

# How AI leaders architect differently


A closer look at AI-ready infrastructure from the Cisco AI Readiness Index 2025

The Cisco AI Readiness Index measures how prepared organizations are to scale AI, drawing on a double-blind survey of over 8,000 senior leaders from large companies in 30 markets and 26 industries. The Index assesses readiness across six pillars: Strategy, Infrastructure, Data, Governance, Culture, and Talent.

These insights from the Infrastructure pillar compare the most AI-ready companies globally—the **Pacesetters**—to the average for all Global companies we surveyed.



**Pacesetter advantage**  
Deployed AI at the scale and speed necessary to realize value



## How do they do it?

Four architectural bets on infrastructure

### Four architectural bets Pacesetters are making to build AI-ready infrastructure

#### ⏻ Bet 1: Power

96% Global Pacesetters

43% Global Avg

Built dedicated infrastructure to optimize power consumption

#### 🌐 Bet 2: Network

81% Global Pacesetters

22% Global Avg

Built networks optimal for AI workloads

#### ⚙️ Bet 3: Operations

72% Global Pacesetters

24% Global

Continuous monitoring of AI deployments with automated alerts and scheduled retraining of models

#### 🔒 Bet 4: Security

84% Global Pacesetters

30% Global

End-to-end encryption with continuous monitoring and instant threat response

Note: All the data in this report comes from the Cisco AI Readiness Index 2025. And any mention of 'Pacesetters' or companies refers only to respondents from that study.

### ⏻ Power

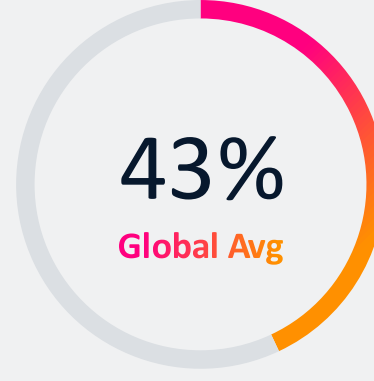
#### Bet 1: Build for power constraints before hitting them

Companies are planning for massive workload growth. But workloads need power to grow. And extra power can be challenging to provision quickly.




**The power infrastructure gap**

Built dedicated infrastructure to optimize power consumption in AI deployment.




#### The workload-power collision

Who's expecting workloads to increase by 50%+ in 3-5 years?



Who's building new data center capacity in the next 12 months?



What this means: Pacesetters have already solved for the power required to run AI at scale (96%). Meanwhile, Global companies on average expect massive workload growth (51%) and some are building capacity (43%)—but only 43% have the power infrastructure to run it. That's a collision waiting to happen.

#### The other side of the coin: AI for power efficiency

Expect power efficiency gains from AI within a year



What this means: Pacesetters aren't just getting the energy infrastructure in place to power AI. They're more likely to use AI to drive power efficiency. And they expect to see results within a year.

#### Key insight



The constraint: Power infrastructure could take 18-36 months to provision. Compute can be added in weeks. Power cannot.

The bet: Pacesetters build for the constraint that could be difficult to retrofit. Global companies risk hitting the power wall at full speed.

### 🌐 Network

#### Bet 2: Treat network as foundation, not afterthought

Many companies are talking about GPUs. But when workloads double, networks could become the bottleneck before compute does. And unlike compute, it may not be as easy to fix network infrastructure when it's under production load.

#### The GPU vs Network reality

Most organizations have adequate GPU resources for current projects.

But the network is the restraint companies might not see coming: While GPU availability varies, network infrastructure gaps present a more universal constraint for AI workloads at scale. GPU resources mean nothing if your network can't move data between them efficiently.

'Just enough' or 'robust' GPU infrastructure




#### The network performance gap

	Global Pacesetters	Global Avg
Networks "optimal" for AI workloads	81%	22%
High-level network integration for efficient data flow and seamless operations	90%	26%
Networks "fully flexible and adaptable" for growth	71%	15%

What this means: Global Pacesetters have built networks that are already optimal for AI and can flex as workloads grow. Global companies on average have networks that are struggling with current AI workloads and that won't scale gracefully.


#### How AI integrates with existing infrastructure



What this means: Pacesetters achieve high integration across their infrastructure stack. Global companies average only 30% integration—meaning AI deployments operate in silos, disconnected from existing systems. Without integration, AI can't scale, can't be monitored effectively, and can't deliver sustained value. Integration isn't a nice-to-have. It's what separates AI experiments from AI at scale.

#### The near-term pressure


Expecting 30%+ workload increases within a year



What this means: Pacesetters expect massive workload increases and soon—but they've already built network infrastructure that can handle it (71% fully adaptable). Global companies face similar pressures but with networks that aren't ready.

#### The long-term confidence

Infrastructure built for AI offerings 3-5 years out



What this means: Pacesetters aren't just ready for today's AI. They've architected network infrastructure for AI that hasn't been deployed yet. That's not incrementalism—that's architectural foresight.

#### Key insight



The constraint: Network capacity and integration must be established before workloads scale. Adding compute is fast. Rewiring data center networks under production load is not.

The bet: Pacesetters built network as the foundation layer—with both eyes on the future. Global companies are discovering network bottlenecks while already trying to scale.

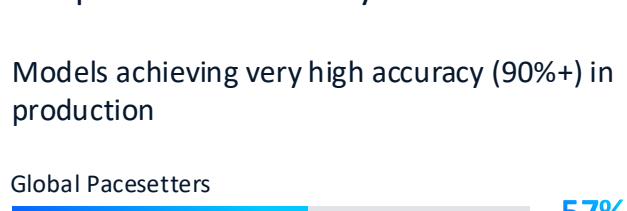
### ⚙️ Operations

#### Bet 3: Obsess over continuous optimization from day one


Getting AI models into production is worth celebrating. But here's the reality: model performance can degrade. The winners aren't those who deploy fastest—it's those who can optimize fastest *after* deployment.

#### The performance reality

Models achieving very high accuracy (90%+) in production



Both experience significant performance degradation over time



What this means: While accuracy targets vary by use case, Pacesetters are 3.5x more likely to report their production models meeting their accuracy goals (90%+), even though both groups experience model degradation at similar rates. The difference isn't the models—it's what happens after deployment.

#### How Pacesetters deal with the degradation dilemma

	Global Pacesetters	Global Avg
Operations Capability for AI Deployments		
Continuous monitoring with automated alerts & scheduled model retraining	72%	24%
Fully automated deployment and model quality tracking	78%	24%
Fully integrated with model registries or MLOps platforms	59%	22%
Fully integrated with monitoring and observability tools	60%	25%


What this means: Pacesetters have built the invisible infrastructure for continuous optimization—monitoring systems, retraining pipelines, deployment automation. These systems enable them to respond to model degradation 3x faster than the average Global company.

## The infrastructure investment pattern


Notice the pattern: Global Pacesetters invest 3x more in infrastructure nobody sees—monitoring, tracking, registries, observability tools—than the average Global company. This invisible infrastructure is what makes their visible results (57% high accuracy vs 15%) possible.

### Update velocity as competitive advantage

Can update models in under 1 hour with minimal downtime



Can update with near-instant, zero downtime



What this means: When models start degrading (which happens to everyone), Pacesetters can retrain and push updates in under an hour. Most Global companies need 1-4 hours with planned downtime.

The compound effect: Over a year, that's 50+ optimization cycles for Pacesetters vs. 12-15 for Global companies. More cycles = better sustained accuracy = higher ROI from each deployed model.

### Key insight

The constraint: Model degradation is inevitable. Your competitive advantage isn't avoiding it—it's responding faster than competitors.


The bet: Pacesetters build monitoring, CI/CD, and observability infrastructure that nobody sees. That invisible infrastructure makes their visible results possible. Deployment is day one, not the finish line.

## Bet 4: Infuse security into infrastructure


Many are aware that AI comes with distinct security risks. But there's a massive gap between knowing security matters and actually building infrastructure to deliver it. And that gap becomes even more critical as you move to deploying autonomous AI agents.

### Awareness vs. Implementation

High awareness of cybersecurity threats specific to AI and ML systems




End-to-end encryption with continuous monitoring and instant threat response




What this means: Security matters. But only 30% of Global companies have actually built the infrastructure to deliver it. Awareness and architecture are still dangerously low for Global companies.

### The AI Agent security exposure

Deploying AI agents (autonomous systems that act independently)



Fully equipped to control and secure agentic AI systems



What this means: Most Global companies are deploying autonomous AI agents that can make decisions and take actions. But only 31% have the infrastructure to properly secure them. That's a 52-point exposure: deploying automation faster than the ability to govern it. For Pacesetters, it's 96% deploying and 75% securing—still a 21-point gap, but much more manageable when security is infused into the infrastructure.

### Security implementation details

Security Capability	Global Pacesetters	Global Avg
Detect/prevent unauthorized tampering	81%	29%
Dynamic access controls with automatic threat-based adjustments	74%	27%
Security and identity systems fully integrated	62%	29%
Governance systems fully integrated	62%	26%


What this means: Detecting tampering means nothing if you can't dynamically adjust access in response. Governance fails without integrated identity. Security capabilities only work when they work together—as an integrated system fused into the infrastructure, not bolted on as isolated tools.

### Same challenge. Different outcome.

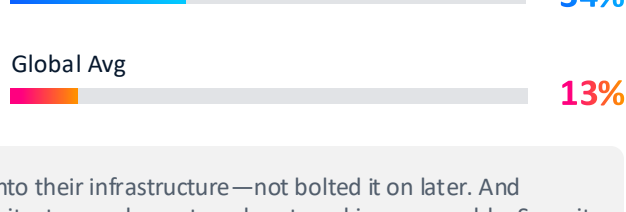
When asked about infrastructure challenges, 'cybersecurity risks posed by AI workloads' ranks as a top concern for both groups. The difference: Pacesetters cite it as a challenge they've invested heavily to address. Global cites it as a challenge most are still figuring out how to solve.

### Security as velocity enabler

Deployed AI at the scale and speed necessary to deliver value



Fully trust third-party AI models to handle security



What this means: Pacesetters deploy faster because they infused security into their infrastructure—not bolted it on later. And they're more likely to trust third-party AI models because their security architecture makes external partnerships governable. Security built into infrastructure accelerates innovation. Security bolted on creates friction.

### Key insight

The exposure: 83% of global companies are deploying autonomous AI agents. Only 31% can properly secure them. That's 52 points of ungoverned automation.

The bet: Pacesetters build security into their infrastructure from day one—as foundation, not as gates. This enables them to deploy faster and more securely than others.

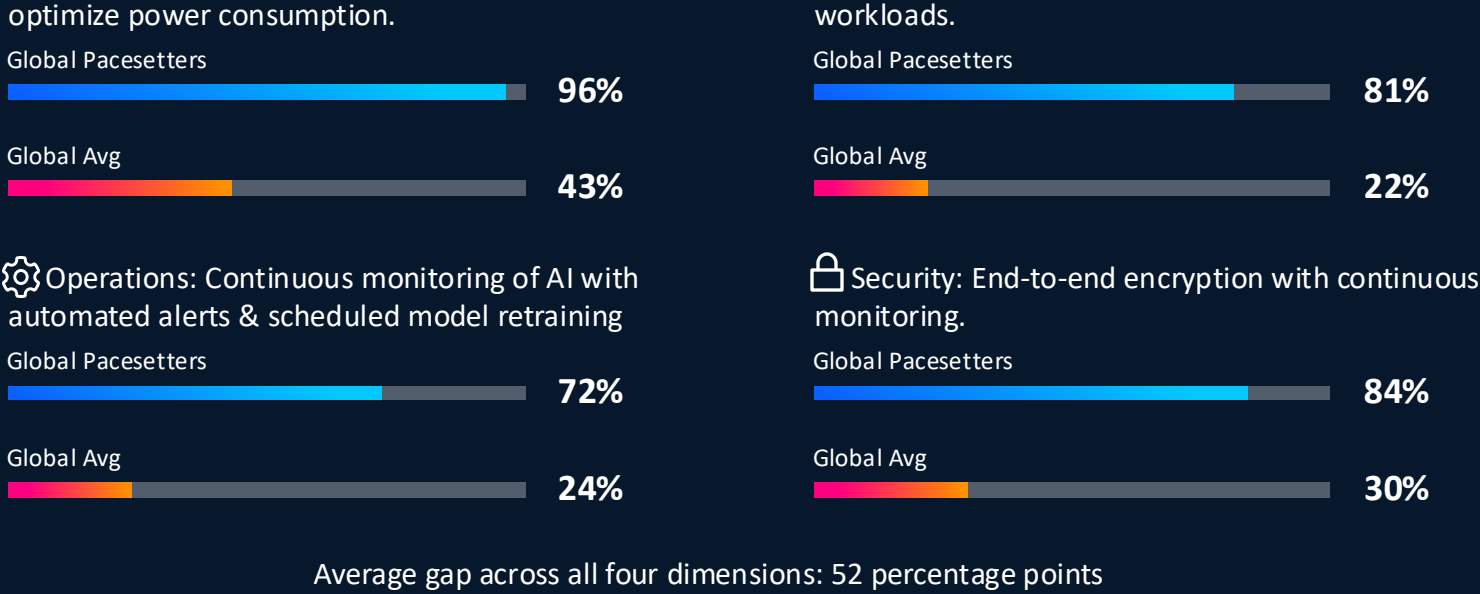
## Four bets. Compounding advantage.

### The strategic foundation that enables these bets

Strategic Element	Global Pacesetters	Global Avg
Have well-defined AI strategies	99%	58%
Have clear leadership/ownership of AI strategy	94%	81%
Building new infrastructure within 12 months	77%	43%

What this means: Despite facing the same challenges—budget constraints, procurement delays, board hesitation—Pacesetters move faster. Why? Because strategic clarity and ownership give them the ability to make different architectural choices. They know what to build and who's accountable for building it.

### The four bets in summary

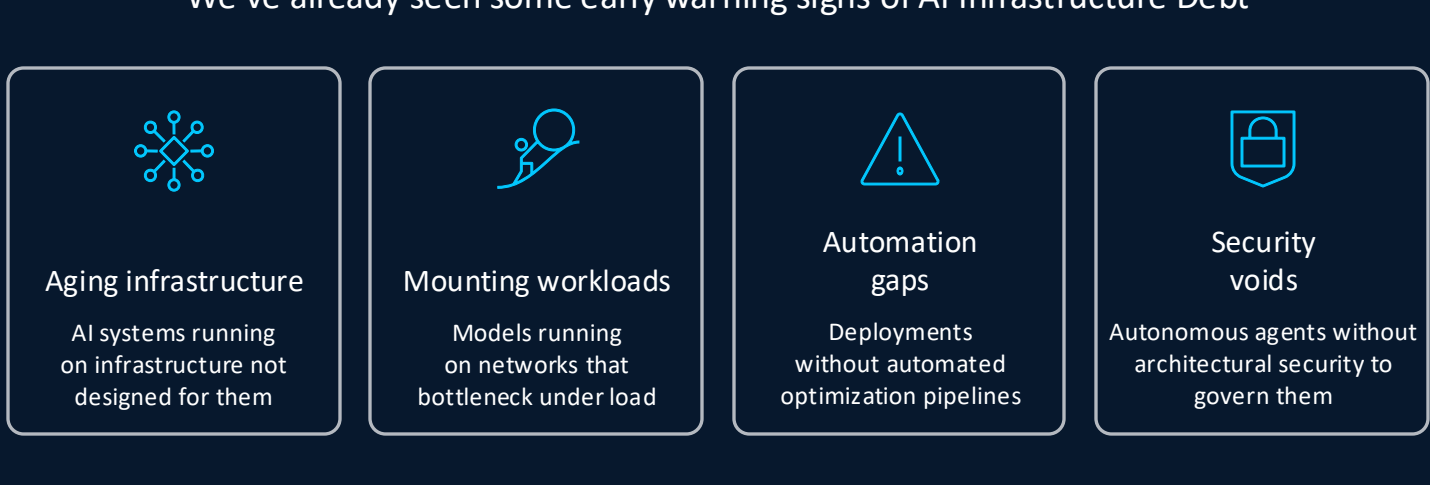


### Architectural bet or AI Infrastructure debt?

The difference between Pacesetters and the average company comes down to a choice: Make architectural bets upfront, or accumulate AI Infrastructure Debt over time.

AI Infrastructure Debt is the technical debt of the AI era—the shortcuts, gaps, and compromises that can quietly accumulate as companies rush to deploy AI without the right foundation. Left unchecked it could become the silent bottleneck to value: slowing innovation, inflating costs, and eroding returns.

### We've already seen some early warning signs of AI Infrastructure Debt



This debt can compound with every deployment. And infrastructure becomes much harder to refactor once you're running AI at scale.

### Which Pacesetter architectural bet are you ready to make?



### The bottom line

The architectural bets you make on your infrastructure today determine the AI value you can realize tomorrow. Pacesetters are making these bets. They're deploying at scale and speed while others struggle.

**Are you ready to build AI infrastructure like a Pacesetter?**

### Take the Infrastructure Readiness test

[Take Test](#)

[Explore](#) the Cisco AI Readiness Index 2025

[Back to Home Page](#)

About the Cisco AI Readiness Index

The Cisco AI Readiness Index 2025 is a global study, now in its third year, based on a double-blind survey of 8,000+ senior IT and business leaders responsible for AI strategy at organizations with over 500 employees across 30 markets and 26 industries.

© 2025 Cisco and/or its affiliates. All rights reserved.