NetApp

# SCALE AI TO VALUE

## The infrastructure strategy for production-first CIOs

# Table of contents

# Executive overview

**Most AI pilots never make it to production. Here's why: Enterprise AI teams spend as much as 70% of their time on data preparation and pipeline management.[1]**

Leading organizations solve this major roadblock by building infrastructure that's ready for enterprise scale from day one with:

### Unified data platforms
that connect information across your entire organization

### Hybrid cloud strategy
that gives you flexibility without adding complexity

### Enterprise cyber resilience
that protects without slowing anything

**The outcomes:**
Faster value realization, lower costs, and IT transformed from a cost center support role to a driver of business advantage—with clear, measurable ROI that commands executive attention.

**Now is the moment for decisive leadership.** Production-first leaders who build AI-ready infrastructure today will shape and dominate tomorrow's AI-driven markets.

**The strategic question for technology leaders**

How do you scale AI from lab experiments to real business impact—flexibly, efficiently, and securely?

**Data management is the #1 AI adoption roadblock.**

CIOs who solve data management complexity while proving clear ROI from AI infrastructure investments gain decisive advantages in the AI transformation race.[2]

1. Forrester, via Berkindale 2. IDC via CIO.com

# 1 Successful AI starts with these three priorities.

## AI changed everything. Is your infrastructure able to adapt?

AI demands that we rethink infrastructure from the ground up, optimizing AI pipelines from data ingestion to training to inference.

### Here's what AI leaders prioritize:

**Unified data platforms**
Make all data instantly available and easily manageable across the pipeline—from on-premises to the cloud to the edge—all under one view.

**Hybrid cloud flexibility**
Use any major cloud—compute meets data, not the other way around. Stop costly data movement, make use of the best cloud data services, and simplify management—without vendor lock-in.

**Cyber resilience**
Build in data protection and security at the data layer, eliminating the compliance, security, and governance gaps that slow AI deployments.

### The data problem slowing AI innovation

# 60%

of organizations don't have the right data management practices to support enterprise-scale AI, putting the success of their AI efforts at risk. In fact, Gartner predicts that through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data.[3]

3. Gartner

# 2 Your biggest AI barrier isn't algorithms. It's a lack of unified data.

## AI poses a fundamental IT challenge: Your most valuable data is scattered across systems, clouds, and edge locations—making it nearly impossible for AI to access and use.

Unified data solves this. Instead of wrestling with fragmented systems, a unified platform creates seamless, secure access to all your enterprise information from a single control plane, while optimizing for AI workloads.

### What you can do with unified data:

- **Move data where it's needed**
  Access data efficiently across environments in hours, not weeks, based on AI workload requirements.
- **Optimize resource usage**
  Automatically allocate high-throughput storage for data preparation, elastic compute for model training, and low-latency access for inference—without over-provisioning.
- **Simplify governance**
  Apply consistent policies, access controls, and compliance frameworks across all environments automatically.

### The business impact:

- **Faster development** – Data prep happens in days, not months
- **Seamless scaling** – Go from pilot to production without rebuilding
- **Consistent performance** – Every AI project gets the resources it needs, when needed
- **Automatic governance** – Security and compliance work across all data, automatically

### Meet the biggest AI killer: data fragmentation

IT organizations manage a staggering volume of data copies—all from valid activities like backup and analytics.

**6.4** data silos on average

**13** data silos across primary, secondary, cloud, and edge storage

This creates much more than an IT management burden. Obsolete data copies can cause AI models to "go backward," learning from outdated information that degrades performance. Intelligent data versioning that maintains authoritative datasets and eliminates obsolete copies ensures AI models work from a single source of truth—protecting the data integrity that determines AI success.[4]

# 3 Run AI wherever it performs best—without cloud complexity.

## The AI and cloud relationship is flawed: AI workloads don't respect cloud boundaries.

Training happens where compute is cheapest. Inference runs where users are. Data stays where regulations require. AI is inherently a hybrid multicloud workload, yet most storage vendors only work in one cloud. This creates a costly problem. Multiple platforms mean different tools, vendors, APIs, and management overhead. This complexity slows AI development and inflates costs.

**Hybrid cloud flexibility solves this.** One unified storage operating system works across all major clouds—AWS, Azure, Google Cloud—plus on-premises infrastructure. Manage everything from one place, and:

- **Put workloads where they perform best**
  Train large models in cost-efficient environments, then deploy inference at cloud scale where usage-based pricing aligns with business value.
- **Scale resources based on demand**
  Burst to cloud for intensive training, scale back for inference—paying only for what you use, when you use it.
- **Deploy globally, optimize locally**
  Run AI at the edge of business operations while keeping centralized governance and consistent policies everywhere.

### The business impact:

- **Better performance, lower cost** – Optimize workload placement
- **Best-of-breed data services** – Access leading AI capabilities across all major cloud platforms
- **Vendor independence** – Keep data moving freely without lock-in or migration costs
- **Global scalability** – Expand across regions while maintaining compliance automatically

### The multicloud integration gap

# 60%

By 2030, more than 60% of organizations will conduct intensive AI model activity across multiple clouds.[5] Yet organizations struggle to realize business value from their multicloud deployments because they lack integration across systems.

5. Gartner

# 4 Stop bolting on security. Build it in.

## Traditional security approaches create a painful tradeoff: slow AI development or accept higher data risk.

**Here's the problem:** AI data moves constantly—from on-premises storage during training to cloud environments for inference to edge locations for real-time decisions. Point solutions only protect individual systems, not entire pipelines. Security gaps emerge. AI projects stall. Built-in data security changes this equation. Protection happens automatically at the data layer. Security teams get the protection they need. Data scientists get the speed they want.

### How it works:

**Secure access everywhere**
Identity-based controls ensure only authorized users access AI data—on-premises or in the cloud. Data science teams collaborate securely with business units and external partners.

**Intelligent data governance and proactive threat detection**
Automatically identify and protect sensitive data used in AI training and inference, ensuring compliance while giving data scientists access to what they need. Detect anomalous AI behavior instantly—spotting model poisoning, data corruption, or adversarial attacks before they impact business outcomes.

### The business impact:

- **Faster deployment** – End security review bottlenecks that delay AI projects
- **Protected investments** – Safeguard AI models and training data from ransomware and attacks
- **Market expansion** – Enter regulated industries with automated compliance and governance
- **Improved cyber resilience** – Maintain availability with automated threat prevention and rapid recovery

### The hidden costs of 'secure it later'

**ONLY 28%**

of companies embed security controls from the start.

Prioritizing speed and innovation over security, most organizations add on point solutions after deployment that create more complexity and cost.[6]

Here's the result: Adding security controls later leads to higher risk and costly retrofitting.

6. Accenture

SUMMARY

# Start with data. Win with AI.

## Here's the reality: AI success starts with data infrastructure. Not algorithms. Not compute. Data.

The challenge most CIOs face? Data lives everywhere—on-premises, in clouds, at the edge. It's hard to manage. It's sensitive, requiring strict compliance. It's vulnerable to attack. And traditional infrastructure wasn't built for AI at enterprise scale.

Production-first leaders solve this with a unified data platform that connects everything—securely mobilizing data across any environment and simplifying management while maintaining governance and protection.

## Unlock AI scale and value now

Leaders who deliver long-term business value start with an AI-ready data foundation. Build your intelligent data infrastructure advantage today.

---

**Is your data infastructure AI-ready? →**

Take the IDC self-assessment.

---

**Learn more about NetApp AI solutions →**

### This is your moment

Build the foundation now. Lead the AI transformation. The organizations dominating tomorrow's AI-driven markets are the ones making smarter infrastructure decisions today.

## Discover the topics in this series

- Data infrastructure modernization →
- Cloud →
- Cyber resilience →
- Artificial intelligence

# About NetApp

For more than three decades, NetApp has helped the world's leading organizations navigate change—from the rise of enterprise storage to the intelligent era defined by data and AI. Today, NetApp is the Intelligent Data Infrastructure company, helping customers turn data into a catalyst for innovation, resilience, and growth.

At the heart of that infrastructure is the NetApp data platform—the unified, enterprise-grade, intelligent foundation that connects, protects, and activates data across every cloud, workload, and environment. Built on the proven power of NetApp ONTAP, our leading data management software and OS, and enhanced by automation through the AI Data Engine and AFX, it delivers observability, resilience, and intelligence at scale.

Disaggregated by design, the NetApp data platform separates storage, services, and control so enterprises can modernize faster, scale efficiently, and innovate without lock-in. As the only enterprise storage platform natively embedded in the world's largest clouds, it gives organizations the freedom to run any workload anywhere with consistent performance, governance, and protection.

With NetApp, data is always ready—ready to defend against threats, ready to power AI, and ready to drive the next breakthrough. That's why the world's most forward-thinking enterprises trust NetApp to turn intelligence into advantage.

Learn more at www.netapp.com or follow us on X, LinkedIn, Facebook, and Instagram.



NetApp