

Validated Designs for HPC Storage

Unlock the value of your data with high performance storage solutions

Table of Contents

Unlock the value of your data with high performance storage solutions	2
Dell Technologies has what you need	2
What are your HPC storage challenges?	3
Why Validated Designs for HPC Storage	4
BeeGFS Storage	5
NFS Storage	7
PixStor Storage	8
Services and financing	9
Why choose Dell Technologies for analytics, HPC and AI	10
Customer Solution Centers	10
AI Experience Zones	10
HPC & AI Centers of Excellence	10
HPC & AI Innovation Lab	11
Proven results	11
Why wait	11

463 exabytes

of medical data created per day by 2025¹

89% accuracy

predicting location, extent, movement,
and intensity of the rain with AI²

“In total, the UT Research Cyberinfrastructure initiative has supported more than 7,000 researchers and students on more than 2,300 projects, using 543 million hours of computing time. It also has been able to contribute significant scientific and societal impacts, including COVID-19 research, hurricane prediction, wind energy design, and dark energy data.”³

Unlock the value of your data with high performance storage solutions

The data-driven age is dramatically reshaping industries and reinventing the future. As vast amounts of data pour in from increasingly diverse sources, leveraging that data is both critical and transformational. Whether you’re working to save lives, understand the universe, build better machines, neutralize financial risks or anticipate customer sentiment, data informs and drives decisions that impact the success of your organization — and shape the future of our world.

Analytics, high performance computing (HPC) and artificial intelligence (AI) are technologies designed to unlock the value of your data. While they have long been treated as separate technologies, they are converging as the industry comes to understand that AI requires the powerful, scalable compute, networking and storage provided by HPC.

This convergence of advanced computing techniques provides the power to accelerate the pace of discovery, creating more opportunities to break new ground, make important discoveries and solve some of the most important challenges of our time.

Dell Technologies has what you need

Expertise and guidance

The technology around data analytics, HPC and AI is emerging quickly, so your team may not have had time to design, deploy and manage solution stacks optimized for new and emerging workloads. And while AI might seem like the latest IT trend, Dell Technologies has been a leader in the advanced computing space for over a decade, with proven solutions, and expertise. Dell Technologies has a team of HPC and AI experts dedicated to staying on the cutting edge, testing new technologies and tuning solutions to your applications to help you keep pace with this constantly evolving landscape.

Validated Designs for HPC

For many organizations, high performance computing is —or is becoming —an important source of competitive advantage. An optimized HPC storage solution delivers the compute, throughput and capacity needed to manage the rapid data growth and increased workload demands presented by advanced analytics, artificial intelligence and other workloads. Validated Designs are workload-optimized rack-level systems with servers, software, networking, storage and services to scale faster with the confidence of an engineering-tested solution while enabling innovation and discovery without boundaries.

Solutions customized for your environment

Dell Technologies uniquely provides an extensive portfolio of technologies to deliver the advanced computing solutions that underpin successful data visualization and AI implementations. With an extensive portfolio, years of experience and an ecosystem of curated technology and service partners, Dell Technologies provides innovative solutions, workstations, servers, networking, storage and services that reduce complexity so you can capitalize on the promise of HPC and AI.

¹ Harmony Healthcare IT, “[Health Data Volumes Skyrocket, Legacy Data Archives On the Rise](#),” accessed November 2021.

² Lifewire, “[AI Breakthroughs Could Improve Weather Forecasts](#),” October 2021.

³ Austin American-Statesman, “[New University of Texas supercomputer to be one of nation's most powerful](#),” November 2021.

1PB/week

of data shipped by national labs⁴

5.7M searches

every minute on Google⁵

What are your HPC storage challenges?

“It’s complex to plan and deploy HPC storage solutions.”

Designing high-throughput, highly scalable HPC storage systems requires a great deal of expert planning and configuration. In addition, working with multiple vendors to acquire, deploy and support HPC storage systems can be difficult and time consuming.

Dell Technologies Validated Designs for HPC Storage can be delivered with hardware, software and support from Dell Technologies. Each system is based on engineering development and tuning in the [Dell Technologies HPC & AI Innovation Lab](#), so you get HPC storage solutions with detailed performance, sizing characterizations and best practices.

Validated Designs for HPC Storage simplify monitoring and overall management of HPC storage without requiring specialized training or expertise to operate, making storage simpler to maintain. Worry-free HPC deployment and management frees researchers, scientists and engineers to focus on core strategic initiatives instead of managing HPC clusters.

“We need stable and secure storage systems for our critical data sets.”

Validated Designs for HPC Storage take the guesswork out of configuration, reducing interoperability issues and improving quality. The systems are built on Dell EMC PowerEdge servers, PowerSwitch and InfiniBand networking, with PowerVault storage. Dell Technologies engineers and industry experts have worked in collaboration with HPC customers and partners to design these systems in the HPC & AI Innovation Lab. The engineering team then devotes hours to rigorously testing and tuning these system for your workloads. The result is storage that’s highly available, with no single point of failure.

“It’s tough to scale capacity and performance to meet user demands — and stay within budget.”

Validated Designs for HPC Storage allow you to meet your capacity and performance needs within budget. First, Dell EMC servers, networking and storage decrease the cost to store and process large HPC data sets. Then, tested and tuned solutions lower the cost of planning and deploying HPC storage. NFS storage is available with virtually all Linux[®] distributions — lowering acquisition costs. PixStor™ consolidates storage and lowers the cost of data retention through cognitive and policy-driven automation. BeeGFS[®] performs well and scales easily right out of the box.

Cracking the HPC storage problem

[The University of Cambridge Research Computing Service](#) is leveraging the [Data Accelerator](#) (DAC) and the Distributed Name Space (DNE) feature in the Lustre file system to optimize the Cumulus cluster for [top I/O performance](#). This optimization work has led to a huge leap forward in storage performance, according to Dr. Paul Calleja, the University’s Director of Research Computing Services.

⁴ Next Platform, “[How National Labs Move Over a Petabyte per Week](#),” June 2021.

⁵ Domo, “[Data Never Sleeps 9.0](#),” accessed November 2021.

Why Validated Designs for HPC Storage

Validated Designs for HPC Storage are delivered in engineering-tested configurations with hardware, software and collaborative support. HPC & AI Innovation Lab engineers develop and tune each solution based on performance characterizations and best practices to simplify installation and provide faster time to results.

Simplified

Reliable

Efficient

“With DNE, the IOPS performance of this solution is amazing... Now we have stable, repeatable and very high performance runs with no error and determinant behaviour, so I think we have cracked the HPC storage problem.”⁶

—Dr. Paul Calleja, Director
Research Computing
Services, University
of Cambridge

Simplicity

Validated Designs for HPC Storage simplify monitoring and management without requiring specialized training or expertise to operate, making storage simpler to maintain. Worry-free HPC deployment and management frees up the team to focus on core business and strategic initiatives.

Reliability

Validated Designs for HPC Storage help take the guesswork out of configuration, reducing interoperability issues and improving quality. The systems are built on Dell EMC servers and storage to provide redundancy. The result is storage that’s highly available, with no single point of failure. With ProSupport Plus, Dell Technologies Services can provide single-point-of-contact support for both the hardware and software.

Efficiency

Tested and tuned solutions lower the cost of planning and deploying HPC storage, while OpenManage lets you manage your data center hardware from anywhere, at any time. Linux NFS is widely available with Linux distributions. PixStor consolidates storage and lowers the cost of data retention through cognitive and policy-driven automation. BeeGFS streamlines performance right out of the box.

<p>Within Validated Designs for HPC Storage, engineering-tested designs include: Linux NFS, PixStor and/or BeeGFS, all created to speed deployment of HPC storage systems with confidence while saving resources.</p> <p>While all of these solutions include Dell EMC servers, networking, storage and services, the specific configurations are different — optimized for the software. There are key differences in the software such as NFS comes with Linux. PixStor focuses on enterprise software-defined features, while BeeGFS streamlines and simplifies for efficiency.</p>		
HPC BeeGFS Storage	HPC NFS Storage	HPC PixStor Storage
<p>For I/O intensive workloads, BeeGFS is easy to set up and performs well right out of the box, making different types of storage devices available within the same namespace.</p>	<p>For those who want to leverage the network file system that comes with Linux, NFS is a reliable, low-cost, easy to administer solution with good performance.</p>	<p>For massive unstructured data growth, PixStor enables scale in capacity and performance, while matching data value to the capabilities and cost of different storage tiers and types.</p>

⁶ Dell Technologies Innovation Exchange, “The Data Accelerator,” November 2019.

BeeGFS Storage

High performance and ease of use for I/O-intensive workloads

If I/O-intensive workloads are slowing your HPC performance, BeeGFS is the solution. BeeGFS is an open-source, parallel-cluster file system, designed specifically to manage I/O-intensive workloads in performance-critical environments.

Choose BeeGFS for its ease of installation, massive scalability, robustness and exceptional flexibility, including converged systems where servers are used for storage and compute. BeeGFS transparently spreads user data across multiple servers. File system performance and capacity can be scaled out to the desired level by simply increasing the number of servers and disks in the system.

BeeGFS high-performance storage solution configuration	
Servers	
Servers	Management Server (Mgmt.): 1x PowerEdge R650 Metadata and Storage Servers (MDS and SS): 6x PowerEdge R750 (8 drives on one server are dedicated for metadata)
Processor	Mgmt: 2x Intel Xeon Gold 6330 per server, 28 cores MDS and SS: 2x Intel Xeon Gold 6348, 28 cores
Memory	Mgmt: 12x 8GB 3200MT/s RDIMMs (128GB) MDS and SS: 16x 16GB 3200MT/s RDIMMs per server (256GB)
Local disks and RAID controller	Mgmt: PERC H745 Integrated RAID, 2x 480GB 2.5" SAS SSDs configured in RAID 1 MDS and SS: Software RAID, 16x Intel 1.6TB, NVMe Mixed Use AG drive U.2 Gen4 with carrier, BOSS card with 2x 240GB M.2 SATA SSDs in RAID1 for OS
InfiniBand HCA (Slot 1 and 8)	2x NVIDIA® ConnectX®-6 HDR 200Gb/s InfiniBand® HCAs per MDS and SS
Networking	
InfiniBand switch	1x NVIDIA QM8790 HDR
Management switch	1x PowerSwitch S3048-ON
Software	
Operating system	Red Hat® Enterprise Linux
Storage software	BeeGFS v7.2.3
Systems management	iDRAC9 Enterprise
Services	
Consulting, education, deployment, support, remote management, cloud options, financing	

“The Commonwealth Scientific and Industrial Research Organisation (CSIRO) has adopted BeeGFS file system for their 2PB all-NVMe storage in Australia, making it one of the largest NVMe storage systems in the world.”⁷

⁷ PacificTeck case study, [CSIRO](#), accessed November 2021.



BeeGFS high capacity storage solution specifications	
Server configuration	
Servers	Management Server (Mgmt.): 1x PowerEdge R650 Metadata Server (MDS): 2x PowerEdge R750 Storage Server (SS): 2x PowerEdge R750
Processor	Mgmt: 2x Intel Xeon Gold 5218, 16 cores MDS and SS: 2x Intel Xeon Gold 6230, 20 cores
Memory	Mgmt: 12x 8GB 2666MT/s RDIMMs MDS and SS: 12x 32GB 2933MT/s RDIMMs per server
Local disks and RAID controller	Mgmt: PERC H740P Integrated RAID, 8GB NV cache 6x 300GB 15K SAS hard drives (HDDs) configured in RAID10 MDS and SS: PERC H330+ Integrated RAID, 2x 300GB 15K SAS HDDs configured in RAID1 for OS
InfiniBand HCA (Slot 8)	1x NVIDIA ConnectX-6 per MDS and SS
Networking	
InfiniBand switch	1x NVIDIA QM8790 HDR
Management switch	1x PowerSwitch S3048-ON
Storage	
Metadata storage configuration	
Storage enclosure	1x PowerVault ME4024 with dual SAS controllers
Hard disk drives	24x 960GB SAS3 SSDs, 12x RAID1 disk groups of 2 drives
Data storage configuration	
Storage enclosure	1, 2, or 4x PowerVault ME4084 with dual SAS controllers
Hard disk drives	84, 168, or 336x 3.5" 7.2K RPM near line (NL) SAS with 4, 8, 10 or 12TB drives Each array has 84 HDDs configured in 8x RAID6 disk groups of 10 drives (8+2) with 4 global spares Max usable capacity is 2766TiB with 336 x 12TB drive and raw capacity is 3667TiB (4032TB)
External storage controllers	SS: 4x SAS 12Gbps HBA for Large configuration, 2x for Small and Medium configurations MDS: 2x SAS 12Gbps HBA
Software	
Operating system	Red Hat Enterprise Linux
Storage software	BeeGFS
Systems management	iDRAC9 Enterprise
Services	
Consulting, education, deployment, support, remote management, cloud options, financing	



NFS Storage

Low cost, high-capacity, scalable HPC storage

Storage solutions based on the NFS protocol are widely used for HPC clusters because NFS is simple and time-tested, and is a standard package in virtually every Linux distribution. If you have clusters running applications with lower I/O needs, NFS storage is reliable, easy to administer and has very good performance within certain boundaries. In clusters with higher I/O requirements, NFS is a good option for a secondary storage repository for home directories, application storage and longer-term storage of application data.

Specifications of a single building block	
Servers	
Server	2x PowerEdge R750
Processor	2x Intel Xeon® Gold 6346, 3.1GHz, 16 cores/processor
Memory	256 with 16x 16GiB 3200MT/s RDIMM per server
Local disks and RAID controller	PowerEdge RAID (PERC) H345 front-embedded with 5x 480GB SAS3 hard drives
Storage controller	2x HBA355e adapters per server
Networking	
	<ul style="list-style-type: none"> • PowerSwitch 10Gb Ethernet management • NVIDIA ConnectX-6 VPI Adapter • NVIDIA QMR8700 HDR100
Storage	
Storage system	1x PowerVault ME4084 Up to 1024TiB of supported usable storage capacity, 50% more than the Red Hat standard supported limit 84x 4, 8, 10, 12 or 18TB NL SAS hard drives
Software	
Operating system	Red Hat Enterprise Linux (RHEL) 8.3
Cluster suite	Red Hat Cluster Suite from RHEL 8.3
File system	Red Hat Scalable File System (XFS) 5.0.0
Systems management	Integrated Dell Remote Access Controller 9 (iDRAC9) OpenManage 10.0.1
Services	
Consulting, education, deployment, support, remote management, cloud options, financing	

Pair **Validated Designs for HPC Storage** with these and other solutions.

[Validated Designs for HPC Life Sciences](#)

Deliver high throughput and fast turnaround for a diverse range of life sciences fields, including drug design, cancer research, agriculture, biofuels and forensics.

[Validated Designs for HPC Digital Manufacturing](#)

A flexible building-block approach for computer aided engineering, modeling and simulation, including structural analysis and computational fluid dynamics.

[Validated Designs for HPC Research](#)

Solutions that enable research centers to quickly develop HPC systems that match the unique needs of a wide variety of workloads involving scientific analysis.



PixStor Storage

High performance, scalable parallel file system with data tiering and simplified management

Delivering data-driven insights requires storage that can handle massive unstructured data growth with security, reliability and high performance. PixStor is a high-performance shared-disk file management software that provides fast, reliable access to data from multiple servers. It can share data using multiple protocols including NFS, SMB, S3 and Rest. It enables seamless storage scaling, advanced search and analytics, tiering and unified management through a single storage namespace.

Specifications	
Servers	
Servers	Management: 1x PowerEdge R650, 2x for Large configuration Gateway/Ngenea: 1x PowerEdge R750, or more as needed Optional high-demand metadata server node: (MD): 2x PowerEdge R750 for Large configuration Storage node (SN): 2x PowerEdge R750
Processor	Management: 2x Intel Xeon Gold 6336Y Gateway/Ngenea, MD and SN: 2x Intel Xeon Gold 6336Y
Memory	Management: 12x 16GB DDR4 3200MT/s DIMMs 256GB Gateway/Ngenea MD and SN: 16x 16GB DDR4 3200MT/s DIMMs 256GB
Local disks	3x 480GB SSDs on RAID1 for OS + hot spare Management only: 5x 480GB SSDs on RAID5 for analysis and monitoring
InfiniBand HCA (Slot 8)	2x NVIDIA ConnectX-6 HDR VPI per MD and SN 3x with optional Gateway
External storage HBA adapters	2-4x SAS HBA355e on each MD server 2-4x SAS HBA355e on each SN 4x for Large, 2x for Small and Medium standard configurations
Networking	
InfiniBand switch	2x NVIDIA QM8700 HDR
Management switch	PowerSwitch S3048-ON
Storage	
Optional high-demand metadata storage configuration	
Storage enclosure	Up to 4x PowerVault ME4024 with 2x or 4x SAS HBA355e controllers
Hard disk drives	24x 960GB SAS3 SSDs per ME4024 enclosure, 12x RAID1 LUNs (options 480GB, 960GB, 1.92TB, 3.84TB)
Object storage configuration	
Storage enclosure	1, 2, or 4x PowerVault ME4084 with 2x SAS HBA355e controllers
Hard disk drives	80x 12TB 3.5" NL SAS3 HDD drives in 8x RAID6 LUNs Options: 900GB @15K, 1.2TB @10K, 1.8TB @10K, 2.4TB @10K, 4TB NLS, 8TB NLS, 10TB NLS, 12TB NLS, 18TB NLS 4x 1.92TB or 3.84TB SAS3 SSDs for metadata in standard configuration 2x RAID1, or 4x hot spares for optional high-demand metadata configuration
Software	
Operating system	Red Hat Enterprise Linux
File system	PixStor 6
Systems management	iDRAC9 Enterprise OpenManage
Services	
Consulting, education, deployment, support, remote management, cloud options, financing	

43 total STAR Awards from the Technology Services Industry Association (TSIA) for commitment to outstanding innovation, leadership and excellence in the technology services industry.

2021 Star Awards for Best Practices in Service Offer Development and innovation in support Employee Engagement Experience & Loyalty.⁸

Services and financing

Dell Technologies partners with you every step of the way, linking people, processes and technology to accelerate innovation and enable optimal business outcomes.

- [Services for High Performance Computing](#) are delivered by certified experts to help you get the business value of advanced computing. The services include assessment, workshop, testing, proofs of concept and production implementation. These experts help determine where advanced computing is a good fit for your organization. They also help you build your own internal team of experts through knowledge transfer.
- [Deployment Services](#) help you streamline complexity and bring new IT investments online as quickly as possible. Leverage our 30-plus years of experience for efficient and reliable solution deployment to accelerate adoption and return on investment (ROI) while freeing IT staff for more strategic work.
- [Support Services](#) driven by AI and deep learning will change the way you think about support with smart, ground-breaking technology backed by experts to help you maximize productivity, uptime and convenience. Experience more than fast problem resolution—our AI engine proactively detects and prevents issues before they impact performance.
- [Payment Solutions](#) from Dell Financial Services help you maximize your IT budget and get the technology you need today. Our portfolio includes traditional leasing and financing options, as well as advanced flexible consumption products.
- [Dell Technologies APEX](#) offers a simple approach that gives you a wide range of consumption models, payment solutions and services so you can optimize for a variety of factors while realizing more predictable outcomes.
- [Managed Services](#) can help reduce the cost, complexity and risk of managing IT so you can focus your resources on digital innovation and transformation while our experts help optimize your IT operations and investment.
- [Residency Services](#) provide the expertise needed to drive effective IT transformation and keep IT infrastructure running at its peak. Resident experts work tirelessly to address challenges and requirements, with the ability to adjust as priorities shift.

“It really is a partnership. It’s not always even Dell equipment. Dell has the ability to work holistically, to take a big picture engineering approach. It’s not just about the hardware. They work to identify the right type of resources, connections and services that we will need. But most importantly, they are a partner who helps us think through problems, and find ideal solutions. In today’s world, no one person is an island. You need a team to get things done.”⁹

—Curtis A. Carver Jr., Ph.D., UAB Vice President
and Chief Information Officer

⁸ “[Dell Technologies Key Facts](#),” September 2021.

⁹ Dell Technologies case study, “[Teamed for success](#),” May 2021.

“We needed more than a technology provider, we needed a partner. And that’s we got, unlike any other. We’ve collaborated, strategized and innovated with thought leaders across the organization, and validated solutions to ensure compatibility with our environments.”

—VP, Global Fortune
500 Company

Why choose Dell Technologies for HPC and AI

We’re committed to advancing, democratizing and optimizing data analytics, HPC and AI.

- Schedule an [executive briefing](#) and collaborate on ways to reach your goals.
- [Dell Technologies Customer Solution Centers](#) are staffed with computer scientists, engineers and subject matter experts in a variety of disciplines to help you find the solution that's right for you.
- We are committed to [providing customers with choice](#). We want you to get what you need and have a great experience working with us. If we don’t have what you need, we’ll tell you who does. We believe in being open, and we publish our performance results.
- Dell Technologies is the only company in the world with a portfolio that spans from workstations to supercomputers, including solutions, servers, networking, storage, software and services.
- Because Dell Technologies offers such a wide selection of solutions, the focus remains on finding the right solution for your needs or problems. That range of solutions also brings the expertise to understand a broad spectrum of challenges and how to address them.

Customer Solution Centers

Our global network of dedicated Dell Technologies [Customer Solution Centers](#) are trusted environments where world-class IT experts collaborate with you to share best practices, facilitate in-depth discussions of effective business strategies, and help your organization become more successful and more competitive. Customer Solution Centers reduce the risks associated with new technology investments and can help improve speed of implementation.

AI Experience Zones

Curious about AI and what it can do? AI Experience Zones inside [Dell Technologies Customer Solution Centers](#) can run demos, try proofs of concept and pilot software. Dell EMC experts are available to collaborate and share best practices as your customer can explore the latest technology and get the information and hands-on experience they need for their advanced computing workloads.

HPC & AI Centers of Excellence

As data analytics, HPC and AI converge and the technology evolves, Dell Technology’s worldwide HPC & AI Centers of Excellence provide thought leadership, test new technologies and share best practices. They maintain local industry partnerships; and have direct access to Dell Technologies and other technology creators to incorporate your feedback and needs into their roadmaps. Through collaboration, [HPC & AI Centers of Excellence](#) provide a network of resources based on the wide-ranging know-how and experience in the community.

“The HPC & AI Innovation Lab gives our customers access to cutting-edge technology, like the latest-generation Dell EMC products... Customers can bring us their workloads, and we can help them tune a solution before the technology is readily available.”

— Garima Kochhar,
Distinguished Engineer

HPC & AI Innovation Lab

The [Dell Technologies HPC & AI Innovation Lab](#) in Austin, Texas, is the flagship innovation center. Housed in a 13,000-square-foot data center, it gives you access to thousands of servers, three powerful supercomputers, and sophisticated storage and network systems. It's staffed by a dedicated group of computer scientists, engineers and subject matter experts who actively partner and collaborate with customers and other members of the HPC community. The team engineers HPC and AI solutions, tests new and emerging technologies, and shares expertise including performance results and best practices.

Proven results

Dell Technologies holds leadership positions in some of the biggest and largest-growth categories in the IT infrastructure business, and that means you can confidently source your information technology needs from Dell.

- #1 in servers¹⁰
- #1 in converged and hyper converged infrastructure (HCI)¹¹
- #1 in storage¹²
- #1 cloud IT infrastructure¹³

See [Dell Technologies Key Facts](#).

Why wait

Learn more today about how you can quickly deploy a high performance storage solution ready to support advanced computing, machine and deep learning initiatives. Contact your Dell Technologies or authorized partner sales representative, join the HPC Community at [dellhpc.org](#), and visit [delltechnologies.com/hpc](#) to learn more.

¹⁰ IDC, WW Quarterly x86 Server Tracker, [2Q2021](#), Vendor Revenue & Shipments, September 9, 2021.

¹¹ IDC, WW Quarterly Converged Systems Tracker, [4Q2020](#), Vendor Revenue, March 2021.

¹² IDC, WW Quarterly Enterprise Storage Systems Tracker, [2Q2021](#), September 9, 2021.

¹³ IDC, [WW Quarterly Enterprise Infrastructure Tracker: Buyer and Cloud Deployment](#), [2Q2021](#), Vendor Revenue, October 1, 2021.

Contact us

To learn more, visit [delltechnologies.com/hpc](#) or [contact](#) your local representative or authorized reseller.

