industry is accelerating across all fronts. From a workload perspective, traditional workloads such as seismic processing, life sciences, and weather analysis continue to provide researchers and engineers the necessary data and insights to continue their discoveries and drive evolutionary HPC system-level advancements. Once-emerging workloads such as artificial intelligence (AI) and high performance data analysis (HPDA) have now become mainstream and are delivering unprecedented scientific and business value to researchers and companies in both traditional HPC and commercial enterprise markets while also exponentially increasing the amount of data required to deliver that value.

From a consumption perspective, the cloud has developed into a viable alternative environment for a number HPC sites, especially those that require resources beyond what their current on-premise datacenter has to offer or those that for business necessities require rapid spin up or shut down of HPC infrastructure. The growing adoption of hybrid and multi-cloud solutions creates new opportunities for a wide section of HPC resource alternatives for users, including HPC storage. That said, on-premises HPC infrastructure is not going away.

From a technology perspective, between the growing number of compute options (CPUs, GPUs, FPGAs) and emphasis on delivering exaFLOPs-class performance, computational capabilities receive the lion's share of attention. Networking and interconnects also receive their fair share of the limelight. Storage, however, operates in relative obscurity. Traditionally viewed by some as a necessary evil, storage is the linchpin and common denominator for all the above. Without reliable, performant 24/7 access to secure and trusted data whenever, wherever, and however it's required, the scientific discoveries and business value for HPC/AI/HPDA solutions would not be possible.

Secure, reliable, and performant storage doesn't just happen. Best-in-class storage solutions require a deep understanding of a wide range of parameters including I/O profiles, workloads, use cases, data types and datacenter types.

Given their #1 market position in both worldwide servers and worldwide data storage, according to IDC, Dell Technologies is well positioned to address the challenging needs of HPC storage, benefiting from both projected robust growth in the traditional HPC market and from the even-faster predicted growth from the commercial enterprise market's adoption of HPC/AI/HPDA techniques and solutions.

SITUATION OVERVIEW

The HPC storage market is multi-faceted. Market dynamics influence budget and investment prioritization areas. Traditional and emerging workloads drive new product requirements. Technology evolution and integration provide the means with which HPC storage is incorporated into overall HPC architectures. Solution delivery, service and support enable new markets to take advantage of what HPC has to offer. Each of these areas are examined in this paper.

Market Perspective

While the storage segment of the broader HPC market has traditionally been second in market size behind the HPC server market, it is the fastest growing HPC market segment. Table 1 summarizes the historical and forecast data for the broader HPC market segments.

TABLE 1

Revenues by the Broader HPC Market Areas ($Millions)

	2019	2020	2021	2022	2023	2024	CAGR 19-24
Server	$13,713	$10,860	$12,313	$14,793	$16,810	$18,262	5.9%
Add-on Storage	$5,427	$4,375	$5,010	$6,097	$7,098	$7,767	7.4%
Middleware	$1,614	$1,286	$1,459	$1,778	$2,034	$2,222	6.6%
Applications	$4,690	$3,724	$4,126	$4,917	$5,492	$5,860	4.6%
Service	$2,239	$1,741	$1,889	$2,213	$2,423	$2,535	2.5%
Total Revenue	$27,683	$21,987	$24,797	$29,798	$33,858	$36,646	5.8%

Source: Hyperion Research, October 2020

Examining the server and storage segments further in Table 2 reveals the increasing impact AI and HPDA workloads are contributing to the overall HPC storage market.

TABLE 2

Post-COVID Worldwide HPC-Based AI Storage Revenues vs Total HPDA Storage Revenues ($ Millions)

	2019	2020	2021	2022	2023	2024	CAGR '19-'24
HPC Add-on Storage Revenues	$5,427	$4,375	$5,010	$6,097	$7,098	$7,767	7.4%
HPDA Add-on Storage Revenues	$1,532	$1,515	$1,966	$2,417	$3,008	$3,551	18.3%
HPC-Based AI Add-on Storage Revenues (ML, DL & Other)	$391	$450	$655	$889	$1,241	$1,730	34.7%

Source: Hyperion Research, 2020

Although the scale and size of an HPC storage system varies between HPC competitive segments, the general makeup and architecture of the storage solution is typically consistent. This widely used design approach allows for different configurations of the same basic system, typically architected as storage building blocks, to be leveraged across the competitive server segments. While the Supercomputer segment accounts for the largest revenue, the Divisional and Departmental segments also drive significant opportunities and should not be dismissed when architecting and designing HPC storage solutions. Table 3 details the server and storage forecast for the HPC competitive segments.

TABLE 3

Worldwide Total Technical Computer Market Add-on Storage Revenue Forecast by Competitive Segment ($ Millions)

Segment	2019	2020	2021	2022	2023	2024	CAGR 19-24
Supercomputer	$2,318	$1,902	$2,389	$2,789	$3,420	$3,846	10.7%
Divisional	$751	$616	$695	$880	$967	$1,048	6.9%
Departmental	$1,660	$1,310	$1,369	$1,774	$2,022	$2,181	5.6%
Workgroup	$698	$547	$557	$654	$690	$691	-0.2%
Totals	$5,427	$4,375	$5,010	$6,097	$7,098	$7,767	7.4%

Source: Hyperion Research, October 2020

HPC/AI/HPDA Workloads and Use Cases

AI and HPDA workloads have been driving storage requirements beyond those of traditional HPC workloads. Traditional HPC storage for conventional modeling and simulation typically consists of project, scratch, persistent and archive workloads. AI workflows present a different set of workloads: ingest, data preparation, training, inference, and archive. Some possess storage attributes like those of traditional HPC workloads and others drive new or more aggressive and extreme aspects.

HPC and AI workloads often exhibit different I/O profiles. Traditional HPC workloads are typically based on large sequential I/O while AI workloads demand a mix of large sequential and small random I/O. Metadata management for AI dataset tagging and labeling requires fast small random I/O.

Use cases also drive a variety of durability and resiliency solution needs. Archiving requires extremely cost-effective solutions without demanding performance requirements. Traditional scratch applications require high performance with the ability to offload interim results to durable storage to protect against failures. AI and HPDA solutions require a mix of storage needs for both high performance, transient storage and durable, resilient storage, including a balanced intermix of large block sequential and small block random I/O profiles.

Lastly, data types drive requirements for different types of storage systems. Structured and unstructured data employ varying degrees of file, block and object access methods.

Table 4 summarizes the relationship between workloads, use cases, I/O profiles and data types.

TABLE 4

Traditional HPC and AI/HPDA Workloads

Workload	Use Case	Description
Traditional HPC	Project	- Sometimes referred to as home directories or user files - Used to capture and share final results of the modelling and simulation - Mixture of bandwidth and throughput needs, utilizing hybrid flash and HDD storage solutions
	Scratch	- Workspace capacity used to perform the modelling and simulation - Includes metadata capacity (high throughput [IOs/sec] and flash-based) and raw data capacity and checkpoint writes for protection against system component failure during long simulation runs (high bandwidth [GB/s], traditionally HDD-based but now largely hybrid flash and HDDs)
	Archive	- Long-term data retention - Scalable storage without a critical latency requirement - Largely near-line HDD-based systems with a growing cloud-based element - Typically file or object data types
AI & HPDA	Ingest	- Quickly loading large amounts of data from a variety of different sources such that the data can be tagged, normalized, stored and swiftly retrieved for subsequent analysis - Requires very high bandwidth (GB/s) performance at scale to sustain retrieving data rates, typically object-based from high-capacity HDD-based storage and increasingly cloud-based

TABLE 4

Traditional HPC and AI/HPDA Workloads

Workload	Use Case	Description
	Data Preparation	- Often times referred to as data classification or data tagging, requires a balanced mix of throughput and bandwidth (hybrid flash and HDD storage systems)
	Training	- Utilizing Machine Learning (ML) and/or Deep Learning (DL) to build an accurate model for researchers, engineers and business analysts to use for their research, design and business needs - Requires high throughput (IOs/sec) and low latency for continuous and repetitive computational analysis of the data, typically flash-based storage
	Inference	- Utilizing the model for experimentation and analysis to derive and deliver the targeted scientific or business insights - Also requires high bandwidth and low latency and typically flash-based, often with a caching layer
	Archive	- Long-term data retention - Scalable storage without a critical latency requirement - Largely near-line HDD-based systems with a growing cloud-based element - Typically file or object data types

Source: Hyperion Research, October 2020

Anatomy of an HPC Storage System

HPC storage generally consists of the elements of an HPC system required to deliver a complete external add-on storage system. These elements include:

- Systems that house the controllers (RAID) and physical devices (HDDs, SSDs) that respectively provide the storage services (replication, snapshots, redundancy) and storage media that manage, store and maintain the data
- Expansion storage enclosures to provide additional storage media that scales out from a storage server
- File systems and servers dedicated to running the file system inclusive of primary storage, metadata storage and archive storage
- Storage interconnect switches and cabling that provide the connectivity between HPC compute servers and storage servers, and between storage servers and enclosures

HPC storage systems had long been HDD-based. Flash-based tiers were introduced to support low-latency needs driven by metadata as part of hybrid HDD/flash systems. Flash storage adoption has been increasing in recent years to the point that all-flash storage systems are beginning to appear for HPC storage. Still, differing workloads drive different performance, scale and budget requirements and will continue to drive demand for all HDD, hybrid HDD/flash and all flash HPC storage solutions.

File systems bridge the gap between the applications consuming the data and the physical devices where the data resides. They distribute the data being generated and analyzed by the HPC applications running on the server across the storage media and manage its layout, performance and

resiliency. There are a variety of file systems that address varying degrees of scale, performance, data services, redundancy, resiliency, maintenance and support.

The server-storage interconnect infrastructure is one of three HPC system interconnects (the others being the compute-memory and server-server) within an HPC system. All three interconnects need to be properly balanced to optimize the performance of the system. There are several server-storage interconnects available, with the most broadly deployed being different generations and variations of Ethernet and InfiniBand.

Complicating the discussion is the crossover occurring within elements of an overall HPC system. On the surface it's increasingly difficult to tell the difference between servers and storage. Product categories such as storage servers and computational storage have emerged to address specific needs such as object-based storage targets and edge computing devices, respectively.

Beyond the Hardware

The line between enterprise and HPC market segments is blurring. Ease of management from the enterprise world is being leveraged into HPC storage systems, particularly in the area of metadata management and proactive diagnostics and maintenance, while HPC-enabled AI is making its way into enterprise systems, particularly in the areas of business intelligence and data analytics.

Delivering the infrastructure that meets the demanding requirements of HPC users is necessary but not sufficient for a vendor to be viewed as a major player in the HPC market. Historically, HPC vendors have targeted datacenter managers and users possessing primarily technical and scientific backgrounds:

- HPC datacenter and system managers:
  - Technical aspects of the solution, e.g. performance (flops, IOPs, bandwidth), connectivity, capacity (PBs), core count, memory size, etc.
  - Power consumption and cooling requirements
  - Support and maintenance costs
- Domain-specific engineers and researchers
  - Time to "science" in respective domain areas (e.g., bioscience, genomics, weather, climate, manufacturing, autonomic driving)
  - Length of time to train an AI model
  - Duration to achieve reliable AI inferencing results

HPC techniques are now being leveraged in AI architectures, and commercial enterprise datacenters are adopting AI to deliver appropriate resources for their business's HPDA and business analytics. To address this growing market, vendors' conversations today must evolve to target an increasingly business-oriented audience, typically with fewer technical resources at their disposal to evaluate and implement the complex infrastructure:

- Enterprise datacenter and system managers:
  - Consumption model choices (outright purchase, lease, pay-as-you-go, cloud)
  - Power consumption and cooling requirements
  - Support and maintenance costs

- Business unit managers:
  - ROI analysis
  - Domain-specific solutions
  - Turnkey solutions targeted at specific domain areas
  - Fully integrated, tested, certified and delivered solutions
- Professional services:
  - Assistance for installation and tuning
  - Options for continuing maintenance

The need for turnkey solutions and professional services for installation, tuning and on-going maintenance and support in the traditional enterprise market cannot be overstated. When compared with the traditional HPC datacenter community, typical enterprise datacenters generally have fewer resources with the technical depth required for their increasingly complex AI infrastructures. Expertise is required not only for the HPC hardware and software but also for the domain-specific knowledge (e.g., weather, bioscience, geoscience, autonomic driving) that can translate the respective domain requirements into solutions that deliver breakthrough science and business results.

One last item to consider is data locality. Data locality presents new and challenging data management needs. Data is often generated from many different sources and needs to be shared globally. Depending on the type of data and the processing that needs to occur with it, it may remain at the edge, be transferred to a central datacenter (either on-prem or in the cloud), or both.

HPC STORAGE AND DELL TECHNOLOGIES

A one-size-fits-all solution does not exist for HPC storage. There are myriad variables to consider when evaluating solutions and vendors for your HPC storage requirements:

- **Workloads:** Are you running traditional HPC Mod/Sim jobs? Is AI your primary focus? What about HPDA? All the above?
- **Use Cases:** How is the data being used? Is there only one use case or will there be multiple use cases?
- **Consumption Model:** Is your infrastructure on-prem? Are you running in the cloud? Both?
- **System requirements:** Do you need a full system, complete with servers, networking and storage? Are you rolling your own system and looking only for the storage element?
- **Breadth of solutions:** Is a complete turnkey solution required, inclusive of integration testing with domain-specific applications? Will you be requiring technical expertise to assist you in your own testing and validation?
- **Installation, service and support:** Do you have the capabilities to perform your own installation service and support? Are you looking for assistance from your system supplier? Do you require on-going professional services?

As success in the enterprise datacenter market is as much about relationships as it is about the products and solutions themselves, and given the company's market leadership position in the enterprise datacenter for both servers and storage, Dell Technologies is well positioned to take advantage of the adoption of HPC and HPC-enabled AI techniques by enterprises. This is in addition to continuing to serve the needs of the traditional scientific HPC community. Augmented by the product, technology, service and support assets gained with the EMC acquisition several years ago,

Dell's broad HPC storage portfolio is a powerful combination of internally developed solutions and partnered offerings. This portfolio covers the spectrum from individual storage elements to certified turnkey systems targeted for specific scientific and business domains, providing a wealth of options for users. These include:

- The PowerEdge product line is a mainstay of the HPC storage portfolio. The recently announced PowerEdge XE7100 storage server is a follow-on to the DSS7000. With 100 high capacity, toolless drives and a one or two dual-socket compute node (with 2nd Gen Intel Scalable Processors), and the ability to support file, block or object-based data, it is versatile enough to be used for either scratch or archive storage. The PowerEdge R740xd can be configured with up to 32 high-performance, low latency 2.5" SSDs, making it extremely suitable for data-intensive AI/ML/DL workloads. Additionally, the R740xd is the basis for the Data Accelerator, an open source, NVMe-based solution focused on supporting the broad HPC community to promote mitigation of I/O-related performance challenges.
- The PowerScale product line (derived from the former EMC Isilon product family), powered by the latest release of its OneFS file system, addresses scalability by efficiently storing, managing, securing, protecting and analyzing unstructured file data. OneFS combines three layers of traditional storage architectures – file system, volume manager and data protection – into a unified software layer, creating a single intelligent file system that spans all nodes within a cluster. OneFS also supports a wide range of data types and diverse workloads with built-in multi-protocol capabilities including NFS, SMB, HDFS, S3, HTTP and FTP protocols, and can store data anywhere – at the edge, in the datacenter or in the cloud.
- Dell EMC PowerVault product line of block-storage-based solutions is targeted as a general-purpose building block with high-capacity and high-performance configurations suitable for project, scratch, and persistent storage.
- Dell EMC ECS Enterprise Object Storage product line is a family of scalable storage solutions suitable for both traditional HPC and data-intensive AI workloads. Built upon PowerEdge servers, ECS is available as either a software-defined storage building block or as a turnkey appliance to support cloud-based infrastructure and long-term data retention.
- Defined as Ready Solutions for HPC Storage consisting of a server, file system, networking and storage, these solutions support multiple file systems (NFS, Lustre, BeeGFS, ArcaStream [GPFS]) to address a wide range of performance, scalability and data management needs, including project and persistent storage.

TABLE 5

Mapping Dell Technologies HPC Solutions to HPC Workloads

Workload	Use Case	Product Line			
		PowerEdge	PowerScale	PowerVault	ECS
Traditional HPC	Project		✓✓✓	✓	✓
	Scratch	✓✓✓	✓✓	✓✓✓	

TABLE 5

Mapping Dell Technologies HPC Solutions to HPC Workloads

Workload	Use Case	Product Line			
		PowerEdge	PowerScale	PowerVault	ECS
	Archive		✓✓		✓✓✓
AI & HPDA	Ingest	✓✓	✓✓	✓	✓✓
	Data Preparation	✓	✓✓✓	✓	
	Training	✓	✓✓✓	✓	
	Inference	✓	✓✓✓	✓	
	Archive		✓✓		✓✓✓

Source: Hyperion Research, October 2020

As products and solutions alone are not enough to address the needs of a complex HPC ecosystem, Dell Technologies offers several additional tools for customers to leverage:

- **Dell Technologies on Demand:** The cloud is increasingly being adopted as an HPC resource incremental to on-premises infrastructure. Dell partners with leading CSPs to offer HPC cloud-based solutions to support customers with cloud-native, hybrid-cloud, and multi-cloud applications.
- **HPC & AI Innovation Lab:** This team of engineers and subject matter experts collaborates with customers and partners to move beyond individual products and develop targeted solutions HPC & AI workloads. The Lab is available directly to evaluate new technology or develop focused solutions for a specific outcome, or virtually via access on-line resources for best practices and benchmark results.
- **Customer Solution Centers:** Resourced with Dell personnel, these centers provide customer and partners free hands-on access to Dell infrastructure and the opportunity to interact directly with Dell for demos and testing before buying. Interaction with the HPC & AI Innovation Lab for advanced solution engineering and performance testing is also available through these centers.
- **HPC & AI Centers of Excellence:** With almost a dozen locations around the world, these third-party centers develop and maintain local partnerships, test new technologies, share best practices and function as entry-points for customers to provide feedback and influence future product roadmaps.
- **Dell HPC Community:** Pre-COVID-19, Dell facilitated several in-person gatherings throughout the year for worldwide community networking and collaboration. Successfully evolving this to an on-line virtual activity, the Dell HPC Community event is a vibrant weekly gathering led by a

combination of industry subject matter experts and Dell HPC experts to provide insight and education across a wide variety of HPC topics, including HPC storage.

Leveraging the tools above as a whole, coupled with extending the highly regarded, world class service and support organization obtained from EMC, will be instrumental for Dell to successfully accelerate the enterprise datacenter's adoption and integration of HPC solutions to address AI workloads and ultimately deliver the resulting business value for their customers.

FUTURE OUTLOOK

HPC storage is a critical element of leading HPC system architectures. Understanding and balancing performance, availability, resiliency, capacity and budgetary requirements will continue to determine the overall technical and business success of an HPC storage system.

Adding to the challenge and complexity of delivering HPC storage solutions is the emergence of AI innovations being adopted by both traditional HPC and commercial enterprise datacenters. As more enterprises realize the value and benefits of HPC-enabled AI techniques, vendors will find an increased TAM (Total Available Market) for their HPC solutions. Enterprise users will rely on vendors to not only provide the physical products and solutions but also HPC and scientific/business domain expertise to allow them to fully exploit the capabilities of HPC-enabled AI technology.

As the global server and storage leader in the traditional enterprise market and the fastest growing HPC server vendor, Dell Technologies is well positioned to serve the HPC storage community. Dell's array of HPC storage-related tools (broad product portfolio, growing suite of tested and certified solutions, HPC and scientific domain expertise, cloud offerings, service and support) merit consideration from users who are deciding on their HPC solutions partner in general, and HPC storage partner in particular.

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
