

Dell EMC Ready Solutions for Artificial Intelligence

Go from AI-possible to AI-ready

Table of Contents

| | |
|---|-----------|
| The future of artificial intelligence is now. | 2 |
| Dell Technologies has what you need.. | 2 |
| AI, ML and DL use cases. | 3 |
| Do any of these challenges sound familiar? | 4 |
| Dell EMC Ready Solutions for AI | 4 |
| Reduce complexity. | 4 |
| Get deeper insights, faster. | 4 |
| Trust proven AI expertise. | 5 |
| Customer success stories | 5 |
| Dell EMC Ready Solutions for AI technical specifications | 6 |
| Deep learning with NVIDIA | 6 |
| DataRobot | 6 |
| Grid Dynamics | 7 |
| Iguazio | 7 |
| Kubeflow on OpenShift | 8 |
| H2O.ai | 8 |
| Intelligent video analytics | 9 |
| Retail loss prevention | 9 |
| Domino Data Science Platform | 10 |
| Services and financing | 11 |
| Why choose Dell Technologies for AI? | 11 |
| Customer Solution Centers | 12 |
| AI Experience Zones. | 12 |
| HPC & AI Innovation Lab | 12 |
| HPC & AI Centers of Excellence | 12 |
| Proven results | 12 |
| Take the next step, today.. | 12 |

The future of artificial intelligence is now.

Artificial intelligence (AI) and its supporting computing models — machine learning (ML) and deep learning (DL) — are decades-old technologies that are just now beginning to take off. Why is AI so hot right now? The reason is likely a convergence of multiple forces.

First, the industry is making incredible breakthroughs in AI, especially in DL. Second, the mainstreaming of high performance computing (HPC) is making advanced computing power available and affordable for a much wider range of companies. And third, there is a lot more data available to fuel AI — with more being produced every second.

This perfect storm is creating an opportunity for you to quickly identify trends and patterns that otherwise would be difficult and time-consuming to detect. Whatever your industry vertical, AI, ML and DL can change everything. Whether you're just getting started or you've been doing AI, ML and DL for some time, Dell Technologies can help you capitalize on the latest technological advances, saving you time and money while reducing risk.

Dell Technologies has what you need.

Expertise and guidance

Technology is evolving quickly, so your team may not have the resources to design, deploy and manage solution stacks optimized for AI, ML and DL. While advanced data analytics and AI might seem like the latest IT trends, Dell Technologies has been a leader in the advanced computing space for over a decade, with proven products, solutions and expertise. Dell Technologies has a team of data analytics, 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.

Dell EMC Ready Solutions for AI

Dell Technologies is at the forefront of AI, providing the technology that makes tomorrow possible, today. Dell Technologies has invested to create a portfolio of Dell EMC Ready Solutions for AI, simplifying the IT infrastructure to provide faster, deeper insights. You can rely on the Dell Technologies team of experts to help you adapt as AI, ML and DL evolve.

Solutions customized for your environment

Dell Technologies uniquely provides an extensive portfolio of technologies to deliver the advanced computing solutions that underpin successful data analytics, 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.

“We’re changing how we connect devices, how we connect biological systems, and really how we connect people.”

— April Agee Carroll,
Vice President of R&D,
AeroFarms¹

“In the future, we believe data will guide every medical decision. That’s why technology will be key for every healthcare company.”

— Kiyotaka Fujii,
President of Global Healthcare,
Konica Minolta²

AI, ML and DL use cases

The use cases for AI are incredibly diverse and rapidly evolving, but there are common patterns across industries and verticals. Here is a sampling of possible use cases.

| Healthcare and life sciences | Financial services | Government security and defense | Media and entertainment |
|--|---|---|--|
| <ul style="list-style-type: none"> • Population health • Cancer detection • Chronic illness prediction • Drug discovery • Drug interactions • Genetics • Bioinformatics • Sanitation | <ul style="list-style-type: none"> • Fraud prevention • Risk management • Investment predictions • Customer service • Digital assistants • Network security | <ul style="list-style-type: none"> • Facial recognition • Video surveillance • Cybersecurity • Satellite imagery • Event prediction • Emergency services • Weather and climate predictions | <ul style="list-style-type: none"> • Video captioning • Content-based search • Real-time translation • Language processing • Recommendation engines |

| Manufacturing | Energy | Transportation | Retail |
|--|---|--|---|
| <ul style="list-style-type: none"> • Smart manufacturing systems • Factory and demand analytics and optimization • Preventive maintenance • Relationship intelligence • Product and service quality | <ul style="list-style-type: none"> • Wind power generation • Solar forecasts • Oil production optimization • Weather prediction • Demand/consumption predictions | <ul style="list-style-type: none"> • Autonomous vehicles • Pedestrian and object detection • Lane tracking and traffic patterns • Preventive maintenance • Risk assessment • Logistics | <ul style="list-style-type: none"> • Supply and demand planning • Predicting buying behavior • Loss prevention • Upsell, cross-sell opportunities • Customer and product movement tracking |

AI, ML and DL: What's the difference?

These three concepts are closely intertwined. AI is an umbrella term, while ML and DL are the techniques that make AI possible.

ML refers to the process of “training” the machine, feeding large amounts of data into algorithms that give it the ability to learn how to perform the task. DL is an ML technique that uses neural networks as the underlying architecture for training models.

Training DL models is compute-intensive, and results are only as fast as the slowest component of any solution. Fast compute and storage with a lot of memory and high-bandwidth networking will allow models to train faster.

¹ Dell Technologies case study, [A harvest full of insights](#), accessed February 2021.

² Dell Technologies case study, [Realizing X-ray that moves using technology that transforms](#), accessed February 2021.

“We have to control all of the hardware. We have to make sure that we’re using the latest and greatest and that we’re getting the right I/O. We obsess about images per second and about performance and latency, but we also obsess about our cost ... and we can’t get that in the cloud. We need to have hardware that enables us to put our foot on the pedal.”

— Ben Taylor,
Chief AI Officer,
ZIFF.ai⁵

“Using Dell EMC PowerEdge servers, we are able to increase our platform capacity in a matter of weeks versus a matter of months.”

— Jun Chen, Senior Vice
President of Technology
Services, Epsilon⁴

Do any of these challenges sound familiar?

“AI is too complex.”

Customers report that it can take several months to architect a new solution with software, hardware, proof of concept (PoC) and vendor selection. Because each vendor focuses on its own piece, with little integration between hardware and software, you must know servers, GPUs, networks and storage. From there, you must choose an operating system (OS) that is compatible with the ML and DL libraries and frameworks data scientists use. Once it’s assembled, improper integration and tuning can lead to failed jobs, software version incompatibilities or unbalanced configurations resulting in underutilized resources. And each data scientist may have a unique combination of frameworks and libraries that they need for their work, making ongoing maintenance of the environment a complex and time-consuming task.

“Results are too slow, and we haven’t been able to take full advantage of our data.”

Data is growing at an astronomical rate, and it’s impossible to take full advantage of it manually to get insights. The speed of automated image and pattern detection can help provide faster insights. And with historical data sets, you can get deeper insights into, for example, buying behavior. While most organizations know that automation can provide faster, better and deeper data insights, some don’t know how to effectively leverage and scale existing resources while maximizing utilization for AI workloads.

“We don’t have the in-house expertise.”

AI and related computing paradigms are complex and emerging quickly. Trained and experienced staff members are hard to come by, and many organizations do not have the time to develop the in-house skills required to succeed with AI, ML and DL architectures. Unfortunately, this may become evident only as you attempt to evolve from AI PoCs to production-level environments. It’s little wonder that 78% of AI/ML projects stall before deployment — one-third of them at the PoC stage.³

Dell EMC Ready Solutions for AI

Dell Technologies is at the forefront of AI, providing the technology that makes tomorrow possible, today. You can rely on our extensive portfolio of technologies — spanning workstations, servers, networking, storage, software and services — along with the Dell Technologies team of experts to help you adapt as AI, ML and DL evolve over time.

Reduce complexity.

Dell EMC Ready Solutions for AI are validated hardware and software stacks optimized to accelerate AI initiatives, shortening the time to architect a new solution. Customers report that Dell EMC Ready Solutions for AI enable them to increase platform capacity in a matter of weeks vs. months.⁴ They increase data-scientist productivity by offering self-service workspaces, allowing each data scientist to configure their environment from a library of AI models and frameworks in just five clicks. IT operations are also simplified through a single console for monitoring the health and configuration of the cluster.

Get deeper insights, faster.

Dell EMC Ready Solutions for AI deliver unprecedented performance for DL at scale with leading Dell EMC PowerEdge servers with accelerators, high-speed networking and Dell EMC Isilon all-flash scale-out network-attached storage (NAS). Dell EMC Ready Solutions for AI are designed to accelerate performance, eliminate I/O bottlenecks, deliver fast access to larger data sets, and help improve model accuracy. Inferencing at scale can drive actionable responses in real time.

³ Robotics Business Review, [Almost 80% of AI and ML Projects Have Stalled](#), Survey Says, May 2019.

⁴ Dell Technologies case study, [Epsilon](#), accessed February 2021.

⁵ Dell Technologies case study, [AI startup ZIFF.ai revs up its business](#), accessed February 2021.

“As we grew from 600 to over 2,000 clients, we were able to minimize growth in management overhead while delivering the scalability and performance of increased workloads required, thanks to our Dell Technologies infrastructure.”

— David J. Brzozowski Jr,
Chief Technology Officer,
Medacis⁶

Trust proven AI expertise.

To enhance customer success, Dell Technologies consulting services provide expert guidance to bridge the gap between IT, data scientists and the lines of business. In addition to a world-class engineering team, the [Dell Technologies HPC & AI Innovation Lab](#) has three HPC clusters and extensive industry partnerships, bringing together a community of the brightest minds focused on AI, ML and DL.

Customer success stories

Medacis⁶

5 minutes

instead of 24 hours for delivering analytics results

Millions of dollars

in savings due to upholding 99.99% uptime service-level agreements

13% savings

using Dell EMC PowerEdge for AI

Read the case study: [Medacis Advances Healthcare Analytics with AI running on Dell EMC PowerEdge and PowerScale](#).

HPC Ready Architecture for AI and Data Analytics

\$1.4 million

TCO savings

40% fewer

people required

\$300K less

power required

Read the brief: [TCO Analysis—HPC Ready Architecture for AI and Data Analytics](#).

AeroFarms⁶

Millions

of data points collected, 24/7

390X more

productivity than a commercial farm

95% less

water used for the same crop yield

Read the case study: [A Harvest Full of Insights](#).

- [The University of Cambridge UK Science Cloud](#) uses Hadoop to expand the boundaries of AI.
- [Epsilon](#)[®] sends billions of emails daily, with campaign adjustments in real time.
- [Mastercard](#)[®] applies 1.9 million rules to 165 million transactions per hour in a matter of milliseconds.
- [Konica Minolta](#)[®] integrates Internet of Things (IoT), AI and ML to process up to 300 medical images in a single scan and animate them in minutes.
- [ZIFF.ai](#) got up and running in one hour instead of weeks and months of configuration, decision making and troubleshooting.
- [Zenuity](#)[®] can run 50 simulations an hour to advance autonomous driving.
- [OTTO Motors](#)[®] uses AI to process gigabytes of data per hour.

Read more [customer stories](#).

⁶ Dell Technologies case study, [Medacis Advances Healthcare Analytics with AI running on Dell EMC PowerEdge and PowerScale](#), January 2021.

Dell EMC Ready Solutions for AI technical specifications

Deep learning with NVIDIA

Complete solution, including GPU-aware cluster management and modeling frameworks

There has been an explosion of interest in DL, and the plethora of choices makes designing a solution complex and time consuming. Dell EMC Ready Solutions for AI – Deep Learning with NVIDIA is a complete solution, designed to support all phases of DL. It incorporates the latest CPU, GPU, memory, network, storage and software technologies with impressive performance for both training and inference phases. This level of throughput and efficiency makes scaling out ML and DL more feasible.

- **High-performance** — Built around PowerEdge servers with NVIDIA® V100 GPUs, NVIDIA NVLink™ connects multiple V100 GPUs at up to 300GB/s to deliver higher inference performance.
- **Save time and risk** — A GPU-optimized solution stack can shave valuable time from DL projects.
- **Services ecosystem aids success** — Dell Technologies engineers can help you configure, test and tune GPU-enabled hardware and software, with included services to help data scientists discover insight from data more quickly.

| Configuration options | | | |
|---|---|---|--|
| Dell EMC | | Networking | Software |
| PowerEdge servers | Storage | | |
| <ul style="list-style-type: none"> • R740xd (master node) • C4140 with NVIDIA V100 GPUs (compute nodes) | <ul style="list-style-type: none"> • Ready Solution for HPC NFS Storage • Isilon F800 all-flash scale-out NAS | <ul style="list-style-type: none"> • Dell EMC PowerSwitch S3038-ON • Mellanox® SB7800 switch • Mellanox SX6036 InfiniBand® / VPI Gateway | <ul style="list-style-type: none"> • Data Science Provisioning Portal • Bright Cluster Manager® for Data Science |

DataRobot

Everything you need for automated ML

One of the most important elements of any ML platform is its ability to democratize data access — empowering users to quickly and easily build predictive models with full transparency. That’s why many organizations are choosing DataRobot®, an advanced enterprise AI software platform that encapsulates the knowledge, experience and best practices of the world’s leading data scientists into an automated ML solution so you can quickly and easily build highly accurate predictive models without previous coding and ML skills.

- **Accelerate AI success** — Succeed with the team you have in place because DataRobot automates many tasks required to create AI and ML applications.
- **Flex as needed** — Engineering-validated designs can be optimized for your specific workloads and use cases.
- **Leverage best-in-class solutions** — Dell Technologies works closely with DataRobot to bring you the Dell EMC Reference Architecture for DataRobot.

| Configuration options | | |
|-----------------------|---|--|
| Dell EMC | | Software |
| PowerEdge servers | PowerSwitch networking | |
| C6420 | <ul style="list-style-type: none"> • S3148-ON (1GbE) • S5224F-ON (10/25GbE) | <ul style="list-style-type: none"> • DataRobot • VMware® ESXi™ |

Grid Dynamics

Purpose-built ML and DL infrastructure based on the Kubernetes® ecosystem

Grid Dynamics® is at the forefront of cloud-native big data platforms. With expertise in implementing reliable data lakes, data pipelines, data warehouses, ML platforms and end-to-end analytical data platforms, Grid Dynamics can help reduce time to market and costs with a cloud-native stack. This reference architecture provides an optimized foundation for those looking to migrate an existing data platform based on open-source technologies, move from a legacy platform based on proprietary products, implement a new cloud-native platform, or implement a specific capability on a cloud-native stack.

- **Simplify AI and ML** — Dell Technologies collaborates with Grid Dynamics to create powerful solutions that make it easier to adopt and use open-source AI and ML application stacks.
- **Optimize your solution** — Modular building blocks can be combined to build a system that’s optimized specifically for your unique workloads and use cases.
- **Trust the industry leader** — Modernize, automate and transform the data center using industry-leading converged infrastructure, servers, storage and data protection technologies.

| Configuration options | | | |
|---|----------------------------------|--|--|
| Dell EMC | | | Software |
| PowerEdge servers | Storage | PowerSwitch networking | |
| <ul style="list-style-type: none"> • R640 (for high availability Kubernetes) • R740xd (application nodes) | R740xd servers for Ceph® storage | <ul style="list-style-type: none"> • S3148-ON 1GbE • S5224-ON 10/25GbE | <ul style="list-style-type: none"> • Kubeflow on SUSE® CaaS Platform • Kubeflow on Canonical® Charmed Kubernetes Cluster |

Iguazio

Speed and simplify deployment of ML applications.

Iguazio® enables real-time processing of streaming data for rapid time to insight. By unifying the data pipeline, Iguazio reduces the latency and complexity inherent in many advanced computing workloads, effectively bridging the gap between development and operations. Data scientists can run queries on large data sets and securely share data and algorithmic models with authorized users during the training phase. Once the models are ready for production, containerized ML models are easy to move from development to operational environments.

- **Speed and simplify deployment of AI and ML applications** — Essential frameworks, such as Kubeflow, Apache® Spark® and TensorFlow™ along with well-known orchestration tools like Docker® and Kubernetes are built in.
- **Transition smoothly to production** — A complete toolset for developing and deploying ML models and allows both training and production models on one platform, closer to the data source.
- **Reduce risks** — Rely on fine-grained security using multi-layered network, identity, metadata or content-based policies.

| Configuration options | | |
|---|---|-------------------------------|
| Dell EMC | | Software |
| PowerEdge servers | PowerSwitch networking | |
| <ul style="list-style-type: none"> • R640 (application nodes) • R740xd (data nodes) | <ul style="list-style-type: none"> • S3148-ON (1GbE) • SN2700-ON (100GbE) | Iguazio Data Science Platform |

Kubeflow on OpenShift

Accelerate ML/DL workloads using Kubeflow and Dell EMC PowerEdge servers.

Integrating production-grade AI technologies in well-defined platforms within the protection of the data center can facilitate wider adoption of advanced computing, extending investments by supporting AI use cases as well as augmenting the resources available to data science teams. Running Kubeflow on Red Hat® OpenShift® offers several advantages.

- **Streamline AI pipelines** — A composable, scalable, portable stack includes components and automation features to integrate ML tools, so they work together to create a cohesive pipeline that makes it easy to deploy ML applications at scale.
- **Develop and deploy anywhere** — Running Kubeflow on OpenShift makes models portable, so ML/DL engineers can develop models locally and easily deploy the application to a production Kubernetes environment.
- **Enhance simplicity and control** — The ability to run ML/DL workloads in the same environment as the rest of your enterprise applications increases control and reduces complexity for IT teams.

| Configuration options | | | |
|--|--------------------------------------|---|---|
| Dell EMC | | | Software |
| PowerEdge servers | Storage | PowerSwitch networking | |
| <ul style="list-style-type: none"> • R640 (bastion, control plane and infrastructure nodes) • R640 servers with NVIDIA T4 GPUs (application nodes) | R740xd servers with NVIDIA V100 GPUs | <ul style="list-style-type: none"> • S5232F-ON • S3048-ON | <ul style="list-style-type: none"> • Kubeflow • Red Hat OpenShift |

H2O.ai

Improve customer experiences, streamline processes, decrease waste and fraud.

Emerging technologies, such as IoT and mobile technologies and applications, are creating data at speeds and volumes never imagined. This data needs AI to be converted into the fuel that enables you to make better decisions, faster. However, deploying systems capable of running AI workloads can be complex, requiring extensive integration and testing of the hardware and software. That's why Dell Technologies works with H2O.ai to bring you engineering-validated designs for H2O.ai.

- **Trust an industry leader** — Hundreds of thousands of data scientists in over 18,000 global enterprises trust H2O® because it supports some of the most widely used statistical and ML algorithms — including gradient boosted machines, generalized linear models, DL and more.
- **Deploy flexibly** — H2O works on bare metal, with existing Apache Hadoop® or Spark clusters. It can ingest data directly from Hadoop Distributed File System (HDFS), Spark, S3, Microsoft® Azure® Data Lake or other data sources into its in-memory distributed key value stores.
- **Accelerate ML** — Take advantage of the computing power of distributed systems and in-memory computing to accelerate ML using industry parallelized algorithms, which take advantage of fine-grained in-memory MapReduce.

| Configuration options | | |
|-----------------------|--|--|
| Dell EMC | | Software |
| PowerEdge servers | PowerSwitch networking | |
| R740xd | <ul style="list-style-type: none"> • S3148-ON 1GbE (management) • S5224F-ON 10/25GbE (cluster) | <ul style="list-style-type: none"> • H2O.ai Sparkling Water • H2O.ai Driverless AI |

Intelligent video analytics

Create safer and smarter environments with Dell Technologies, NVIDIA and Deep Vision.

New AI capabilities have opened a whole new world of possibilities for monitoring and analyzing video feeds in real time for intelligent insights that can help improve customer service, efficiency, health and safety and much more. However, systems operating at the edge, at scale, require experience and deep application knowledge to determine the right balance of hardware, software and media processing algorithms. That's why Dell Technologies, in partnership with Deep Vision (an IntelliSite™ company) and NVIDIA, is offering a Ready Solution for AI that has been validated by Dell Technologies engineering teams for intelligent video analytics at the edge.

- **Reduce complexity** — Proceed directly to deploying a validated solution on a fast track with lower risks and cost than do-it-yourself approaches.
- **Get the right price/performance** — A wide choice of configuration options enable the right balance of compute, memory and GPU performance with lower costs.
- **Trust proven AI expertise** — The combined expertise of Dell Technologies, NVIDIA and Deep Vision provide the foundation for deploying AI at the edge more easily and with less risk.

| Intelligent video analytics configuration options | | | |
|--|---|---|--|
| Compute | Networking | AI platform | AI modules |
| NVIDIA GPU Cloud (NGC)-Ready for Edge Systems based on Dell EMC PowerEdge R7515, XE2420, R740 or R7525 Servers with NVIDIA T4 GPUs | Dell EMC PowerSwitch N3048ET-ON, N3024EP-ON, N3048EP-ON and S3148P-ON | Deep Vision deployed on: <ul style="list-style-type: none"> • NVIDIA EGX™ software stack for edge AI workloads (Kubernetes, Helm, Tiller and the NVIDIA GPU Operator) • NVIDIA Metropolis, (DeepStream SDK and TensorRT™ SDK) • NVIDIA CUDA® | Each video stream can be configured to one or more Deep Vision module(s): <ul style="list-style-type: none"> • Facial recognition • Vehicle identification • People counting and demographics • Alert zone entry • Thermal analysis |

Retail loss prevention

Cost-effectively reduce inventory loss with AI-driven software and hardware technology.

Preventing scan fraud activities such as scanning errors and UPC barcode switching at the point of sale (POS) is critical to reducing the major source of inventory loss in the retail sector. Dell Technologies, in partnership with Malong Technologies and NVIDIA, offers a self-learning computer vision solution that uses AI to reduce inventory loss caused by accidental and intentional behavior during checkout. AI-based automated systems using Dell Technologies hardware and special purpose DL models from Malong provide cost-effective tools to reduce inventory loss for retail stores.

- **Reduce losses** — Protect a wide range of stock keeping units (SKUs) by detecting mis-scans and ticket switching in near real time.
- **Reduce costs** — Reduce reliance on cost-intensive human security personnel.
- **Deploy and scale quickly** — Validated configurations enable easy deployment at scale for thousands of retail stores at self-checkout and staffed checkout lanes.

| Configuration options | | | |
|--|------------------------|---|--|
| Dell EMC PowerEdge servers | Storage | Networking | Software |
| R740, R7515 or R7525 with NVIDIA T4 GPU and associated software (CUDA, DeepStream SDK, etc.) | 2TB SAS SSDs in RAID 6 | Broadcom 57416 Dual Port 10 GbE BaseT Network LOM Mezz Card | <ul style="list-style-type: none"> • RetailAI® Protect by Malong Technologies • Docker container runtime |

Domino Data Science Platform

Speed development, delivery and monitoring of AI models with an optimized technology platform.

One of the most complex tasks for getting AI workloads into production is model development. Data scientists need to be able to streamline development to deployment by using a platform that accelerates the data science lifecycle with open access to tools, compute and data. Enterprises require an open and flexible data science platform that automates tasks while allowing teams to collaborate efficiently, enabling the use cases of today while not locking out those of tomorrow.

- **Speed time to value** — Standardized building blocks simplify designing systems optimized for specific workloads and speed configuration and ordering.
- **Enhance efficiency** — Accelerate the data science lifecycle with open access to tools, compute and data that streamline development and deployment.
- **Augment internal skills** — Dell Technologies AI experts are available to assist with designing a solution for specific needs.

| Configuration options | | | |
|--|---------------------------------|---|---|
| Dell EMC PowerEdge servers | Networking | Storage | Software |
| <ul style="list-style-type: none"> • R640 (bastion, control plane and infrastructure nodes) • R640 with NVIDIA T4 GPUs (application nodes) | 25 GbE Mellanox ConnectX®-5 NDC | 2x 240GB or 480GB SATA SSD 1TB plus SSD storage for project files | <ul style="list-style-type: none"> • Domino® Data Science Platform • Kubernetes |

“Our partnership with Dell Technologies allows us to take advantage of the full breadth and depth of their compute, storage, networking and security solutions.... The PowerEdge R740 provides the agility to adapt to ever-changing business needs, including running Kubernetes. We’re able to run our machine learning workloads on NVIDIA GPUs within the same platform, resulting in cost savings on both hardware and operational overhead.”

— David J. Brzozowski Jr,
Chief Technology Officer,
Medacis⁶

Solution highlights

Dell EMC PowerEdge servers: Dell EMC PowerEdge servers are engineered to deliver unmatched performance and versatile configurations to meet the demands of AI workloads. Flash storage, the latest processors, greater memory bandwidth and flexible local storage make Dell EMC PowerEdge servers a foundational choice for AI.

Dell EMC PowerSwitch networking: Today’s AI workloads call for new thinking about network architecture. Based on open standards, Dell EMC PowerSwitch networking frees the data center from outdated, proprietary approaches. Our future-ready networking technology helps you improve network performance, lower networking costs and remain flexible to adopt new innovations. Take control of your network’s future and learn how the Dell Technologies strategy for open networking can dramatically transform your business.

Dell EMC PowerScale storage: AI environments require large, scalable, reliable and efficient storage. With support for multiple workloads and enterprise-grade data and file management capabilities out of the box, Dell EMC Isilon scale-out NAS is the leading storage for AI. You can take advantage of the high capacity of Isilon to reduce the acquisition and ownership cost for managing and monetizing data using advanced or predictive analytics and ML.

Dell EMC Omnia software: Dell EMC Omnia is an open-source, Ansible®-based software stack designed to automate the deployment of mixed-workload clusters, giving IT the agility to run AI, HPC and data analytics workloads in the same environment, with a single pane of glass for cluster provisioning, deployment and management, with easy-to-use point-and-click templates for building environments.

Services and financing

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

- [Data Analytics Consulting Services](#) help you create a competitive advantage for your business. Our expert consultants work with companies at all stages of data analytics to help you plan, implement and optimize solutions that enable you to unlock your data capital and support advanced techniques, such as AI, ML and DL.
- [Deployment Services](#) help 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 (ROI) while freeing IT staff for more strategic work.
- [Support Services](#) 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 just 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 On Demand](#) 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.

“With [Dell Technologies], we get world-class support, so we can avoid the finger-pointing you get with competing vendors. This is a key to our relationship.”

— Barry Carter, Chief Information Officer, RealPage⁷

Why choose Dell Technologies for AI

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

- Schedule an [executive briefing](#) and collaborate on ways to reach your business goals.
- [Dell Technologies Customer Solution Centers](#) are staffed with computer scientists, engineers and 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 at [HPCatDell.com](#).
- Dell Technologies 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 Technologies 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.

⁷ Dell Technologies case study, [RealPage Boosts Property Management Performance and Results with Data Analytics](#), July 2019

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 information technology needs from Dell Technologies.

- #1 in servers⁸
- #1 in converged and hyperconverged infrastructure (HCI)⁹
- #1 in storage¹⁰
- #1 cloud IT infrastructure¹¹

See [Dell Technologies Key Facts](#).

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

Are you 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 the flagship innovation center. Housed in a 13,000-square-foot data center, it gives you access to thousands of Dell EMC 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 and have direct access to Dell and other technology creators to incorporate your feedback and needs into their roadmaps. Through collaboration, Dell Technologies HPC & AI Centers of Excellence provide a network of resources based on the wide-ranging know-how and experience in the community.

Take the next step, today.

Don't wait to harness the benefits of AI on optimized solutions designed from the ground up to be AI-ready. Contact your Dell Technologies representative to find out more today.

Contact us

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



Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

Medacast® is a registered trademark of Medacast Solutions Group, LLC. AeroFarms® is a registered trademark of Just Greens, LLC. Epsilon® is a registered trademark of Epsilon Data Management, LLC. Mastercard® is a registered trademark or service mark of Mastercard or its subsidiaries in the United States. Konica Minolta® is the registered trademark of Konica Minolta, Inc. Zenuity® is a registered trademark of Zenuity AB. OTTO Motors® is a trademark of Clearpath Robotics, Inc. NVIDIA®, EGX™, NVLink™, TensorRT™, CUDA®, Mellanox®, InfiniBand®, and ConnectX® are trademarks or registered trademarks of NVIDIA Corporation and/or Mellanox Technologies in the U.S. and other countries. Bright Cluster Manager® is a trademark of Bright Computing, Inc. DataRobot® is a registered trademark of DataRobot Inc. in the United States and/or other countries. Kubernetes® is a registered trademark of The Linux Foundation. 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. Grid Dynamics® is a trademark or registered trademark of Grid Dynamics International, Inc. in the United States and other countries. Red Hat®, Ceph®, Ansible®, and OpenShift® are registered trademarks of Red Hat, Inc. or its subsidiaries in the U.S. and other countries. SUSE® and the SUSE logo are registered trademarks of SUSE LLC. Canonical® is a registered trademark of Canonical Ltd. Iguazio® is a trademark of Iguazio Systems, Ltd. Microsoft® and Azure® are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apache®, Hadoop® and Spark® are trademarks of the Apache Software Foundation. TensorFlow™ is a trademark of Google, Inc. Docker® is a trademark or registered trademark of Docker, Inc. in the United States and/or other countries. H2O® is a trademark of H2O.ai. RetailAI® is a trademark of Malong Technologies. DOMINO® is a registered trademark of Domino Data Lab, Inc. IntelliSite™ is a trademark of Koninklijke Philips N.V. (Royal Philips). Other trademarks may be the property of their respective owners. Published in the USA 03/21 Solution overview RS-AI-SO-106

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