

Omnia open-source software

Run AI, HPC and data analytics on the same systems, with greater flexibility.

Table of Contents

AI and HPC Infrastructure as Code	2
Omnia use-cases.	3
Why Omnia.	3
Deploy faster.	4
Get instant gratification.	3
Technical specifications	4
Support Matrix.	4
Omnia 2.1 Architecture	5
Omnia Kubernetes Stack.	5
Omnia Slurm Stack.	5
Dell Technologies has what you need	6
Customer Solution Centers	7
AI Experience Zones	7
HPC & AI Innovation Lab	7
HPC & AI Centers of Excellence	7
Proven results	7
Take the next step, today.	8

AI and HPC Infrastructure-as-Code

The rapid convergence of artificial intelligence (AI) and high-performance computing (HPC) is accelerating the evolution of advanced computational infrastructures as organizations push toward exascale-class performance. Exascale systems—capable of exceeding one quintillion operations per second—are reshaping scientific simulation, engineering workflows, and large-scale AI model training by leveraging heterogeneous architectures and extreme parallelism.

To support these demanding workloads, modern HPC and AI environments are increasingly adopting rack-scale architectures that integrate compute, storage, and high-bandwidth networking within tightly coupled rack systems. These architectures are designed to accommodate rising compute density and the performance needs of GPU-accelerated nodes used in both simulation and AI model training. They also support efficient scaling, predictable performance, and streamlined deployment across large cluster installations.

This shift toward exascale-oriented and rack-scale environments increases the complexity that systems administration teams must manage. Workload diversity - from physics-based simulations to large-scale deep learning - drives constant change in cluster configuration. System administrators are frequently asked to reconfigure and redeploy subclusters in a repeatable fashion - quickly. Omnia's infrastructure-as-code approach allows administrations to eliminate repetitive manual steps by orchestrating Slurm® and/or Kubernetes® environments along with the drivers, libraries, and applications required by modern scientific and AI workflows.

As the demand for compute density increases, clusters are incorporating liquid-cooling systems. This brings the additional complexity of coolant distribution units (CDUs), power distribution units (PDUs) and integrated rack controllers. These elements introduce additional considerations for monitoring and operational stability, such as tracking thermal conditions, adjusting cooling flow, and observing power fluctuations.

As an open-source project originating from the Dell Technologies HPC & AI Innovation Lab, Omnia encourages community participation and continuous evolution to meet the needs of next-generation HPC and AI environments.





Omnia use-cases

Omnia lends itself to a number of use cases, including:

- Training large-scale generative AI models such as LLMs and image/video-generation
- AI inference clusters serving multiple users and applications
- HPC clusters for tightly and loosely coupled parallel computation
- General-purpose accelerated clusters using RDMA-enabled fabrics

Why Omnia

- Omnia is an open-source platform for deploying and managing high-performance clusters across HPC, AI, and analytics workloads. For Dell PowerEdge servers running RPM-based Linux® distributions, Omnia installs Kubernetes and/or Slurm, along with additional frameworks and services, enabling diverse workloads to run on a unified infrastructure.
- Continuous development expands Omnia's ability to quickly deploy new infrastructure into resource pools that can be dynamically allocated to different workloads—helping IT deliver the right environment at the right time.

Deploy faster

- Reconfiguring clusters to support both HPC and AI workloads can be slow and labor-intensive.
- Omnia automates cluster deployment, creating a flexible pool of shared resources.
- It applies the appropriate software stack to each server - whether for HPC simulations, neural-network training, or in-memory analytics - reducing deployment time from weeks to hours.

Get instant gratification

- As an open-source project, Omnia can be customized instantly and evolves with community contributions.
- Dell Technologies integrates leading open-source tools and supports a collaborative ecosystem that accelerates innovation and simplifies troubleshooting—backed by enterprise support when needed.



You don't need to become an expert when using Omnia abstracts away the manual steps that can lead to configuration mistakes.

Technical specifications

Omnia is a collection of Ansible playbooks for automating OS provisioning, deployment of open source Kubernetes and Slurm services; installation of hardware drivers, optimization libraries, and machine learning frameworks/platforms; and serving pre-trained AI inferencing models from the internet or running of research or commercial HPC applications Kubernetes. This is accomplished by installing software from a variety of sources, including RHEL/CRB/AppStream, DockerHub® and Helm® repositories, and Source code compilation. Omnia can be downloaded from GitHub® at: <https://github.com/dell/omnia>

Support matrix: Software and Hardware requirements

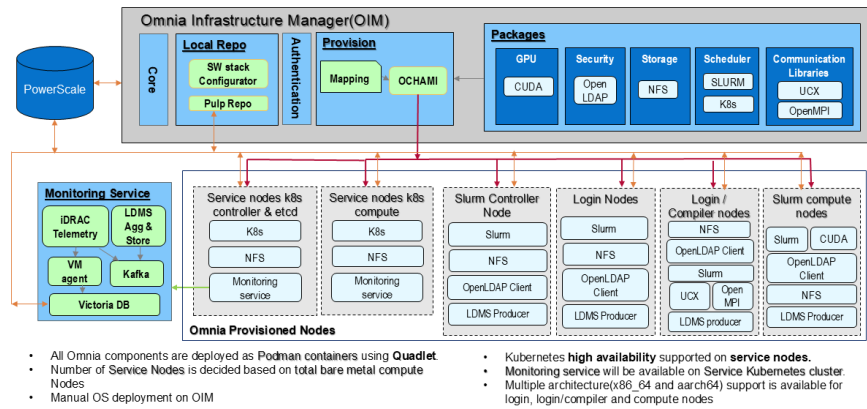
Dell PowerEdge Nvidia Servers 17G	XE8712
Dell PowerEdge AMD Servers 17G	R6725, R7725, R6715, R7715
Dell PowerEdge AMD Servers 16G	R6625, R7625, R6615, R7615
Dell PowerEdge Intel Servers 16G	R660, R760, R760xa, R760xd2, C6620, XE8640, XE9640, XE9680
Nvidia Accelerators	A10, A100, A30, H100, H100NVL, H200SXM, L40, L40s, T4, GB200
Dell Networking Switches	PowerSwitch S3048-ON, S5232F-ON, S5232F-ON, S4148T-ON PowerSwitch N3248TE-ON PowerSwitch Z9332F-ON, Z9264F-ON, Z9264F-ON, Z9664F-ON, Z9432F-ON, Z9864F-ON
Nvidia InfiniBand Switches	NVIDIA SB7890-* EDR Infiniband Switch NVIDIA QM8700-* Quantum HDR InfiniBand Switch NVIDIA QM9700-* NDR Infiniband Switch
Dell PowerVault Storage	PowerVault ME4084, ME4024, ME4012, ME5012, ME5024, ME5084 Storage Arrays
Dell PowerScale Storage	PowerScale F710, F600, H7000, H5600
Operating system on the Omnia Infrastructure Manager	RHEL 10.0 Server with GUI
Operating system deployed by Omnia on bare-metal servers	RHEL 10.0 Minimal
Kubernetes Control Plane on Omnia Infrastructure Manager	1.34.1
Kubernetes Control Plane on Head and Compute Nodes	1.34.1
Slurm Workload Manager	25.05.2
Helm	3.19.0
Full Hardware and Software list can be found on: https://omnia.readthedocs.io/en/latest/Overview/SupportMatrix/index.html	

Solution overview

Telemetry and Cluster Visualization

- Omnia fetches key metrics and logs indicators from iDRAC (power, temperatures) and at the OS level (CPU/GPU/RAM utilization) in the cluster
- Omnia also supports fetching software health and job data from Kubernetes managers and Slurm controllers
- The telemetry data is plotted on [Grafana](#) to provide better visualization capabilities.

Omnia 2.1 Architecture



Internal Use - Confidential

Copyright © Dell Inc. All Rights Reserved.

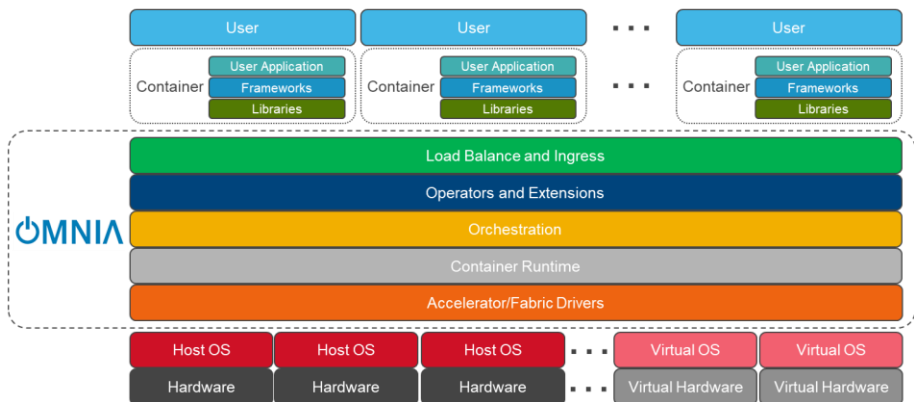
Dell Technologies

Omnia installs Kubernetes, Slurm or both, along with additional drivers, services, libraries, and user applications

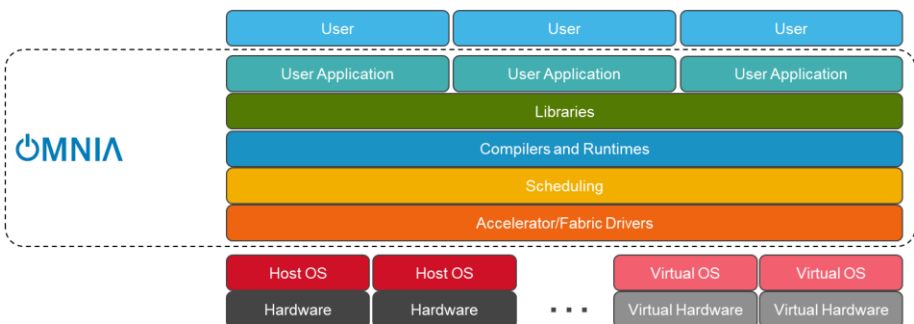
“HPC systems deployment can be difficult, and the addition of AI and data analytics makes it even more complicated. We created Omnia to make it easier, incorporating expertise from Dell's HPC & AI Solutions engineers, our HPC & AI Centers of Excellence, and across the HPC Community. And as the HPC landscape changes, whether that be on premises or in the cloud, Omnia will continue to grow and evolve to meet our customers' and the Community's needs.”

— John Lockman,
Technologist, HPC & AI

Omnia Kubernetes stack



Omnia Slurm stack



Solution overview

“We choose Dell because it’s the best in quality and the best in support. I am not joking. We now have around 600 servers in our data center, including different generations from Dell, and we have statistics that show us that Dell is the best in quality and support.”

— Maurizio Davini
Chief Technology Officer
University of Pisa¹

Arizona State University: [Accelerating scientific research with HPC](#)

Read more [customer stories](#)

Dell Technologies has what you need.

Expertise and guidance

Technology is evolving quickly, so your team may not have time to design, deploy and manage solution stacks optimized for multiple workloads. Dell Technologies has been a leader in the technical computing space for over a decade, with proven products, solutions and expertise. Dell Technologies has a team of AI and HPC 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.

Solutions customized for your environment

Dell Technologies uniquely provides an extensive portfolio of technologies to deliver the advanced computing solutions that underpin successful AI and HPC implementations. With 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 and enable you to capitalize on a universe of data.

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

- [Consulting](#) helps you create a competitive advantage for your business. Our expert consultants work with companies at all stages of analytics to help you plan, implement, and optimize solutions that enable you to unlock your data capital and support advanced techniques, such as AI and ML.
- [Deployment](#) helps you streamline complexity and bring new IT investments online as quickly as possible. Leverage our 30+ years of experience for efficient and reliable solution deployment to accelerate adoption and return on investment while freeing IT staff for more strategic work.
- [Support](#) driven by AI and DL 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.
- [APEX](#) delivers cloud services for a range of data and workload requirements, enabling you to simplify transformation, adapt to evolving conditions, and stay in control of your data. APEX is based on innovative Dell Technologies infrastructure built with flexibility and performance.

“Our partnership with Dell Technologies allows us to take advantage of the full breadth and depth of their compute, storage, networking and security solutions.”

— David J. Brzozowski Jr
Chief Technology Officer
Medacis²

¹ Dell Technologies case study, [Capitalizing on virtualization](#), August 2020

² Dell Technologies case study, [Medacis Advances Healthcare Analytics with AI running on Dell EMC PowerEdge](#) and PowerScale, January 2021.

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 business become more successful and competitive. Dell Technologies 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 for your business? Run demos, try proofs of concept and pilot software in Singapore, Seoul, Sydney, Bangalore and other Customer Solution Centers. Dell Technologies 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.

HPC & AI Innovation Lab

The [Dell Technologies HPC & AI Innovation Lab](#) in Austin, Texas, is a flagship innovation center. Housed in a 13,000-square-foot data center, it gives you access to thousands of Dell servers, three powerful HPC clusters, 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.

HPC & AI Centers of Excellence

As data analytics, HPC and AI converge and the technology evolves, Dell Technologies worldwide [HPC & AI Centers of Excellence](#) provide thought leadership, test new technologies and share best practices. They maintain local industry partnerships to incorporate your feedback and needs into their know-how 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 customers can confidently source information technology needs from Dell Technologies.

- #1 in servers³
- #1 in converged and hyperconverged infrastructure (HCI)⁴
- #1 in storage⁵
- #1 cloud IT infrastructure⁶

See [Dell Technologies Key Facts](#).

³ IDC WW Quarterly Server Tracker, Units & Revenue.

⁴ IDC WW Quarterly WW Quarterly Converged Systems Tracker, Vendor Revenue.

⁵ IDC WW Quarterly Enterprise Storage Systems Tracker.

⁶ IDC WW Quarterly Cloud IT Infrastructure Tracker, Vendor Revenue.

Solution overview

Start working with Omnia.

<https://github.com/dell/omnia>

Learn more.

[HPC/AI Engineering](#)

[Dell.com/HPC](#)

[Dell.com/AI](#)

[Dell.com/Analytics](#)

[InfoHub](#)

Join the HPC Community: [DellHPC.org](#)



Take the next step, today.

Don't wait to reap the benefits of an open-source software solution designed to help you deploy faster, leverage fluid pools of resources, and integrate complete lifecycle management for unified data analytics, AI and HPC clusters. Visit [Omnia on Github](#) and contact your Dell Technologies representative to find out more, today.

Contact us.

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

DELLTechnologies

Copyright © 2026 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

Slurm[®] is a registered trademark of SchedMD LLC. Kubernetes[®], Helm[®], Prometheus[®], and OpenHPC[™] are trademarks of The Linux Foundation. Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat[®], CentOS[®], and Ansible[®] are registered trademarks of Red Hat, Inc. in the United States and other countries. BeeGFS[®] is a registered trademark of Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. PixStor[™] is a trademark of Arcapix Holdings. VMware[®] products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware[®] is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. GitHub[®] is an exclusive trademark registered in the United States by GitHub, Inc. Other trademarks may be the property of their respective owners. Published in the USA 8/22 Solution overview OMNIA-SWSTK-SO-102

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.