The Enterprise AI ROI Era Has Arrived

March 16, 2026

From Investment to Business Driver: Three Moves to Break the AI ROI Barrier

By Arthur Lewis, President, Infrastructure Solutions Group, Dell Technologies

SAN JOSE, Calif., March 16, 2026 /PRNewswire/ -- In March 2024, we launched the Dell AI Factory with NVIDIA to help operationalize AI and prove value. Today, with more than 4,000 customers deploying the Dell AI Factory and early adopters seeing up to 2.6x ROI within the first year[1], we're sharing what we've learned about what it really takes to make enterprise AI successful.

The Value Problem

The enterprise AI conversation has shifted since 2024. Two years ago, customers were asking "How do we get access to AI technology?" Now, they're asking "How do we make our data AI-ready? How do we operationalize this at scale? And, most importantly, how do we prove ROI?"

When our customers ask us about the ROI of AI, we understand the instinct — it's how we've evaluated technology investments for decades. But I'd challenge us all to reframe the question. AI isn't a point solution you bolt on and measure in isolation. It's a fundamental shift in how work gets done, how insights are generated and how value is created across every part of the business. The real risk isn't that you invest in AI and the returns disappoint — it's that you wait for a spreadsheet to tell you it's safe while your industry transforms around you. The better question isn't 'what's the ROI of AI?' It's what is the cost of not being AI-ready when your competitors are.

Why This Moment is Different

We just had our best enterprise AI quarter ever in Q4 2025, and the momentum reflects a fundamental shift in enterprise strategy. As AI code assistants and agentic workflows drastically lower the cost and time to build custom applications, the traditional "build vs. buy" equation is being rewritten. CIOs are increasingly choosing to develop AI capabilities in-house—and that shift is driving the need for owned infrastructure.

The logic is straightforward: building custom AI applications requires training, fine-tuning and running inference on proprietary corporate data. With 83% of the world's data sitting on-premises, the economics favor bringing compute to the data. But here's the challenge: most of this data is in cold backup, is dark or otherwise unprepared for ingest by AI agents and engines. Data not available to AI is value lost.

Three Requirements for Success

Working with over 4,000 customers—from neoclouds to sovereign entities to enterprises to research institutions—has revealed three critical requirements for achieving measurable returns from AI:

1. Making enterprise data AI-ready

The first place that transformation has to start is your data. Here's the reality most enterprises are living with today: massive amounts of information — valuable information — sitting in backup, in silos, in formats that AI simply can't touch. We talk about data as the fuel for AI, but for most organizations, that fuel is frozen in place. It's time to thaw it out.

That's exactly why we built the Dell AI Data Platform with NVIDIA — to address the full data lifecycle, from preparation through high-performance storage. At the heart of it is our new Data Orchestration Engine. Think of it as the intelligence layer that discovers, labels, enriches, and transforms your data — structured, unstructured, multimodal — into governed, AI-ready datasets at scale. No code required. And it gets smarter over time through active learning and human-in-the-loop workflows, so your data quality and your model accuracy keep improving together.

Now, once your data is AI-ready, you need infrastructure that can move it at the speed AI demands. Dell's Lightning File System delivers up to two times greater throughput per rack unit — the world's fastest parallel file system. Pair that with Dell Exascale Storage at up to six terabytes per second of read performance per rack, and you start to see what's possible: customers are achieving up to 12x faster vector indexing and 19x faster time-to-first-token compared to traditional approaches.

These infrastructure innovations are paired with NVIDIA CUDA-X accelerated data and AI libraries to dramatically accelerate the most data-intensive stages of AI pipelines—from ingest and transformation to embedding and retrieval.

This isn't incremental improvement. This is removing what has become the primary technical barrier to AI deployment — getting your data ready, and getting it there fast.

2. Scaling infrastructure from desktop to data center

You need infrastructure that scales efficiently from pilot to production and keeps workloads running at full speed without bottlenecks. Our new Dell Pro Precision workstations deliver the power and expandability AI developers need. We're the first OEM to ship the GB300 Grace Blackwell Ultra Desktop Superchip with our Dell Pro Max desktop, bringing enterprise-grade AI computing directly to developers' desks.

Here's what makes this different: with NVIDIA NemoClaw and OpenShell, Dell Pro Max desktops let developers build and deploy autonomous,

self-evolving AI agents that run for hours or days, learning and adapting as they work—all locally, on sensitive data, without ever touching the cloud. That's frontier-level intelligence at your desk.

At the data center level, our liquid-cooled Dell PowerEdge servers – including the flagship XE9812 with NVIDIA Vera Rubin NVL72 platform – deliver the performance enterprises need for massive training and inference workloads. This infrastructure enables everything from training large language models to running real-time inference for autonomous AI agents. High-performance AI networking, including our PowerSwitch SN6000-series with 1.6TbE liquid-cooled switches, ensures data moves at the speeds AI demands, keeping GPU resources fully utilized rather than sitting idle waiting for data.

3. Compressing deployment timelines

You need solutions, software and services that accelerate your time to value. Our modular architecture, combined with the Dell Automation Platform, enables rapid deployment of validated AI workloads—compressing timelines from months to days.

This is where technology becomes business value. Our knowledge assistant gives employees instant access to institutional knowledge, while our Agentic AI Platform—developed in collaboration with Cohere North and DataRobot—lets autonomous AI agents handle complex workflows from customer service to supply chain optimization. Dell Accelerator Services bridge the gap from experimentation to enterprise-wide deployment, closing skill gaps that often delay ROI.

The Only Integrated Approach

What makes this work is integration. Data platforms, infrastructure and services aren't separate purchases—they're components of a system designed to work together with NVIDIA AI infrastructure and software at the core. The Dell AI Factory with NVIDIA brings all of these pieces together, so the customer doesn't have to: the compute, the network, the storage, the data platform, the SW ecosystem and the services. That integration is what creates the proven path from AI investment to business outcome.

It's also changed our relationship with customers. We're now brought in at question one—"Where do I start?"—not question five: "Send me a quote.' We're a strategic advisor throughout the journey, not just a vendor at the end.

The Path Forward for Enterprise AI

CEOs are seeing 20-30% productivity gains and asking every department to match them. Our advice: organize, be methodical, but move. You'll make mistakes — experiment boldly, fail fast, and keep going. The cost of waiting is higher than the cost of learning.

Two years in, we're more convinced than ever that enterprise AI success isn't just about the most advanced technology — it's more about an integrated approach that turns technology into measurable business results. The over 4,000 customers deploying the Dell AI Factory with us prove the model works.

For enterprises still stuck between pilot and production, the lesson is simple: integration matters, data readiness matters, deployment expertise matters. A partner who delivers all three is the difference between AI as an experiment and AI as a business driver.

Learn more about the Dell AI Factory with NVIDIA

[1] Based on Enterprise Strategy Group paper commissioned by Dell, "Analyzing the Economic Benefits of the Dell AI Factory with NVIDIA," comparing the ROI of on-premises Dell and NVIDIA solution, August 2025. Estimated costs were modeled utilizing Llama 3 70B LLM for inferencing and model fine-tuning workloads by organizations over a 4-year period. Server models used were XE9680s with 8 x H100 GPUs. Actual results may vary.

View original content to download multimedia:https://www.prnewswire.com/news-releases/the-enterprise-ai-roi-era-has-arrived-302715128.html

SOURCE Dell Technologies

Media Relations: Media.Relations@Dell.com