

Dell EMC Ready Solutions for HPC Storage

Unlock the value of your data with artificial intelligence and high performance computing

Table of Contents

Unlock the value of your data.	2
Dell EMC has what you need	2
What are your HPC storage challenges?	3
Cracking the HPC storage problem	3
Why Dell EMC Ready Solutions for HPC Storage?	4
Ready Solutions for HPC NFS Storage	5
Ready Solutions for HPC Lustre Storage	6
Services and financing	7
Why choose Dell EMC for data analytics, HPC and AI.	8
Dell Customer Solution Centers	8
Dell EMC HPC and AI Centers of Excellence	8
Dell EMC AI Experience Zones	8
Dell EMC HPC and AI Innovation Lab	9
Proven results	9
Why wait?	9

2,314 exabytes

of medical data predicted by 2020¹

99% accuracy

identifying tropical cyclones, weather fronts and atmospheric rivers using AI and HPC²

“We laid out our requirements and the people at Dell EMC took those requirements and developed exactly what we needed.”³

—Thomas McCauley,
Engineering Manager,
Caterpillar Autonomous
Mining Program

Unlock the value of your data with artificial intelligence and high performance storage

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.

Data analytics, HPC and AI are technologies designed to unlock the value of your data. While they have long been treated as separate technologies, the three are converging as the industry comes to understand that analytics and AI are both essentially big-data problems that require 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 EMC 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 develop the skills required to design, deploy and manage solution stacks optimized for new workloads. While AI might seem like the latest IT trend, Dell EMC has been a leader in the advanced computing space for over a decade, with proven products, solutions and expertise. Dell EMC has a team of data analytics, 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.

Dell EMC Ready Solutions for HPC

The advantage in today's marketplace goes to the data driven enterprise. For many organizations, high performance computing is —or is becoming —an important source of competitive advantage. An optimized HPC solution delivers the compute, throughput and capacity needed to manage the rapid data growth and increased workload demands presented by advanced data analytics and other enterprise workloads. Dell EMC Ready Solutions for HPC simplify design, configuration and ordering of clusters with standardized building blocks that are tested for high performance computing applications.

Solutions customized for your environment

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

¹ Dell EMC ebook, “[Making digital transformation in healthcare a reality](#),” February 2018.

² Earth Institute at Columbia University, “[Artificial Intelligence—A Game Changer for Climate Change and the Environment](#),” June 2018.

³ Dell EMC case study, “[Autonomous Mining](#),” August 2017.

>30 billion

requests for global weather information every day⁴

30% higher

crop yields using HPC and AI insights for preparing land, applying fertilizer and timing sowing⁵

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 an HPC storage system can be difficult and time consuming.

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

Dell EMC Ready Solutions 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 business and strategic research initiatives instead of managing HPC clusters.

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

Dell EMC Ready Solutions for HPC Storage take the guesswork out of configuration, reducing interoperability issues and improving service quality. The systems are built on Dell EMC PowerEdge servers and Dell EMC PowerVault storage to provide redundancy. Dell EMC engineers and industry experts have worked in collaboration with Dell EMC HPC customers and partners to design these systems in our Dell EMC HPC and AI Innovation Lab. The Dell EMC engineering team then devotes hours to rigorously testing and tuning the 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.”

Dell EMC is leading the charge to mainstream HPC by bringing costs down so this technology is available for all. Dell EMC Ready Solutions for HPC Storage allow you to meet your capacity and performance needs with compelling total cost of ownership (TCO) benefits. First, cost-optimized, industry-standard Dell EMC servers, storage and networking decrease the cost to store and process large HPC data sets. Then, pre-configured, tested and tuned solutions lower the costs of planning and deploying HPC storage. Cost-effective network file system (NFS) storage is available with virtually all Linux® distributions — lowering acquisition costs. It’s easy to configure, deploy and maintain — lowering operational costs. Lustre® features user interface–driven configuration, monitoring and overall management to lower management complexity and cost.

Cracking the HPC storage problem

[The University of Cambridge Research Computing Service](#) is leveraging the Dell Data Accelerator (DAC) and the Distributed Name Space (DNE) feature in the Lustre file system to optimise the Cumulus cluster for top I/O performance. This optimisation 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 AccuWeather, “[AccuWeather Exceeds Record Milestone in Big Data Demand, Answering More than 30 Billion Requests Daily](#),” October 2017.

⁵ Earth Institute at Columbia University, “[Artificial Intelligence—A Game Changer for Climate Change and the Environment](#),” June 2018.

Why Dell EMC Ready Solutions for HPC Storage?

Dell EMC Ready Solutions for HPC Storage are delivered with hardware, software and support from Dell EMC. HPC and AI Innovation Lab engineers develop and tune each design based on performance characterizations, tuning and best practices to simplify installation and provide faster time to results.

Simplified

Reliable

Cost-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

Dell EMC Ready Solutions 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 instead of managing HPC clusters.

Reliability

Dell EMC Ready Solutions for HPC Storage help take the guesswork out of configuration, reducing interoperability issues and improving service quality. The systems are built on Dell EMC PowerEdge servers and storage to provide redundancy. The result is storage that’s highly available, with no single point of failure. Dell EMC support services can cover both hardware and software, to help prevent and address potential issues with a single source of solution support.

Cost-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. NFS is widely available with Linux distributions. It’s also easy to configure, deploy and maintain — reducing complexity while lowering operational costs.

Choose from two designs, along with the flexibility to customize a solution for your specific needs.	
HPC NFS Storage	HPC Lustre Storage
Low-cost, performant HPC storage for lower I/O needs	High-performance, massively scalable and cost-effective HPC storage
For clusters running applications with lower I/O needs, network file system (NFS) storage offers simplicity, reliability and lower costs.	For extreme scalability without losing performance, Lustre® allows scaling efficiently both up and down to suit workloads without losing performance or capacity.
This solution is for those who need an industry-standard storage solution that is reliable and easy to administer, and has very good performance within certain boundaries.	This solution delivers parallel storage with enterprise-tested performance, capacity and extreme scalability.
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.	The supported, commercial-ready Lustre file system has been rigorously tested in demanding enterprise environments. It is a single, simplified system that handles both big data and HPC workload demands.

⁶ Dell EMC case study, “UK Science Cloud,” November 2018.



Ready Solutions for HPC NFS Storage

Low cost, high-capacity, scalable and performant HPC storage for lower I/O needs
 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	
Server	2x PowerEdge R740 Servers
Processor	Dual Intel® Xeon® Gold 6240, 2.6GHz, 18 cores/processor
Network connectivity	Choice of: <ul style="list-style-type: none"> • Dell EMC 10Gb Ethernet • Dell EMC H-Series based on Intel Omni-Path • Mellanox® ConnectX-5 InfiniBand® EDR
Memory	12x 16GiB 2933 MT/s RDIMM per server
Local disks and RAID controller	PowerEdge RAID (PERC) H730P with 5x 300GB 15K SAS hard drives
External storage controller	2x 12Gbps SAS Dell EMC HBAs per server
Storage system	1x PowerVault ME4084 Up to 768TB of supported usable storage capacity, 50% more than the Red Hat®-imposed limit 84x 4, 8, 10 or 12TB NL SAS hard drives
Systems management	Integrated Dell Remote Access Controller 9 (iDRAC9) OpenManage
Operating system	Red Hat Enterprise Linux (RHEL 7.6) x86_64
Cluster suite	Red Hat Cluster Suite from RHEL 7.6
File system	Red Hat Scalable File System (XFS) v4.5.0-18
Network driver	Mellanox OFED 4.5-1.0.1

Ready Solutions for HPC

Small and medium organizations can achieve profit or discovery goals faster with modular HPC solutions from Dell EMC.

[Dell EMC Ready Solutions for HPC Life Sciences](#)

Infrastructure architected to deliver high throughput and fast turnaround for a diverse range of fields, including drug design, cancer research, agriculture, biofuels and forensics.

[Dell EMC Ready Solutions for HPC Digital Manufacturing](#)

A flexible building-block approach to building HPC systems for design simulations, including structural analysis and computational fluid dynamics.

[Dell EMC Ready Solutions for HPC Research](#)

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



Ready Solutions for HPC Lustre Storage

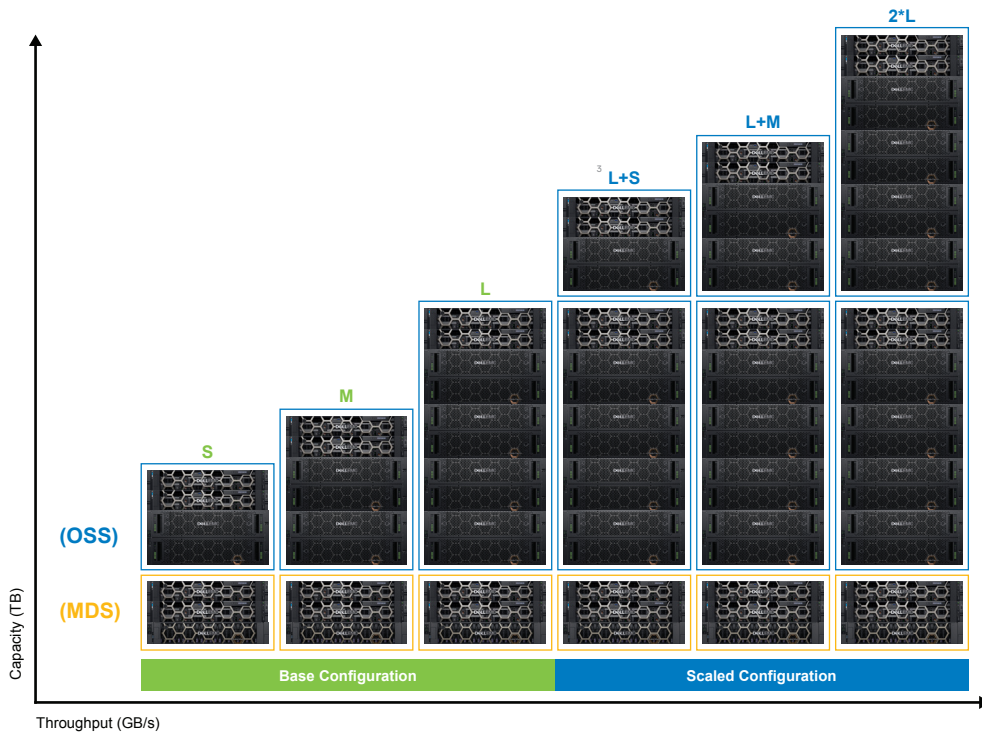
High-performance, massively scalable and cost-effective HPC

Today's HPC workloads require storage infrastructure that scales endlessly and delivers unmatched I/O levels. Dell EMC Ready Solutions for HPC Lustre Storage allow you to scale efficiently both up and down to suit your workloads without losing performance or capacity, so you can tap into the power and scalability of Lustre with simplified installation, configuration and management features — on efficient, scalable Dell EMC systems.

Specifications	
Servers	Management server (IML): 1x PowerEdge R640 Lustre metadata server (MDS): 2x PowerEdge R740 Lustre object storage server (OSS): 2x PowerEdge R740
Processor	IML: 2x Intel Xeon Gold 5218 MDS and OSS: 2x Intel Xeon Gold 6230
Memory	IML: 12x 8GiB 2993 MT/s RDIMMs per server MDS and OSS: 12x 32GiB 2933 MT/s RDIMMs per server
Local disks and RAID controller	IML: PERC H740P Integrated RAID, 8GB NV cache 6x 300GB 15K SAS hard drives (HDDs) configured in RAID10 for object storage (OS) and IML database storage MDS and OSS: PERC H330+ integrated RAID 2x 300GB 15K SAS HDDs configured in RAID1 for OS
Software	Lustre Community Edition — Whamcloud
Systems management	iDRAC9 Enterprise OpenManage
Operating system	CentOS™ 7.6
Network connectivity	Choice of: <ul style="list-style-type: none"> • Dell EMC H-Series based on Intel Omni-Path • Mellanox InfiniBand EDR 10/40GbE
OST storage array	1, 2, or 4x PowerVault ME4084 with 84, 168, or 336x 3.5" 7.2K RPM near-line (NL) SAS with 4, 8, 10 or 12TB drives Max usable Lustre capacity is 2766 TiB with 336x 12TB drives
MDT storage array	1x PowerVault ME4024 with 12x or 24x 960GB SAS SSDs
External storage controllers	Object storage target (OST): 4x SAS 12Gbps Dell EMC HBA MDS: 2x SAS 12Gbps Dell EMC HBA

Ready Solutions for HPC Lustre Storage sizing flexibility

Dell EMC Ready Solutions for HPC Lustre Storage are available in scalable building blocks for 4, 8, 10 and 12TB of estimated usable storage. Each system requires one PowerEdge R640 management server — not represented in the graphic below. Metadata servers (MDS) consist of 2x PowerEdge R740s. Metadata storage consists of 1x PowerVault ME4024s. Object store servers (OSS) consist of 2x PowerEdge R740s. Object storage expands by adding PowerVault ME4084s to create small, medium and large base configurations that can be scaled modularly.



The total rack units (Total U) below represents solution sizes with metadata servers, metadata storage, object store servers and object storage with the required PowerEdge R640 management server, 2x networking switches and 2x PowerVault ME4084 storage.

Total U, # of ME4084	18U, 1	23U, 2	33U, 4	42U, 5	47U, 6	58U, 8
Estimated usable space 4TB/8TB/10TB/12TB¹	231 TiB	461 TiB	922 TiB	1153 TiB	1383 TiB	1844 TiB
(7.2 K RPM NL SAS HDD)	461 TiB	922 TiB	1844 TiB	2305 TiB	2766 TiB	3688 TiB
	576 TiB	1152 TiB	2305 TiB	2881 TiB	3458 TiB	4610 TiB
	691 TiB	1383 TiB	2766 TiB	3458 TiB	4149 TiB	5532 TiB
Peak read performance⁴	≈ 5.6 GB/s	≈ 11.3 GB/s	22.56 GB/s	≈ 28.2 GB/s	≈ 33.8 GB/s	≈ 45.1 GB/s
Peak write performance⁴	≈ 5.3 GB/s	≈ 10.6 GB/s	21.27 GB/s	≈ 26.6 GB/s	≈ 31.9 GB/s	≈ 42.5 GB/s
Sustained performance^{2,4}	≈ 5 GB/s	≈ 10 GB/s	≈ 20 GB/s	≈ 25 GB/s	≈ 30 GB/s	≈ 40 GB/s

¹ Estimated Lustre usable space in TiB ≈ 0.99 * #Arrays * 80 * 0.8 * HDD size in TB * 10¹²/2⁴⁰

² Sustained performance (steady state performance over a longer period of time/thread counts after the peak is attained) of this solution for read as well as write is very similar

³ The L+S configuration shows the maximum performance and density per 42U rack, depending on datacenter power and weight restrictions

⁴ Performance for L configuration is measured. The performance numbers for rest of the configurations are an estimation/extrapolation based on L configuration

“Our turn-key system removes the complexity from the installation, management and use of artificial intelligence frameworks, and has enabled CSIRO to speed up its time to market for scientific outcomes, which will in turn boost Australia’s competitiveness in the global economy.”⁸

—Andrew Underwood,
Dell EMC Field CTO

Services and financing

Dell EMC is there every step of the way, linking people, processes and technology to accelerate innovation and enable optimal business outcomes.

- [Dell EMC Big Data Vision Workshop](#) focuses on big data for business leaders. We have a unique methodology to identify and prioritize a single use case with a combination of implementation feasibility and business value. It’s a three week engagement that applies research, interviews and data science expertise and techniques to the organization — culminating in a one-day workshop for your team to identify and agree on a use case and path forward. This approach sets Dell EMC apart from the “bring in a bunch of technology and see what it can do” approach that’s pushed by many vendors.
- [Dell EMC Consulting Services](#) are delivered by certified experts to help you get the business value of advanced computing. The services include an 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 at each step.
- [Dell EMC Education Services](#) offers courses and certifications including data science and advanced analytics through self-paced online labs and instructor-led workshops.
- [Dell EMC Deployment](#) experts have the experience, expertise and best practices to enhance your success with data analytics, HPC and AI solutions. With a proven track record of success in thousands of engagements worldwide, you can rely on Dell EMC as your partner.
- [Dell EMC Support](#) experts can provide comprehensive hardware and collaborative software support 24x7 for optimal system performance and minimized downtime. ProSupport includes next-business-day on-site service with four- and eight-hour parts-and-labor response options, and escalation management with customer-defined severity levels. You can also opt for ProSupport Plus to get a technology service manager, who serves as a single point of contact for your support needs.
- Once the HPC cluster is deployed, [Dell EMC Remote HPC Cluster Management](#) services help keep it running smoothly with proactive monitoring and management of the entire HPC solution.
- [Dell Financial Services](#) offers a wealth of leasing and financing options to help you find opportunities when your organization faces decisions regarding capital expenditures, operating expenditures and cash flow.

⁸ CSIRO press release, “[CSIRO powers bionic vision research with new Dell EMC PowerEdge based artificial intelligence capability](#),” accessed March 2019.

Winner of the coveted HPCwire Editor's Choice Award for Best Use of High Performance Data Analytics⁹

“The goal is always to push the boundaries of knowledge and inspire the next generation of scientists. With this new system from Dell EMC, I believe we can do that.”¹⁰

—Jarrod Hurley,
Professor of
Astrophysics, Swinburne
University of Technology

Why choose Dell EMC for data analytics, HPC and AI

We're committed to advancing data analytics, HPC and AI, and we've dedicated a great deal of resources toward that goal.

- Come in for an [executive briefing](#) and collaborate on ways to reach your business goals.
- [Dell Customer Solution Centers](#) are staffed with computer scientists, engineers and P.h.Ds who are subject matter experts in a variety of disciplines.
- We are committed to [providing you 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 EMC is the only company in the world with a portfolio that spans from workstations to supercomputers, including servers, networking, storage, software and services.
- Because Dell EMC offers such a wide selection of solutions, we can act as your trusted advisor without trying to sell you a one-size-fits-all approach to your problem. That range of solutions has also given us the expertise to understand a broad spectrum of challenges and how to address them.
- To protect, detect and recover from cyberattacks, [security is built into the Dell EMC PowerEdge server design](#), not bolted on after the fact.

Dell Customer Solution Centers

Our global network of 21 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 business become more successful and competitive. Dell Customer Solution Centers reduce the risks associated with new technology investments and can help improve speed of implementation.

Dell EMC HPC and AI Centers of Excellence

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

Dell EMC AI Experience Zones

Curious about AI and what it can do for your business? Run demos, try proofs of concept and pilot software in Singapore, Seoul, Sydney, and Bangalore. Dell EMC experts are available to collaborate and share best practices as you can explore the latest technology, get the information and hands-on experience you need for your advanced computing workloads.

⁹ HPCwire, “[2018 HPCwire Awards – Readers' & Editors' Choice](#),” November 2018.

¹⁰ Dell EMC case study, “[Inspiring A New Generation of Scientists with Large-Scale Compute](#),” October 2018.

“The HPC and 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,
Dell EMC Distinguished
Engineer

¹¹ IDC [WW Quarterly Server Tracker](#), Vendor Revenue, March 2019.

¹² IDC [WW Quarterly Converged Systems Tracker](#), Vendor Revenue, September 2018.

¹³ IDC [WW Quarterly Enterprise Storage Systems Tracker](#), Vendor Revenue, March 2019.

¹⁴ IDC [WW Quarterly Cloud IT Infrastructure Tracker](#), Vendor Revenue, January 2019.

Dell EMC HPC and AI Innovation Lab

The [Dell EMC HPC and 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 Dell EMC servers, two powerful supercomputers, and sophisticated storage and network systems. It's staffed by a dedicated group of computer scientists, engineers and Ph.D. 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 EMC 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 EMC.

- #1 in servers¹¹
- #1 in converged and hyper converged infrastructure (HCI)¹²
- #1 in storage¹³
- #1 cloud IT infrastructure¹⁴

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 EMC or authorized partner sales representative, join the HPC Community at [dellhpc.org](#), and visit [dell EMC.com/hpc](#) to learn more.

Contact us

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

