

AI-enabled  
modernization

From cloud-native to intelligent, adaptive, and intent-driven applications:  
**Why AI-enabled modernization is now the strategic imperative**

AI is not only transforming business outcomes and services but also changing the way modern applications are built.

**F**rom major disruption to customers in banking, travel, and logistics, to reputational damage due to data leaks and lack of compliance, recent high-profile system failures have dramatized the fragility of legacy computer systems. Not to mention the downscaling of business due to major economic upheavals and the inability to keep pace with the speed of innovation from competition. Applications built using legacy technologies and development practices have proven ill-suited to adapt to the rapidly changing business needs and round-the-clock availability demands of today's businesses.

Aging, inflexible applications that are costly to maintain, difficult to integrate, and resistant to change are frustrating organizations' digital transformation initiatives. The risk of falling further behind is heightened by the recent emergence of agentic AI, which promises to redefine applications from static data processors to adaptive, autonomous platforms that interpret intent and take actions.

Legacy applications impede agility, impose costly maintenance overhead and scale poorly. IDC's 3Q24 Cloud Pulse Survey of 1,700 cloud buyers found that 60% believe their business's IT or digital infrastructure currently requires major transformation.

The same study, one quarter earlier, found that 49% describe their business as being predominantly digital.

This data exposes a paradox: Digital interactions and automation increasingly define the rules of business, but the majority of organizations aren't prepared to respond to these new demands.

The future of business is in agent-powered software orchestrated by sophisticated workflow engines. This modular, responsive design promotes flexibility by enabling seamless interoperability between people and AI agents.

Aging, inflexible applications that are **costly to maintain, difficult to integrate, and resistant to change** are frustrating organizations' digital transformation initiatives.

Although AI has been a mainstream business focus for barely three years, its adoption rate has been breathtaking.

- TheCUBE Research found that **87% of businesses have implemented or soon plan to implement AI.**
- Foundry reported that **49% of organizations now have dedicated budgets for AI projects,** with an average of 23% of total IT spending going to AI initiatives.
- Menlo Ventures says **enterprise AI spending has surged from \$1.7 billion in 2023 to a forecasted \$37 billion in 2026,** making it the fastest-growing software category in history.
- McKinsey's Global AI Survey found that **78% of respondents said their organizations use AI in at least one business function,** up from 55% in 2023.


Routine tasks will be processed automatically, and processes reimagined around the decision-making capabilities of agents acting in concert. Organizations that fail to modernize their applications to support these new demands risk falling further behind their more agile, AI-equipped competitors.

AI is already significantly impacting application modernization initiatives. Such efforts have long been grounded in shifting to cloud-native constructs such as microservices, software containers, APIs, and agile DevOps practices.

Mordor Intelligence predicts that the cloud-native computing market will nearly quadruple to over \$41 billion over the next five years and will come to dominate software development in the long term.

Modernization is no longer the primary goal; it is now a necessary foundation

**Figure 1 | Cloud buyers who believe their business's IT or digital infrastructure currently requires major transformation.**



for intent-aware applications. However, “true modernization requires rethinking architecture for cloud-native capabilities,” says Rohan Braganza, Software & Data Engineering Leader at NTT DATA.

While cloud-native elements are still foundational, their new role is more as modernization enablers than final objectives. Rewriting existing applications in modern languages on cloud-native platforms no longer goes far enough. They must be transformed into intelligent, intent-aware, agentic systems capable of reasoning, adapting, and operating autonomously across a variety of platforms.

Skills must be upgraded as well. Businesses will increasingly rely on AI for decision support and operations, making AI literacy a vital core competency. Resistance to AI initiatives can be minimized if people who worry about threats to their jobs or responsibilities learn to rely upon agents as colleagues rather than rivals.

## Paradigm shifts

Modernization need not be an overwhelming task. Rohan recommends an incremental approach, using the “7 Rs” principle

**True modernization  
requires rethinking  
architecture for  
cloud-native  
capabilities.**

– Rohan Braganza,  
Software & Data Engineering  
Leader at NTT DATA.

(retire, retain, rehost, replatform, refactor, repurchase, rearchitect), tailoring strategies for each. Start with smaller applications that yield fast, measurable payoffs when modernized, and move up the complexity curve as you gain experience.

“Incremental modernization reduces risk, enables learning, and avoids operational disruption,” he says.

Generative AI and modern engineering tools can infuse intelligence into every layer of the application development and modernization life cycle. They can offload much of the coding drudgery, check for security flaws, and even design test beds that ensure applications perform reliably in production.

New approaches to modernization should embody three paradigm shifts.



## Shift #1: Intent-centric, agentic applications

Traditional enterprise applications were built to repeatedly process a known set of data inputs and return predictable results. They are based on well-defined rules and exception-handling processes. The chief objectives are to ensure consistency and reliability at scale. Functionality is limited and is identical for all users. Data must conform to a rigid set of formats and schemas.

The result is software that is large, inflexible, and difficult to maintain. Monolithic applications are simpler to manage because all components are contained within a single codebase, but adapting them to different use cases

is risky and expensive. Complexity can lock organizations into aging software for decades. By most estimates, over 60% of enterprises still rely on legacy applications for core operations.

The new breed of intent-centric applications interprets objectives, reasons over context, and orchestrates workflows autonomously. In contrast to rule-based or workflow-driven systems, intent-centric applications understand business goals and can identify ways to achieve them. This knowledge drives coordinated, autonomous action across the enterprise, with embedded AI agents streamlining processes, ensuring compliance, delivering personalized experiences, and automating decisions to improve efficiency and reliability.

Workflow engines orchestrate AI agents to optimize operations and accelerate throughput, enabling teams to handle more requests efficiently—without adding extra resources.

# +60%

of enterprises still rely on legacy applications for core operations.

**AI-enabled  
modernization**

A simple example of an intent-centric application is a personalized shopping agent. It learns a buyer's preferences by analyzing factors such as past purchases, browsing behavior, and stated preferences. A team of agents can use this data to search multiple e-commerce outlets and suggest purchases that best match a person's interests and budget.

Building intent-centric applications requires new approaches to software development. A key challenge is giving AI agents reliable access to the data they need, as they cannot fill gaps through informal communication as humans do. Development teams must also modernize extract/transfer/load pipelines, data fabrics, and integration frameworks to ensure that AI agents have prompt access to trustworthy, contextual data.

Organizations are addressing this gap by consolidating fragmented data into unified, shared platforms called data lakes or lakehouses to enable agentic collaboration. These repositories have the same scalability and performance benefits as traditional data warehouses but accommodate a wide variety of data types, including unstructured data such as videos and free-form text.

**Workflow engines  
orchestrate AI  
agents to optimize  
operations**

and accelerate  
throughput, enabling  
teams to handle more  
requests efficiently—  
**without adding  
extra resources.**

Because agents can make decisions and take actions autonomously, organizations must apply strong guardrails to protect sensitive information. A multilayer governance approach applies data compliance, agent observability, and security controls using tools like role-based access control, encryption, and audit logs. Agents should also be audited regularly to protect against “model drift,” or performance degradation that can occur over time and compromise accuracy and reliability.

To account for the probabilistic nature of AI models, Rohan advises testing outputs against predefined roles or expected results. “Treat agents as ‘digital insiders,’” he recommends. “Apply least-privilege principles, secure tokens, and runtime isolation to prevent misuse while implementing observability tools to track agent behavior and detect anomalies.”

## AI-enabled modernization

Agent training and inference work best when cloud-native constructs such as microservices, orchestration, and serverless architectures are applied to enhance speed and efficiency. For example, containerized environments allow for faster experimentation with

AI models, with the cloud dynamically allocating resources to optimize compute efficiency and reduce latency. Robust application programming interfaces (APIs) provide consistent, high-quality, real-time access to structured, machine-readable data at scale.

## Intent-centric applications in practice

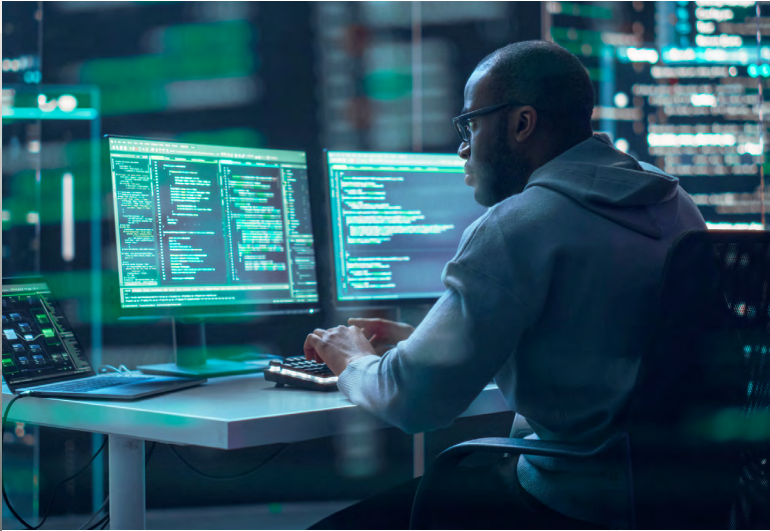
A global technology distributor depended on an ordering ecosystem spanning multiple enterprise resource planning systems and required frequent manual intervention.

To achieve the goal of scaling to resellers and thousands of vendors, a unified, automated platform was needed that could support consistent operations across regions while reducing operational overhead. The platform needed to integrate with multiple partner and reseller operational applications.

The company modernized on an integrated cloud-based platform using industrial accelerators and AI-driven components. External clients can now connect directly to the system to place orders. An agentic AI layer, developed on top of cloud-native Microsoft Azure AI and data services, automates key tasks such as order routing, validation, and workflow execution – with human oversight.

The platform is built on a combination of modernization patterns and AI constructs to be scalable, reliable, and extensible while enabling seamless integration with diverse partner ERP environments. Automated processes have replaced manual handling, supporting faster partner onboarding while reducing processing time.

The result is a streamlined, AI-enabled order management platform that has expanded the organization's capacity to support new channels, introduced products more quickly, and maintained consistent operations globally.




## Shift #2: AI-augmented development

One of the earliest successes of generative AI was in augmenting software development, a domain rich with structured patterns and repetitive tasks. AI copilots have demonstrated that models can reliably suggest code snippets, generate unit tests, and automate boilerplate within familiar developer workflows.

Generative engines excel in this area in part because modern programming languages follow strict syntactic rules, and there is abundant high-quality training data in public repositories, documentation, and coding forums. Generative AI copilots and coding assistants have already seen widespread adoption; 82% of developers use them.

This paradigm applies AI throughout the entire development life cycle to improve developer productivity and deliverable quality. Agents will fuel the move beyond copilots to agentic engineering workflows, characterized by real-time suggestions, autonomous code generation, automated testing, security scanning, architectural recommendations, and cloud configuration suggestions. AI can reduce errors, improve maintainability, and deliver applications faster. Collaboration is enhanced as AI copilots help teams standardize coding practices and optimize performance across diverse projects.

An example of AI's transformative power in software engineering is NTT DATA's aXet, a generative-AI-driven application modernization toolset that analyzes legacy code, maps architectures, and automates conversion to modern frameworks.



**Figure 2**  
**Developers already using Generative AI agents and coding assistants.**

## AI-enabled modernization

aXet reverse engineers monolithic applications, identifies patterns and technical debt, and generates transformation recommendations or code scaffolding for newer stacks such as cloud-native .NET, Java, or JavaScript frameworks. It connects to Microsoft Azure tools such as GitHub, Azure DevOps, and Azure Machine Learning. aXet is a central platform for AI-augmented engineering and a foundation for agentic evolution.

The introduction of agentic engineering workflows in software development “changes the role of software engineers from writing code line-by-line to managing AI agents that autonomously generate, review, and deploy code,” Rohan says. “Agile practices persist but adapt as humans guide agents, set priorities, and validate outputs instead of performing all tasks themselves.”

## AI-augmented development in practice

A large logistics provider used an AI-enabled modernization approach to accelerate the transformation of more than 40 legacy applications.

The legacy systems were monolithic, difficult to scale, and increasingly costly to maintain. The IT team used generative AI tools to analyze each application’s architecture, code base, and dependencies. They produced detailed insights into technical debt, modernization requirements, and refactoring pathways.

The automated analysis replaced weeks of manual code review and provided a consistent blueprint for migration. Generative AI was applied to automated code recommendations. Pattern extraction and transformation routines reduced the development burden and ensured that the applications adhered to cloud-native design principles. A modern stack was built using NTT DATA tools in an NTT DATA industry cloud.

The initiative improved scalability and reliability, optimized cloud infrastructure usage, and significantly reduced technical debt while lowering overall conversion costs. More than 190 applications were ultimately migrated using this AI-driven method.



### Shift #3: Industry-specific intelligent platforms

Organizations requiring functionality specific to their industries have often had an uncomfortable choice to make. General-purpose applications enjoy broad market support, offer flexible features, and typically have lower upfront costs because they spread development investment across many use cases. Support is enhanced by large user communities

**Mature industry clouds embed industry-specific data models, compliance frameworks, and AI-driven workflows, not just bolt-on templates**

– Rohan Braganza, Software & Data Engineering Leader at NTT DATA.

and third-party ecosystems that drive continuous improvement. However, they can be difficult to customize to fit specialized workflows, which can erode value over time.

Industry-specific applications, by contrast, embed domain knowledge and regulatory considerations out of the box, reducing configuration burden and improving alignment with sector-specific processes. The tradeoff comes in smaller ecosystems, higher per-customer pricing, slower innovation cycles, and a risk of vendor lock-in if interoperability is limited. Vendors of industry platforms also tend to be small and more prone to competitive pressures and business downturns than their large peers.

Industry-specific intelligent platforms are cloud-native versions of mainstream enterprise platforms tailored to industry outcomes. They offer numerous advantages over general-purpose applications.

- **Sector-specific data models** encode the semantics, regulations, and relationships unique to a given industry. Benefits include reduced data mapping effort, better interoperability across core systems, and improved compliance.

- **Prebuilt workflows** streamline adoption by codifying industry best practices, allowing organizations to stand up new capabilities with minimal configuration.
- **Domain APIs** expose business functions that are common to each industry, such as claims processing or patient intake. That makes it easier to integrate partners and leverage modern digital channels.
- **Integrated generative AI services** bring context-aware summarization, forecasting, and task automation directly into industry workflows without requiring customers to assemble complex AI stacks.
- **Composable design** builds systems from modular, interchangeable building blocks rather than fixed, tightly coupled components. This allows customers to more easily adapt to change by mixing and matching modular capabilities without rebuilding entire systems.
- **Composable infrastructure, applications, and services** support continuous innovation by enabling rapid reconfiguration and reuse of components as regulations and customer expectations shift.

**+70%**

**of enterprises will use industry clouds by 2027,** up from less than 15% in 2023.

By most accounts, the market for intelligent industry platforms is set for robust growth. Gartner expects that more than 70% of enterprises will use industry clouds by 2027, up from less than 15% in 2023. Foundry's Cloud Computing Study found that 40% of IT decision-makers rated the importance of industry-specific cloud offerings at 7 or 8 on a 10-point scale. The most popular perceived benefits include better security and application performance.

Cloud hyperscalers working with global system integrators are aggressively expanding AI-centered assets in their industry clouds. Modern industry platforms combine cloud-native architecture, AI model training, real-time automation, and domain-specific accelerators.

For example, Microsoft for Manufacturing, Azure cloud now offers AI-powered predictive equipment maintenance, supply chain optimization fueled by AI demand forecasting, quality

## Industry-specific intelligent platforms in practice

A global automotive manufacturer needed to modernize its digital ecosystem to deliver a unified customer experience for electric vehicle charging, service, and ownership. The company relied on multiple backend systems and monolithic legacy applications burdened by technical debt, inconsistent workflows, limited scalability, and slow release cycles.

The company needed an extensible, cloud-based platform capable of harmonizing disparate systems into a single experience layer. It undertook a modernization effort built on an integrated, cloud-native platform supported by modern front-end technologies and AI-enabled constructs.

The resulting service platform built with modern frameworks supports both web and mobile applications, ensuring consistency across all customer touchpoints. A scalable backbone unifies EV charging, service management, and a broad range of vehicle-related features. This integration has simplified customer interactions and enabled faster rollout of enhancements across markets.

The integrated application experience connects customers to products, services, and benefits across the company's entire portfolio, providing faster time-to-market, new revenue opportunities, and a more agile digital foundation.

control leveraging computer vision and machine learning models, and AI-enhanced “digital twins,” or computer-generated replicas of physical assets.

Healthcare providers are using AI-powered diagnostics and patient engagement tools to improve outcomes and reduce costs. Financial institutions are deploying AI to dramatically improve fraud detection and provide hyper-personalized services that enhance customer loyalty.

Industry platforms address enterprises' need to manage outcomes, not infrastructure. “Mature industry clouds embed industry-specific data models, compliance frameworks, and AI-driven workflows, not just bolt-on templates,” says NTT DATA's Rohan. “Instead of monolithic apps, they use modular, cloud-native architectures with AI being the connective tissue for orchestration, personalization, and predictive insights.”

## A new mandate for digital leaders

Traditional applications were created to automate processes and workflows that were expected to change little over time. Their rigid architectures and hard-coded logic made them resistant to adapting to volatile and unpredictable business conditions.

Modern cloud applications are built to meet a different set of requirements. They understand intent, adapt autonomously, make suggestions and decisions, conform to domain models, and operate on cloud-native, modular architectures. Loosely coupled architectures and APIs promote extensibility and foster innovation through third-party ecosystems. Making changes no longer requires taking systems offline for hours.

To achieve the benefits of modernization, organizations need to master domain-driven engineering and embrace unified telemetry and API governance while ensuring data readiness, security, and compliance. The combination of an AI-augmented systems development life cycle and composable services accelerates innovation and dramatically reduces time-to-market.

Platform-led thinking replaces application-by-application adoption.

Modernization is no longer a goal but an enabler of business transformation. AI becomes the fabric connecting systems, data, and workflows.

The NTT DATA and Microsoft partnership is grounded in intelligent modernization using AI. NTT DATA provides engineering accelerators such as Coding by NTT DATA and aXet. Microsoft Azure cloud provides the necessary cloud-native constructs and AI-native capabilities.

The approach combines the three paradigms of intent-centric design, intelligent engineering, and industry-specific platforms, with AI deeply embedded in engineering and platform architectures. Applications are not simply modernized but intelligently reimaged for the coming age of agentic AI. ♦

For more information on how **NTT DATA and Microsoft** can help transform your business through AI-enabled modernization [click here.](#)

© 2026 FoundryCo, Inc.