

Challenges to AI adoption

Legacy infrastructure creates real barriers to progress when it comes to AI, but integrating new technology comes with its own set of challenges. Many organizations find their AI journey stalled by three main obstacles:



Outdated infrastructure

Legacy systems and traditional infrastructure aren't equipped to meet the demands of AI workloads. Aging hardware lacks the performance and capacity needed, leading to slow processing, bottlenecks and an inability to keep up with innovation. Additionally, older systems are not designed to scale efficiently, making it difficult to handle growing data volumes and complex workloads.



Limited processing power

Underpowered IT slows down complex models, delays insights and makes it difficult to take full advantage of data-driven outcomes. Older PCs lack the multitasking ability to juggle modern workloads and tasks, leading to slowdowns and bottlenecks that frustrate employees and get in the way of productivity.



Skills gaps

Managing a complex AI environment can be daunting without a team of dedicated data scientists. IT departments need solutions that are simple to deploy, manage and scale without requiring extensive specialized expertise.



Over time, these limitations can increase costs, create vulnerabilities and hamper innovation. By proactively addressing these challenges with the right technology partner and modernization strategy, you can transform obstacles into opportunities, setting the stage for AI-driven growth and resilience.


The Dell PowerEdge solution

PowerEdge servers with the 5th Gen AMD EPYC™ processor give medium businesses an infrastructure designed for every step of the AI journey. Within this portfolio, two clear categories — efficiency champions and performance powerhouses — help you match workloads with the right technology to drive meaningful outcomes every day.


Efficiency champions

Dell PowerEdge R4715 and R5715

Built for general-purpose computing, virtualization and entry-level AI inference, the R4715 and R5715 servers deliver the performance you need in a straightforward, single-socket design. This results in lower upfront costs, reduced hardware footprint and predictable software licensing, without sacrificing enterprise capabilities. These servers are the smart choice for organizations looking to achieve more with less, and to enter the AI space without overprovisioning or complexity.



The Dell PowerEdge solution



Performance Powerhouse

Dell PowerEdge R6715 and R7715

For businesses ready to scale AI workloads, the R6715 and R7715 servers deliver uncompromising performance. Engineered for high-density computing, model training and complex analytics, these servers offer high core counts, broad memory support and seamless GPU integration. The combination of advanced AMD EPYC™ architecture and engineered Dell reliability means your infrastructure is ready for even the most demanding tasks, now and as your business scales.

Up to 51%
generational CPU performance improvement¹



With every model, built-in features like PCIe Gen5 connectivity and lightning-fast NVMe® storage help keep data moving and insights flowing. Dell OpenManage tools streamline deployment and ongoing management so your team can focus more on outcomes and less on operational overhead. With security integrated at every layer and flexible deployment models, Dell PowerEdge servers with the latest AMD EPYC™ processor empower your organization to innovate with confidence and clarity.

¹ Based on Dell analysis of specifications comparing the Dell PowerEdge R77x5 server with up to 40 E3.s drive slots with the Dell PowerEdge R76x5 servers with up to 36 E3.s drive slots. Data was collected as of 10/2/2024.

CPU vs. GPU

Choosing the right tool for the job

A smart AI strategy starts with understanding which workloads are a good fit for CPU and which require GPU acceleration. While GPUs are essential for certain high-performance use cases, not all AI workloads require specialized, high-cost accelerators. A modern CPU is highly capable and efficient for a wide range of AI tasks.

CPU-based AI



Versatile, efficient and cost-effective

For many AI applications — especially inference, analytics and running smaller or pre-trained models — a modern CPU provides the ideal foundation. The 5th Gen AMD EPYC™ processor delivers exceptional core density and performance, making it well-suited for running both AI and traditional business workloads side by side. This versatility helps maximize IT investment, lowers costs and avoids overprovisioning. Starting AI initiatives on a CPU allows you to scale capabilities at your own pace, making it a practical, efficient entry point into AI.

GPU-based AI



Accelerating advanced, demanding workloads

For compute-heavy tasks such as deep learning model training and large-scale simulations, GPUs are critical. Their parallel processing architecture dramatically accelerates complex mathematical calculations, cutting training times and powering rapid innovation. This makes GPUs indispensable when your organization needs to build or iterate on advanced models quickly.

The Dell PowerEdge advantage: Seamless scalability and flexibility

Building a cost-effective AI infrastructure requires making smart decisions about your hardware, with flexibility as a key guiding principle. With Dell and AMD, you don't have to choose a single path. Dell PowerEdge servers with AMD EPYC™ processor are built to support both CPU and GPU driven AI strategies, giving you flexibility as your needs evolve.

You can start with cost effective, CPU based AI to address your immediate needs, and seamlessly add GPU acceleration as your workloads grow in scope or complexity. This scalable architecture lets you make smart, cost effective infrastructure decisions, so you always have the right foundation for every phase of your AI journey, helping you achieve meaningful results with clarity and confidence.

Real-world applications



Modernizing with Dell PowerEdge servers and the latest AMD EPYC™ processor isn't just about potential. It delivers measurable results for real organizations.

A manufacturing business replaced 75 aging servers with just 15 PowerEdge R7715 servers, leading to an 80% savings in server count and rack space and the potential to save on additional costs, such as admin time to perform routine maintenance and management, power, and cooling. This consolidation gave them the performance of seven legacy servers in a single, energy-efficient unit, drastically lowering total cost of ownership (TCO) and boosting operational efficiency.⁴

In sectors like healthcare, organizations have used Dell PowerEdge servers to accelerate AI-driven analytics and double their image processing speed, leading to faster insights and improved outcomes for patients.

Medium-sized businesses running the Dell R4715 and R5715 servers benefit from right-sized, enterprise-grade performance for analytics and virtualization, achieving predictable costs and scalable growth without overprovisioning.

These examples underscore how Dell and AMD help IT leaders turn modernization into business value by boosting speed, lowering costs and creating a reliable, future-ready foundation for AI success.


⁴ Principled Technologies, "Modernizing your data center with Dell and AMD," accessed July 21, 2025. Disclaimer: Information shown is based on real Dell customer data. Actual results will vary. 5:1 average rate guaranteed across customer applications for reducible data. Rates for individual applications may vary. See terms and conditions for details.

Next Steps

Build your end-to-end AI ecosystem with Dell and AMD

The journey toward AI-powered progress starts with a foundation designed for your long-term goals. Wherever you are on your modernization path, Dell Technologies, together with AMD, delivers secure, scalable solutions to move your organization forward.

Don't let legacy technology hold you back. To unlock the full potential of AI:



With Dell, you gain a trusted partner invested in your success, delivering the tools and expertise to help you confidently achieve your AI goals and future-proof your business.

[Learn More](#)

Dell.com/Servers/AMD