November 2025

From pyramid to diamond: The future of IT operating models in the AI era

What the end of the offshore pyramid means for IT leaders – and how to build an AI-native operating model that lasts

Contents

The evolution of IT operations has always mirrored the evolution of technology itself, each wave driving new efficiencies, new expectations and new possibilities for how we deliver value. Today, we stand at the beginning of another major shift.

Agentic AI is redefining what it means to operate at scale. Where the last decade was about becoming cloud-native, the next will be about becoming AI-native — embedding intelligence into every layer of the stack so that systems, processes and teams can act with autonomy and purpose.

This isn't just about automating tasks; it's about rethinking operating models, talent structures and the very economics of service delivery. The traditional offshore pyramid, which was built on human volume and manual process" is giving way to a "diamond model" built on expertise, intelligent platforms and digital tools that enhance human capability rather than replace it.

At NTT DATA, we see this transformation happening every day. With our Agentic Factory, modern delivery platforms and global communities of practice, we're helping clients design operating models that are adaptive, intelligent and sustainable. The organizations that succeed in this new era will be those that blend technology with culture, combining innovation, governance and reskilling to create value that compounds over time.

This guide outlines what that future looks like. It's a practical roadmap for leaders who are ready to evolve from cloud-native to AI-native, from silos to product-aligned teams and from isolated automation to platform-scale intelligence.

The future of operations isn't coming — it's already here. The question is how quickly we can adapt our models, our teams, and our mindset to match it.

“This isn't just about automating tasks; it's about rethinking operating models, talent structures and the very economics of service delivery.”

Best regards,

Simon Ruelens
Managing Director, Cloud Presales,
NTT DATA

AI is turning the traditional IT outsourcing model on its head

For 20 years, IT outsourcing relied on a pyramid structure. At the wide base, teams of Level 1 engineers at offshore delivery hubs handled repetitive but important tasks such as monitoring system alerts or resetting passwords. Above them, thinner layers of Level 2 and Level 3 experts managed more complex issues. Their customers benefited from lower costs because of global wage differences, but often at the expense of intimacy and agility.

Agentic AI is upending this structure. Intelligent agents are already handling more mundane tasks and will increasingly move into operational areas, taking on much of the monitoring and triage that once consumed entire teams.

This is eroding the base of the pyramid to turn it into a diamond — fewer people overall, but more senior, multiskilled engineers in the middle, supported by a team of digital colleagues, instead of juniors.

Agentic AI also goes beyond automating tasks and forces us to rethink the processes we've built around these tasks. Information Technology Infrastructure Library (ITIL) flows, ticket handoffs, queues and escalations — these were all designed for a human-enabled pyramid.

Although ITIL won't disappear overnight, it will need to evolve into a lighter-touch model that supports a more autonomous, AI-driven way of working.

From cloud-native to AI-native

The first major reinvention came with cloud-native architectures — a shift in how we built, deployed and scaled software. Microservices, containerization and elasticity made IT more agile at scale and broke the dependence on monolithic design. But cloud-native alone won't take us far enough. The next reinvention is AI-native — applications that have intelligent agents woven into their fabric from the start, not bolted on afterwards.

Product-aligned teams are uniquely positioned to make this shift. They have the business knowledge, the technical breadth and the end-to-end ownership needed to design, run and continuously evolve applications that are both cloud-native and AI-native (where intelligent agents are part of the architecture).

This reimagination doesn't stop at the application layer. Product teams and operations squads will collaborate with the same intelligent agents, embedded in every layer of delivery. Agentic AI is not confined to the business or infrastructure layers — it's the connective tissue that links them.

“Product-aligned teams are uniquely positioned to make this shift.”

The rise of product-aligned teams

The way IT is organized is also changing. Traditional technology silos — networks, applications and computing — have always made sense to IT but never to the business. Business leaders want outcomes, not silos. They want teams aligned to products, services and value streams.

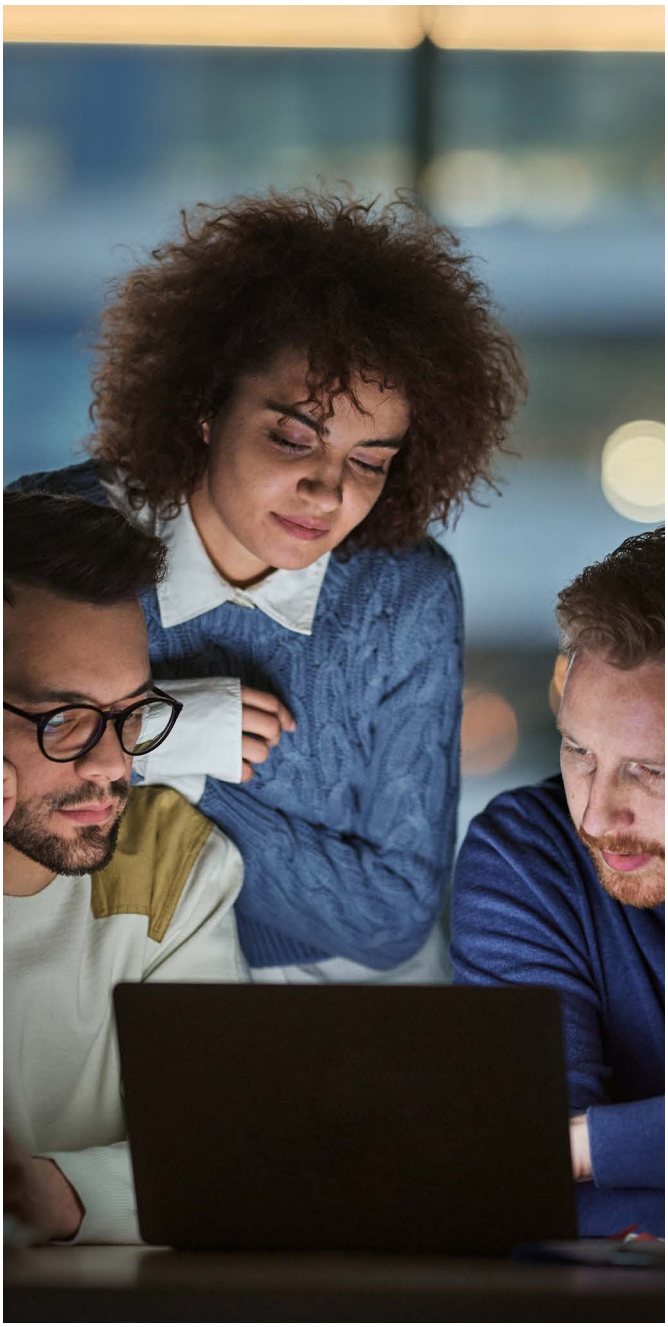

New operating models reflect this shift by introducing full-stack teams that own everything from infrastructure to the business service itself. Within these teams, engineers move easily between day-to-day operations and transformation work, making IT less about keeping the lights on and more about moving the business forward.

But there is a risk, because freedom without guardrails creates fragmentation. This is where centers of excellence (CoEs) play a critical role in providing the consistency that product teams need: standardized tool chains, observability frameworks, security practices, AI platforms and financial governance. CoEs ensure that when teams innovate, they do so in a way that's safe, compliant and scalable across the organization.

Additionally, the focus has shifted from application modernization to application reimagination. For years, the challenge was to lift, shift and refactor workloads for the cloud. In the age of agentic AI, this is no longer enough. Organizations need to reimagine their applications from first principles — not by asking **"How do we migrate this?" but "If we were to rebuild this with AI-native capabilities from the start, what would it look like?"**

AI agents as colleagues, not replacements

It's tempting to see AI as another form of automation. But the real shift is cultural. Intelligent agents will work in engineering teams as specialized digital colleagues with contextual understanding. They have agency and act autonomously to identify and resolve issues.

Picture an operations squad where some members are human and others are digital. The digital colleagues handle the boring but essential work. The humans oversee the whole operation, solve cross-domain problems, engage directly with customers and step in when judgment is required.

Engineers can therefore offload mundane, low-value tasks to their digital colleagues while focusing on high-value tasks that offer them opportunities to learn and grow.

Intelligent agents won't just transform infrastructure operations

These agents are beginning to move higher into the stack. They are supporting continuous integration and delivery (CI/CD) pipelines, automating deployments and testing, optimizing cloud spend through FinOps, and processing observability data at scale. Coding is also changing rapidly, as intelligent agents can generate, test and refactor software alongside human developers.

This changes the IT skills profile completely. We won't need thousands of junior engineers spending years at Level 1 before progressing. Instead, we need critical thinkers and multiskilled problem solvers — engineers who coordinate their digital colleagues rather than execute repetitive tasks themselves.

In this context, reskilling is essential. While we don't have all the answers yet, we're rethinking how people enter IT careers, what skills they will need and how to keep these skills relevant across technologies and customers.

Change management is just as important as technical training, because without the right culture and mindset in the workforce, even the best AI will fail to deliver.

“While we don't have all the answers yet, we're rethinking how people enter IT careers, what skills they will need and how to keep these skills relevant across technologies and customers.”

The economics of agents: Why platforms matter

Intelligent agents aren't free. They require careful design, foundational tools to run on and ongoing training and oversight. Without scale, they don't make sense economically.

This is why platforms matter. When monitoring, incident management, observability and deployment pipelines are standardized across environments, an intelligent agent can be built once and used everywhere. It allows both service providers and their customers to achieve exponential productivity gains.

Custom stacks will always exist. Some organizations will need bespoke intelligent agents. But those who buy into shared platforms will benefit from a growing catalog of reusable agents that improve over time. It's the only way to unlock the full potential of agentic AI.

The Agentic Factory: Industrializing the pipeline

NTT DATA has invested in a dedicated Agentic Factory that designs, tests and deploys intelligent agents at pace. Some of these agents automate infrastructure operations, while others are emerging higher up the stack in CI/CD pipelines, testing, FinOps, observability and coding. Not every use case is fully mature, but the trajectory is clear.

The factory lets us build agents in an iterative way by experimenting, tuning and reusing patterns across environments. We know that agents need continuous training, governance and guardrails. In this instance, industrialization doesn't mean perfection on day one. Rather, it means repeatability, safety and scale.

Balancing control and innovation

As we introduce more automation and agents into the model, an obvious concern arises for our clients: What remains under their control, and what is entrusted to the provider? In other words, how do we avoid creating black boxes?

The way forward is clear. Clients retain ownership where business context is critical — service management, observability of key applications and DevOps pipelines. Service providers take on the commodity layer, providing patching, configuration and AI operations where scale and innovation matter more than local nuance.

The good news is that this balance doesn't have to mean vendor lock-in. Modern platforms integrate with each other, application programming interfaces are everywhere and tooling is far more portable than in the past. Lock-in remains a risk, of course, but the barriers are lower than ever. The real value is not in the individual tools but in how the operating model stitches them together, governs them and runs them at scale.

Squads, guilds and trains: Delivery with flexibility

If the pyramid is no more, what replaces it? Our answer is squads, guilds and trains.

Squads are small, agile technical teams embedded within your operations. Guilds are the global communities that keep these squads fresh, sharing standards, training and innovation across accounts. Trains are the delivery structures that bring squads together into a coherent, client-facing whole.

Some trains are fully dedicated, some are hybrid and others serve multiple clients. Service excellence, client satisfaction and shared accountability are the principles that bind them, with every part of the delivery model contributing to a unified, high-performing experience.

As automation accelerates productivity, even large infrastructure estates may no longer justify fully dedicated teams. A flexible model means clients can evolve their delivery approach without disruption, moving from dedicated to leveraged as needed.

Why the next operating model can't be built on old foundations

The temptation in tough times is always to go for the cheapest option. Many organizations rush back to the old offshore pyramid because it delivers immediate savings. But real value and lasting results come from people and intelligent agents working in tandem, combining judgment, speed and scale.

Agentic AI is reshaping IT operating models at a pace we've never seen before. Yet there's a risk that organizations facing financial uncertainty — due to tariffs, inflation and sector pressures — may be tempted to double down on offshore labor arbitrage. Sign those contracts today, and you may wake up tomorrow locked into an obsolete operating model. Your competitors will grow faster, innovate sooner and build more resilient operations with intelligent AI that learns fast, works smart and scales endlessly.

To succeed in this new world, redesign your operating models for tomorrow. Move from pyramids to diamonds, from silos to product teams, from cloud-native to AI-native and from application modernization to true application reimagination, and combine human judgment with machine intelligence.

> "Agentic AI is reshaping IT operating models at a pace we've never seen before. Yet there's a risk that organizations facing financial uncertainty — due to tariffs, inflation and sector pressures — may be tempted to double down on offshore labor arbitrage."


NTT DATA