


Special Edition

Private Cloud Outlook 2026

From Modernization to AI Readiness:

The New Platform Imperative



Contents

| | |
|--|-----------|
| Introduction | 3 |
| AI Is the Number One Modernization Priority—and Most Enterprises Are Behind | 4 |
| Modern Private Cloud Meets the AI Moment | 5 |
| Platform Engineering Teams Become Standard | 7 |
| Platform and Infrastructure Teams Converge for AI | 8 |
| Recommendations for IT | 9 |
| Conclusion | 11 |

Introduction


The [Private Cloud Outlook 2026](#) report, based on a global survey of 1,800 IT leaders worldwide, examines how enterprises are rethinking cloud strategy as AI moves from experimentation to production. This addendum looks deeper into several related findings from the same research:



Why application modernization remains unfinished for many enterprises



How private cloud has matured for AI-era workloads



Why platform engineering is becoming a core IT function

What emerges is a practical roadmap for AI readiness: modernize what exists, strengthen the private cloud platform, and align the teams responsible for delivering it.



AI Is the Number One Modernization Priority—and Most Enterprises Are Behind

Enterprise application modernization has always been difficult. AI is making it urgent.

The Private Cloud Outlook 2026 data shows that most organizations are still early in their modernization journey: 72% of enterprise organizations have modernized less than half of their applications.

What's standing in the way? According to the IT decision-makers who participated in the survey, their top three application modernization challenges are:

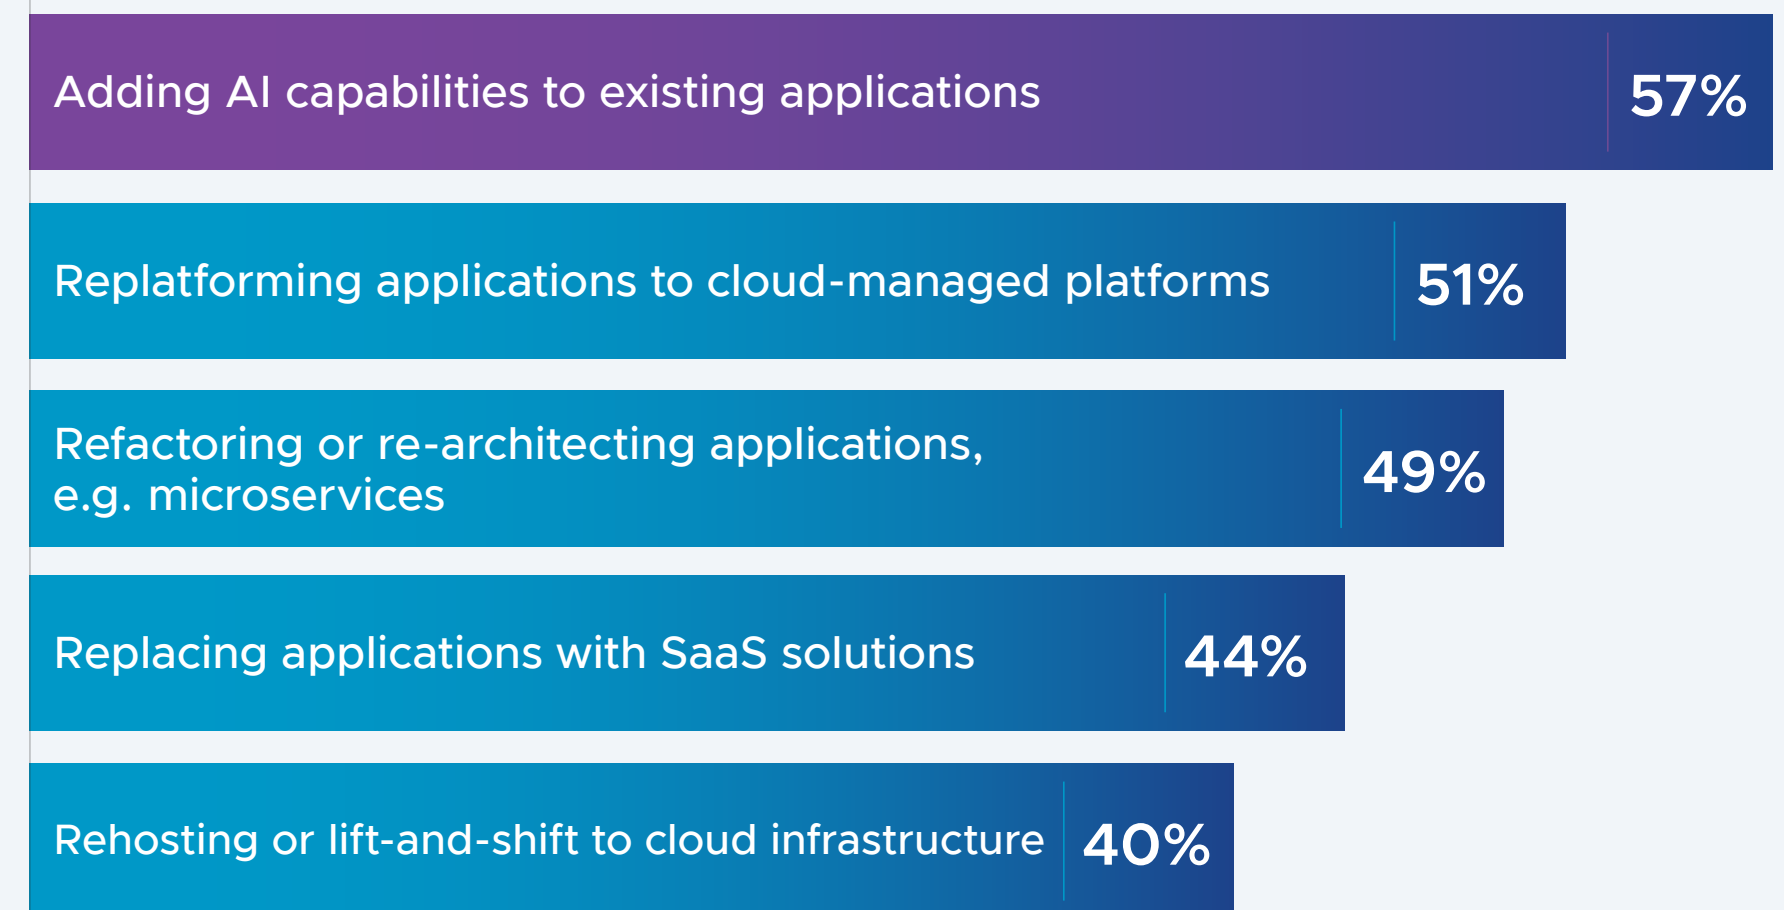
- **Security and compliance requirements, cited by 46% of respondents**
- **Integration complexity, cited by 42%**
- **Skills gaps, cited by 41%**

These issues become more consequential as enterprises try to extend AI across applications that were not originally designed for it. Adding AI capabilities to existing applications is the top modernization strategy for 57% of enterprise IT organizations—not re-architecting applications, replacing them with SaaS, or lifting and shifting them. They want to enhance what they already have with AI.

Many of these organizations are being asked to accelerate toward AI while still carrying a substantial modernization backlog. And they're not starting from a clean slate. They are working with complex, mixed estates that include traditional applications, modern applications, VMs, and containers running across multiple platforms.

Adopting a cloud platform that supports both traditional and modern workloads together allows organizations to modernize incrementally at their own pace.

MOST CRITICAL MODERNIZATION PRIORITY



Q: Application modernization approaches used or pursued; n=1800

Modern Private Cloud Meets the AI Moment

For years, public cloud was seen as the answer to application modernization. It offered the tooling, automation, and self-service capabilities that many private environments lacked. For organizations under pressure to move faster, public cloud often appeared to be the only viable model.

Today the gap between public and private cloud has closed. The capabilities once associated with public cloud are now available in modern private cloud environments.

Enterprise IT leaders are confirming private cloud's advantages in practice. Ninety-three percent agree it delivers the reliability business-critical applications demand, and 92% say it provides the financial transparency and predictable costs needed to govern AI infrastructure spend. For organizations navigating data residency and sovereignty requirements, four out of five (84%) report that private cloud mostly or fully enables deployment in those constrained environments.

THE GREATER MAJORITY OF ENTERPRISE IT TEAMS ARE RUNNING MATURE PRIVATE CLOUDS



Q. Capabilities currently delivered through private cloud; n=1800

Private cloud has matured across capabilities that are essential for modern workloads such as AI.

- **A single platform for all workloads:**

Containers, VMs, and AI workloads can run together with consistent policy and operations. This gives enterprises a stronger foundation for platform teams and reduces the fragmentation created by separate operating models.

- **Automated policy enforcement:**

Security guardrails can be built into the platform, rather than added after deployment. This becomes increasingly important as AI touches more sensitive applications, data, and workflows.

- **Self-service provisioning:**

Developers can access the resources they need without waiting on manual infrastructure processes. For AI-enabled application modernization, that speed matters.

- **Cost visibility and predictability:**

Embedded FinOps and cost control capabilities provide the cost transparency to plan, optimize, and govern AI infrastructure spend.

- **Multi-persona services:**

Developers, security teams, FinOps leaders, operators, and ML engineers can work through role-specific dashboards and experiences, while still operating from a common platform.

- **Centralized management:**

AI workloads interact with databases, CPUs, GPUs, storage, networking, and existing applications. A mature private cloud platform gives enterprises a way to manage this complexity in one place. It brings AI closer to the applications and data it depends on, while giving IT teams the control needed to meet new requirements for data governance, security, and sovereignty.

ADVANTAGES OF PRIVATE CLOUD



93%

agree private cloud delivers the reliability business-critical apps demand

Q: Agree with following statements about private cloud; n=1800



4 out of 5

84% state their private cloud mostly or fully enables deployment in sovereign or data residency constrained environments

Q: Private cloud enables the named capabilities; n=1800



92%

agree private cloud provides financial transparency and predictable costs

Q: Agree with following statements about private cloud; n=1800

Platform Engineering Teams Become Standard

In addition to changing infrastructure requirements, AI is affecting how IT teams are organized. IT organizations now have to manage AI workloads alongside traditional applications, across hybrid environments, with governance, security, and cost controls in place. That level of complexity requires a dedicated team to create common services with consistent policies across the enterprise.

Today, 80% of enterprises have a platform engineering team responsible for building and operating internal platforms. Fifteen percent plan to in the next 18 months and 5% have no plans.


The organizations furthest along in platform engineering maturity are better positioned to operationalize AI at scale. They already have the structures

needed to turn infrastructure capabilities into consumable services for developers, security teams, FinOps teams, and ML engineers.

But the role of platform engineering is also expanding. The first wave of platform engineering was largely built around developer productivity: self-service access, standardized pipelines, reusable templates, and faster delivery. The next wave will be broader. It will be built for AI-native operations and will need to support AI-native infrastructure, role-specific services, embedded FinOps, platform security for AI, and composable architectures that allow enterprises to adapt as AI requirements evolve.

That makes platform engineering a core part of AI readiness. It gives enterprises a way to translate modern infrastructure capabilities into services that teams can use—and govern—at scale.

80% HAVE A PLATFORM ENGINEERING TEAM TODAY



Q: Do you have a dedicated platform engineering team; n=1800

Platform and Infrastructure Teams Converge for AI

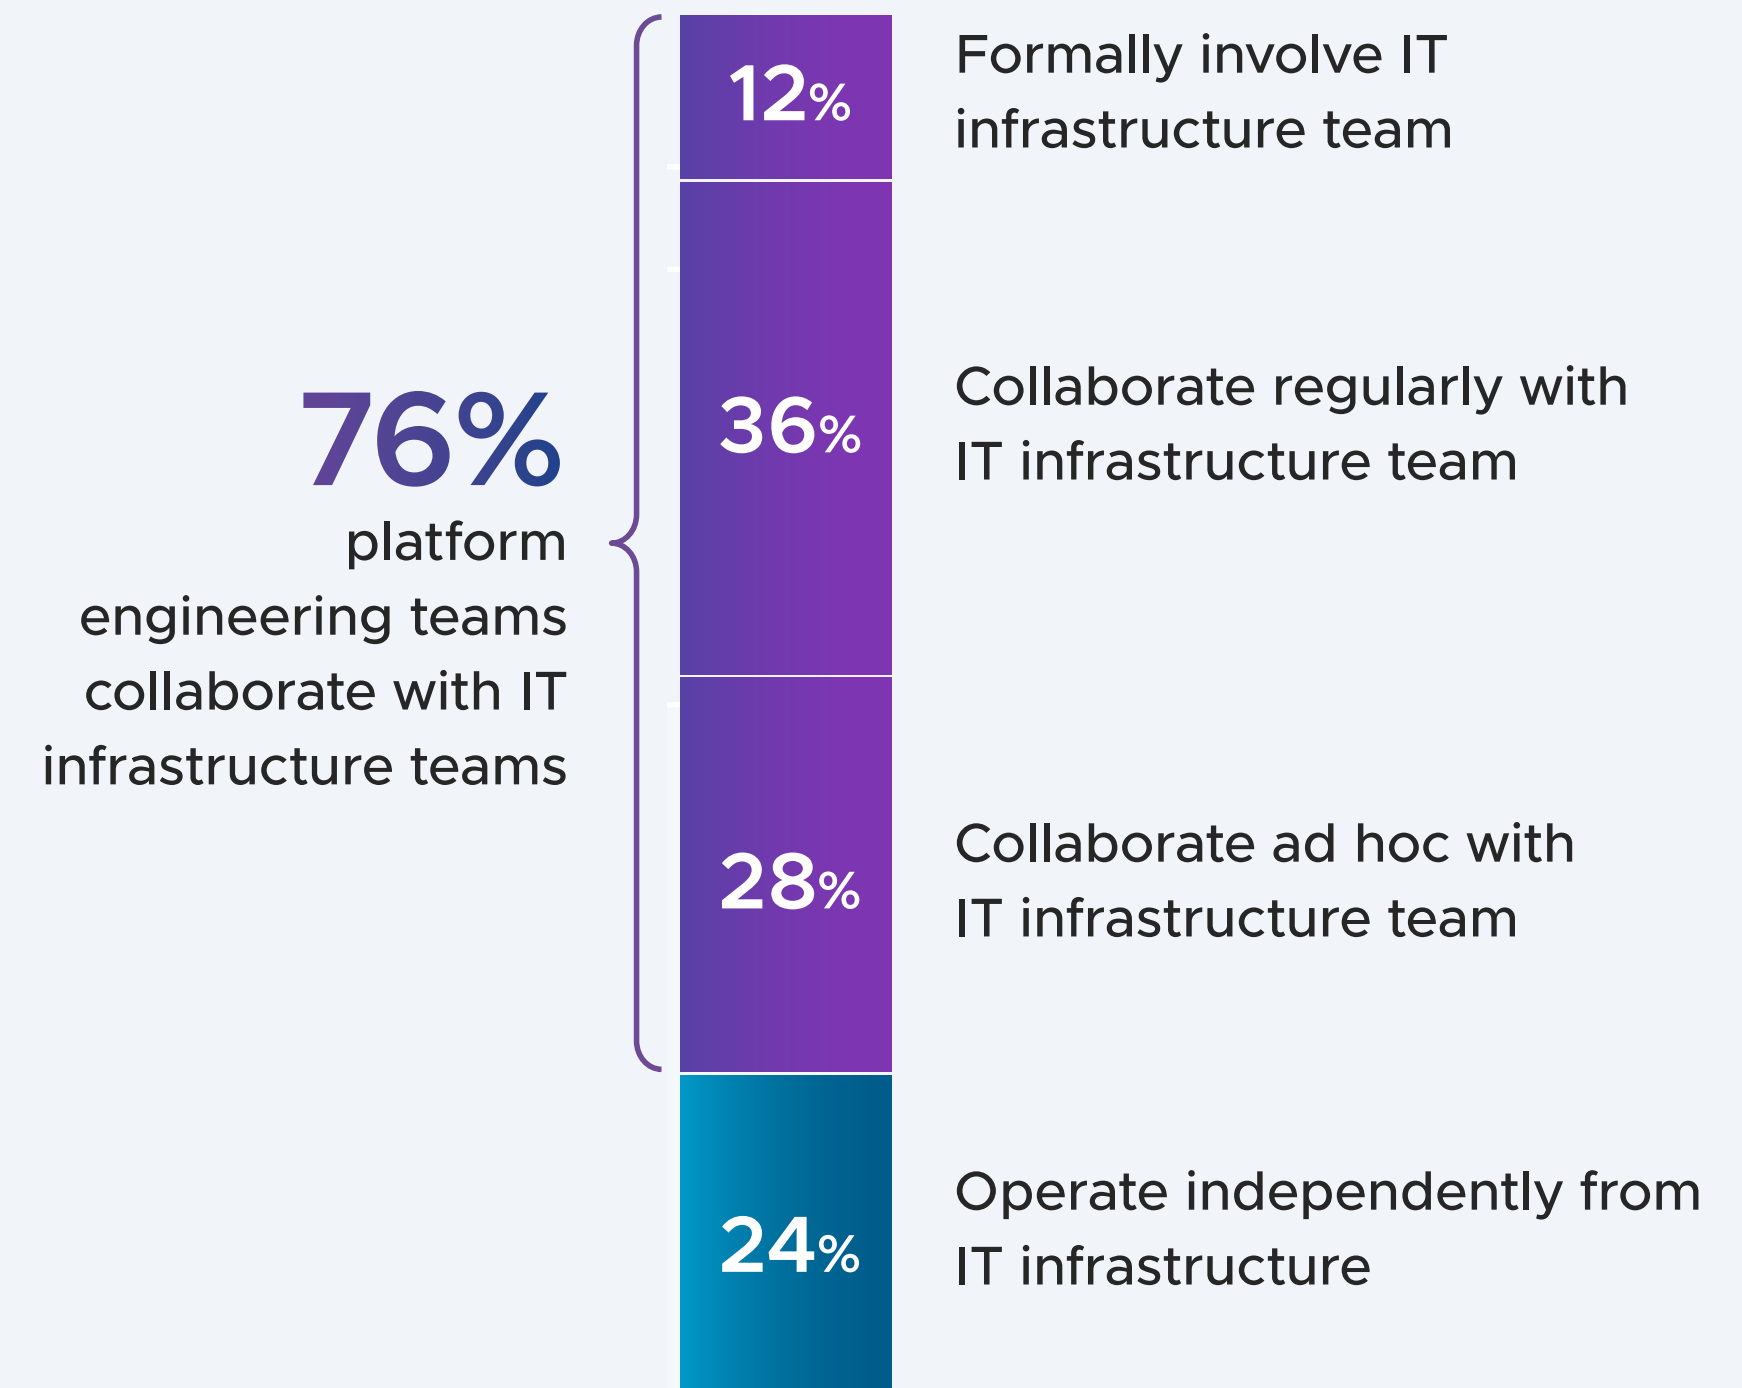
With platform engineering now becoming a standard enterprise function, the next stage is closer alignment with IT infrastructure. That alignment is already taking shape: only 24% of platform engineering teams operate completely independently from IT infrastructure teams, meaning most organizations have already begun to connect the two functions, from ad hoc engagement to formal integration.

Neither team can fully meet AI requirements alone. Platform teams need infrastructure teams to make AI services reliable, governed, performant, and cost-aware. Infrastructure teams need platform teams to make those capabilities consumable across the broader organization.

This reflects the multi-stakeholder reality of modern cloud platforms. Developers, security teams, FinOps teams, and ML engineers all need different services, dashboards, controls, and workflows from the same underlying platform. Without close coordination between platform engineering and IT infrastructure, those needs can easily create new silos.

The organizations still running platform engineering independently from infrastructure will eventually have to close that gap. As AI workloads scale, the connection between teams will only become more important. The enterprises already building that collaboration will be better positioned to move quickly and with less friction.

MOST PLATFORM ENGINEERING TEAMS COLLABORATE WITH IT INFRASTRUCTURE TEAMS



Q: relationship between platform engineering and IT infrastructure ops; n=1456

Recommendations for IT

The modernization backlog, the maturity of private cloud, and the rise of platform engineering are all connected. IT leaders should approach them as part of the same AI readiness agenda.

For Platform Engineering Leaders

The shift from sequential to concurrent modernization is overdue. The model of modernizing first and adding AI second does not fit today's operational reality. Enterprises need a platform that supports both traditional and AI-enhanced workloads in parallel.

Security and compliance guardrails belong at the platform layer. Embedding them into the infrastructure foundation—rather than appending them just at the application level—addresses security and compliance requirements, the top barrier to application modernization cited by IT leaders.

Collaboration with IT Infrastructure should be formalized now, not deferred as a future initiative. Seventy-five percent of organizations describe some form of active collaboration between platform and infrastructure teams, but only 12% have made that relationship formal.

This gap requires structured capability development, not individual hiring. Virtualization administrators already hold the foundational knowledge needed to operate a full cloud operating model.

For CIOs

Infrastructure investment strategy needs to reflect where production workloads are heading. Organizations that plan to grow private cloud are making that shift now, driven by cost, governance, and AI performance requirements. Rebalancing toward private cloud is no longer a future consideration.

Private cloud maturity should be assessed against AI-readiness criteria: unified workload support, automated compliance guardrails, self-service provisioning, showback and

chargeback, platform team structure, and persona-based services. These maturity dimensions determine whether private cloud can function as a production-ready operating environment for AI or whether fragmentation and manual processes will continue to slow progress.

Governance for AI workload placement requires board-level visibility. Eighty-six percent of enterprises say geopolitical and regulatory factors are reshaping their infrastructure strategy. Placement decisions made without formal policy carry compounding risk.

AI infrastructure is a platform decision, not a procurement line item. Organizations that route AI through a unified private cloud platform with cost attribution, security controls, and governance built in will close the cost and compliance gaps that currently concern the majority of enterprise IT leaders.

For Infrastructure VPs

AI-specific capabilities belong inside the private cloud operating model, not on a separate infrastructure track. GPU provisioning, data governance, model inferencing infrastructure, security controls, and cost attribution all need to be managed as part of the enterprise platform.

Joint operating models with Platform Engineering teams should be a near-term priority. Formalizing this relationship—through shared governance, coordinated workload placement policies, and unified tooling—creates the operational foundation that makes platform maturity for AI readiness achievable at scale.

Application modernization is not a project with a defined endpoint. Seventy-two percent of enterprises have modernized less than half their application portfolio. A platform approach that supports both legacy and modern workloads in parallel—rather than requiring sequential transformation—is the model that keeps pace with the current rate of change. The Kubernetes and cloud-native skills gap is

the most direct operational constraint on readiness. At 37% globally, it is the single largest people-related challenge organizations report. Addressing it through structured upskilling of existing infrastructure talent, supported by a mature private cloud operating model, delivers faster results than building net-new hiring pipelines.

VMware Cloud Foundation: Built for This Operating Model

VMware Cloud Foundation is designed to support these recommendations across all leadership roles. Its unified platform supports VMs, containers, and AI workloads under consistent policy and operations, which helps organizations modernize existing applications while adding AI capabilities in parallel. Automated policy-based guardrails address security and compliance requirements at the platform layer, while self-service provisioning enables application teams to move without requiring IT intervention on every deployment.

VCF also delivers the operating model required for mature private cloud at scale. Infrastructure

automation, policy-based governance, built-in FinOps capabilities, and self-service delivery across workloads address the operational dimensions that define AI-ready private cloud. Persona-based service delivery capabilities—offering role-specific dashboards, experiences, and APIs for developers, security teams, FinOps teams, and ML engineers—enable platform teams to serve the broader enterprise from a single operational foundation.

By consolidating compute, storage, networking, and security under a single management plane, VCF also reduces the technology fragmentation that keeps platform and infrastructure teams operating separately. Its automation and infrastructure-as-code capabilities support the shift toward repeatable, scalable operations, helping organizations close skills gaps while making AI-ready infrastructure easier to consume, govern, and operate across teams.

Conclusion

Enterprise AI will be built on the application estates organizations already run: partly modernized, deeply integrated, and increasingly complex. That makes AI readiness an infrastructure and platform challenge as much as an application challenge.

The data shows that private cloud has matured for this moment. At the same time, platform engineering is becoming standard, giving IT organizations a practical operating model for delivering those capabilities at scale.

That is the platform imperative: enterprises need infrastructure that can support AI, and a mature private cloud operating model that can deliver its benefits consistently across teams.



