

TOP IT INSIGHTS FOR 2026

Navigating the Future of Technology and Business

Foreword

In 2026, business and IT leaders face critical decisions with emerging technologies that will determine if they succeed or fall behind. AI deployments will move from experimentation to execution, post-quantum cryptography will shift from future concern to near-term priority, and cloud strategies will evolve from wholesale migration to optimized placement.

Through extensive conversations with industry experts and our own company leaders, we've identified **10 insights** that will shape enterprise technology decisions this year.

Early patterns of how AI is being used and driving results for businesses are taking shape. We see organizations succeeding with **focused AI deployments** rather than transformational projects. The first repeatable high-ROI use cases are emerging, including knowledge assistants, coding agents and service operations tools.

Security needs will change. AI will play a leading role for both attackers and defenders, changing the way businesses approach cybersecurity investments and making recovery and resilience a priority. Leaders will need a **post-quantum cryptography roadmap**, as organizations defend against "harvest now, decrypt later"-style attacks to protect sensitive information.

We will also see **cloud strategies shift** from wholesale migration to workload-by-workload placement. Data sovereignty requirements and infrastructure efficiency will add complexity to existing systems, requiring leaders to prioritize the way they operate to ensure business continuity.

The way technology is impacting the workforce will also shift. Rather than seeing widespread AI-related layoffs and disappearing roles, we will see work evolve, with the greatest disruption impacting entry-level coding positions.

The insights are designed to help you refine your thinking and assist in your planning for the year ahead.

There's real momentum building. Let's capture it.

Mike Thomson

CEO and President, Unisys



Top IT Insights

- 01** Focused AI deployments will outpace transformation projects
- 02** Three AI applications will break through as repeatable high-ROI deployments
- 03** AI investments will shift from cost reduction to quality improvement
- 04** Organizations will train AI models on task-specific datasets rather than pursuing scale
- 05** Mass layoffs from AI automation won't occur, though entry-level coding positions will shrink
- 06** Organizations will need to establish post-quantum cryptography strategies
- 07** AI will accelerate both cyber attacks and defenses
- 08** Organizations will be measured by recovery speed, not breach prevention
- 09** Data sovereignty requirements will trigger the creation of regional and national clouds
- 10** Organizations will optimize workload placement rather than pursue wholesale cloud migration



**Focused AI deployments
will outpace transformation
projects**



The majority of AI deployments will shift from large-scale transformation efforts to smaller task-based integrations into existing processes. Focused deployments use smaller datasets that are easier to clean, require lower investment thresholds, enable smoother change management and deliver quicker payback.

Organizations are refining their AI strategies, moving toward focused implementations that deliver measurable results. Instead of pursuing large-scale transformations, teams are finding success with targeted deployments that fit into existing workflows.

The wins will come from projects that assist rather than replace. These deployments will help employees and customers complete existing work more quickly and easily. They'll deliver several practical advantages:

- Lower initial investment with fewer AI specialists needed and less data preparation required
- Faster deployment timelines
- Smoother change management
- Simpler execution
- Quicker learning cycles
- Higher success rates

Combined, these advantages will outweigh the risk-adjusted returns from larger, more ambitious AI projects. The faster learning cycles will also feed a second generation of deployments tackling more complex problems.

"Smaller, simpler AI use cases, where you do not have to reengineer a whole process or the people who run it — that's what's working."

Anthony Martucci

Vice President, Automation Hub, Unisys

Strategies for success

- Start with processes that have good documentation and clean underlying datasets.
- Integrate AI into processes as opposed to redesigning them.
- Invest in user training so people can maximize new capabilities.
- Build a pipeline of focused deployments that compound for a larger impact.



We are beyond the AI adoption stage; we are moving rapidly toward the pervasive stage."

**Chief Technology Officer,
Major U.S. Retail Company**



**Three AI applications will
break through as repeatable
high-ROI deployments**

Several repeatable high-ROI use cases will break through in 2026 as organizations move from experimentation to standardized deployment. Chatbots for employees and customers, AI coding agents and AI-driven service assistants that help level 1 techs resolve more complex IT issues will become packaged, measurable and fast to deploy — changing how organizations evaluate AI investments.

After years of experimentation, enterprises will converge on a handful of AI applications that can deliver value in relatively pre-packaged forms. Three standouts are:

01

Enterprise knowledge assistants that index governed content and answer with citations inside collaboration tools

02

Development agents tuned to an organization's codebase and policies, increasing development velocity while reducing defects


03

AI service operations that classify and help less experienced staff resolve incidents across IT, HR and finance

What makes these applications different? The infrastructure around them: secure retrieval, role-based access, prompt governance, human review and metrics tied to outcomes (e.g., cycle times, net promoter score). Vendors will productize these tools, and system integrators will deliver them as accelerators, shortening time-to-value from months to weeks. As adoption spreads, expectations will shift from “proof-of-concept AI” to “deployed AI.” Competitive pressure will mount because these apps change cost curves and experience simultaneously, speeding adoption even further.

“We have AI for customer service, but the caveat is that it all comes down to the quality of your data.”

**Senior Vice President of Information Technology,
Large U.S. Utility Company**



Strategies for success

- Prioritize packaged use cases with clear key performance indicators (KPIs) (e.g., cycle times, net promoter score, developer throughput).
- Shift expectations and funding from one-off proofs of concept to rapid deployments.
- Focus initial projects on areas where change management requirements are low to build confidence in both the technology and the ability of the organization to adopt it.



The chatbots in use today will pale in comparison to the much more human-like, conversational interactions we will have tomorrow."

Patrycja Sobera

Senior Vice President and General Manager,
Digital Workplace Solutions,
Unisys



**AI investments will shift
from cost reduction
to quality improvement**

Enterprises will redirect how they use AI from pure cost reductions and efficiency gains to improve quality processes — including accuracy, compliance, reliability and customer experience. Gains in first-time-right rates and reduced rework will outstrip labor savings as the primary value driver.


Early AI programs chased cost reduction because it was easy to model and sell. In 2026, the bigger prize is quality — measuring how AI improves decision confidence, reduces variance and elevates outcomes. In underwriting, AI can surface missing evidence and policy conflicts before coverage begins; in supply chain, it flags anomalous lead times and suggests alternates; in customer operations, it drafts empathetic replies with policy-aware language and cites the source. These improvements compound: Higher first-time-right reduces rework, which shortens cycle times, ultimately boosting customer satisfaction and revenue.

Organizations will build “quality observability” into AI: automated test suites for prompts and tools, red-team libraries and outcome dashboards tied to business controls (e.g., service-level agreements (SLAs), compliance rules). AI-assisted processes can be both more consistent and more tailored at the same time, changing the trade-off between standardization and customization. This can shift the conversation with finance. Instead of focusing on headcount reduction, teams can showcase increased revenue per hour and improvements across both margin and revenue.

“The AI focus shifts from reducing headcount to empowering employees.”

Alan Shen

Chief Architect, Enterprise Solution Strategy, Unisys



Strategies for success

- Add quality KPIs (e.g., accuracy, variance, first-time-right) to AI scorecards alongside cost metrics.
- Define how to measure quality/revenue impacts in ways that will satisfy finance decision-makers.
- Incentivize product owners on quality outcomes, not just cost savings.



The biggest shift with AI is going to be leveraging it as a work productivity tool."

**Chief Information Officer,
U.S. Manufacturing and Energy
Services Company**

Organizations will train AI models on task-specific datasets rather than pursuing scale

Task-specific fine-tuning using small specialized models will dominate enterprise build strategies. These approaches are cheaper to implement, can leverage proprietary datasets more easily and outperform generalist models on narrow jobs.

The pendulum will swing from giant generalists and large language models (LLMs) trained on massive datasets to compact specialists and small language models.

In practice, many enterprise tasks — classifying claims, drafting email responses, extracting fields from forms — will benefit from models tuned to industry- or even company-specific terms and practices. Curating a few thousand high-quality labeled examples often creates AI models that produce more accurate, more context-aware and overall higher-quality outputs than LLMs.

This approach will reduce compute spend and accelerate iteration while increasing control and simplifying risk management. Smaller models trained on cleaner data are easier to test and monitor, and they allow teams to fine-tune outputs with company tone, policy and overall intent.

The outcome: a faster path from idea to production and an easier way to govern the model on an ongoing basis. Organizations will discover that focused data beats massive scale when it comes to getting AI into production and keeping it performing well.

“Companies are not yet seeing the ROI for large-scale AI implementations and will turn to smaller, more focused projects.”

Joel Raper

Senior Vice President and Chief Commercial Officer, Unisys



Strategies for success

- Invest in data quality and active-learning loops rather than chasing scale.
- Define the most fertile areas to begin work and standardize approval workflows for small models.
- Monitor and govern the models after deployment to ensure that they stay on task and maintain quality.



AI works well when it is focused on smaller, even partial, tasks."

**Chief Technology Officer,
Major U.S. Retail Company**

Mass layoffs from AI automation won't occur, though entry-level coding positions will shrink

Despite automation, broad AI-driven layoffs will not materialize in 2026. Organizations will redeploy capacity to growth, quality and resilience while managing change thoughtfully. However, AI agents will automate routine coding, shrinking pure “junior coder” roles.

Enterprises will learn that blunt headcount cuts undermine transformation and jeopardize key controls. Instead, leaders will redirect productivity gains to backlog reduction, customer experience and modernization. Many tasks will be automated, but roles evolve. Analysts will become insight curators; support agents will become case managers; engineers will become system owners assisted by agents.

This shift requires reskilling at scale and transparent communication to maintain trust. Labor relations, brand considerations and regulatory scrutiny also dampen mass layoffs. Change programs that pair automation with upskilling and internal mobility outperform slash-and-burn tactics on both morale and business outcomes.


Within the developer suite, generative tools now draft boilerplates, unit tests and scaffolds, raising the bar for net-new human value. Teams will still hire early-career talent, but expectations will change: understanding systems, shaping requirements, evaluating AI outputs and owning small services end-to-end.

Apprenticeship moves from “ticket factory” to “pair with AI + mentor,” emphasizing debugging, architecture literacy and safe use of automation. Meanwhile, demand grows for roles adjacent to code: security engineering, data engineering and developer platform teams.

“The volume of code that entry-level coders provide will not be needed in the same way.”

Joel Raper

Senior Vice President & Chief Commercial Officer,
Unisys



Strategies for success

- Be wary of business cases built on large-scale headcount replacement — they are hard to realize and may carry invisible costs (e.g., degradation in the customer experience).
- Recognize that employees have valuable institutional knowledge that may be key to enabling processes to work fluidly and with high-quality outcomes.
- Track value-based metrics (e.g., customer outcomes, backlog burn-down) to show return on investment.
- Redesign early-career programs around product thinking, testing and AI-assisted workflows.
- Pair junior coders with mentors and agents; measure learning velocity, not lines of code.



I think we are years away from actually materializing the reduction of the workforce."

**Chief Information Officer,
U.S. Manufacturing and Energy
Services Company**



**Organizations will need
to establish post-quantum
cryptography strategies**

The “harvest-now, decrypt-later” threat makes post-quantum cryptography (PQC) a 2026 priority. Organizations will inventory cryptography, pick migration paths and begin phased rollouts aligned to emerging standards.

While large-scale quantum computers are not yet here, bad actors are stockpiling encrypted data. Standards bodies are finalizing PQC algorithms and guidance, and vendors are adding crypto-agility to stacks.

The practical work: discover where and how cryptography is used (protocols, libraries, certificates), assess sensitivity and shelf life of protected data and plan transitions to PQC-approved algorithms with hybrid modes during migration.

A variety of upgrades may be required, as will testing for performance impacts and compatibility. Smart boards and regulators will increasingly ask for PQC roadmaps and evidence of progress.

Treat PQC like a multi-year Y2K: inventory, prioritize, pilot, migrate and monitor. Organizations that start mapping their cryptographic dependencies in 2026 will be far better positioned than those waiting for quantum computers to arrive or, worse, for the first major breach.

“People will not adopt quantum without a breach to prod them, but that breach is coming.”

Alan Shen

Chief Architect, Enterprise Solution Strategy, Unisys



Strategies for success

- Launch a crypto inventory and classification program (e.g., systems, data, vendors).
- Prioritize protection of secrets/data that have long-lived value (e.g., intellectual property, medical records, national IDs).
- Contractually require PQC readiness from suppliers.



The investment timeframe continues to shift forward on quantum. What was once thought to be years away is starting to creep into the near-term horizon."

Mike Thomson
CEO and President,
Unisys



**AI will accelerate both
cyber attacks and defenses**

AI will play a major role in cybersecurity for both attackers and defenders alike — with tooling that scales reconnaissance, social engineering and malware as well as AI-powered detection, response and deception. Prevention will no longer be the main priority. Instead, speed and adaptation will become defining factors in protecting sensitive information.

In 2026, offensive AI will accelerate phishing personalization, deepfakes and voice spoofing, creating complex new threats of previously unimaginable quality. On defense, AI will augment pattern recognition and anomaly detection as well as threat hunting with natural-language interfaces and automated responses. The battleground shifts in terms of both speed and volume as AI enables attackers to deploy at scales previously unseen.

Security programs will respond with their own automations and planning exercises that include AI attack techniques. Insurance and regulators will ask for evidence of AI-aware controls and incident drills. The goal is not perfect prevention but fast containment, credible forensics and resilient recovery. Organizations that prepare for AI-enhanced attacks now will fare better than those assuming traditional security approaches will suffice.

“Think about what we are spending on how to use AI — the bad actors are spending that and more on how to break AI.”

Mike Thomson

CEO and President, Unisys



Strategies for success

- Add AI-specific capabilities to threat models (e.g., prompt injection, data poisoning, model theft).
- Deploy AI-assisted tooling with rigorous evaluation and human adjudication.
- Implement content provenance/watermark checks for executive communications.
- Expand tabletop exercises to include AI-enabled attacks and deepfake scenarios.
- Require vendors to disclose AI components and security hardening.



Bad actors are becoming more sophisticated — better crafted emails, voice fakes — and the good actors will have to respond in kind."

Sean Tinney

Senior Vice President and General Manager,
Enterprise Computing Solutions,
Unisys



**Organizations will be
measured by recovery speed,
not breach prevention**

With breach likelihood rising, enterprises will invest equally in rapid recovery and business continuity. Immutable backups, clean-room rebuilds and rehearsed failovers become board-level metrics.

Ransomware, supply-chain compromise and cloud misconfigurations will make “assume breach” a practical reality. The fastest path to resilience is preparation: immutable offline backups, golden images, clean-room rebuild capabilities and pre-negotiated crisis vendors and processes.

Observability will extend to recovery SLAs: time to isolate, restore critical services and validate integrity. Architecture choices — blast-radius segmentation, least privilege, secrets hygiene — will determine how bad a day becomes.

Regular game days and cross-functional playbooks (e.g., IT, security, legal, communications, operations) will reduce chaos and reputational damage. Customers and regulators will increasingly ask how quickly and transparently you recovered, not whether you were breached. Investment shifts from more prevention tools to balanced resilience portfolios.

Organizations that can demonstrate recovery capabilities will have competitive advantages in customer trust, insurance rates and regulatory relationships.

“Breaches are a question of when, not if.”

Vice President and Chief Information Security Officer,
Major U.S. Financial Services Company

Strategies for success

- Define recovery SLAs and measure them via periodic disaster-recovery rehearsals.
- Maintain offline immutable backups and test restores to production-like environments.
- Pre-stage clean-room rebuild pipelines and golden images for critical stacks.
- Segment networks/data to reduce blast radius and practice credential reset drills.
- Establish transparent communication protocols for clients, regulators and media.



Resilience and responsiveness will continue gaining traction as organizations increasingly view them as essential, cost-effective strategies for driving operational continuity and fortifying brand reputation against disruption."

Chris Arrasmith

Executive Vice President and Chief Operating Officer,
Unisys



**Data sovereignty requirements
will trigger the creation of
regional and national clouds**

Sovereignty will move from niche requirement to standard expectation in 2026. Governments and regulated industries will require data, keys and sometimes compute to remain within borders — driving regional and national cloud zones and partner ecosystems.

Sovereignty demands will expand beyond storage location to include identity, key custody and, in some cases, compute locality. Hyperscalers will respond with country-tuned offerings like isolated regions, customer-managed keys and local support. Sovereign cloud vendors will fill specific gaps for public sector and defense customers.

Enterprises operating across borders will likely face a sprawl of controls and conflicting contractual requirements. The complexity is operational as much as legal: workload placement, data residency catalogs, cross-border transfer approvals and resilience between sovereign zones.

Vendor lock-in risk will rise if sovereignty features are proprietary. Smart architectures will separate data and identity layers and provide robust tracking/proof of compliance. Organizations that map their sovereignty requirements early and choose platforms with genuine local control will navigate this complexity more successfully than those treating it as a compliance checkbox.

“Data sovereignty requirements are popping up everywhere, and once one jurisdiction adopts, others will follow.”

Patrycja Sobera

Senior Vice President and General Manager,
Digital Workplace Solutions,
Unisys



Strategies for success

- Map residency and sovereignty constraints by regulator/oversight authority, customer, dataset, identity and workload.
- Choose platforms with evidence of local control planes and key custody options.
- Negotiate contracts clarifying regulator access, support SLAs and audit rights.
- Offer clients a menu of sovereignty options: standard, enhanced and fully sovereign.



Regulatory requirements across jurisdictions often lack consistency and clarity, yet the push for greater oversight is undeniable. The EU continues to set the pace as the bellwether to watch."

Chris Arrasmith

Executive Vice President and Chief Operating Officer,
Unisys



Organizations will optimize workload placement rather than pursue wholesale cloud migration

The “lift-and-shift everything” era is over. 2026 strategies emphasize fit-for-purpose placement: private cloud for predictable workloads, sovereign zones for regulated data and selective rebalancing where economics warrant it.

Many large enterprises have completed their major cloud migration efforts and now run hybrid models: Some apps stay on-premises or in private clouds to meet needs for privacy, control, latency, cost predictability or licensing; others take advantage of hyperscaler services; and regulated datasets land in sovereign zones.

New workloads will be evaluated through a “workload-by-workload” lens, balancing cost, performance, risk and vendor concentration. IT operations and finance teams will mature capacity planning for private clouds, while platform teams invest in common pipelines to make hybrid feel seamless to users.

Best practices will standardize golden paths and productize internal platforms so the user experience remains smooth across environments. The shift reflects a broader maturation: Organizations will move from “cloud first” to “cloud smart,” making placement decisions based on actual economics and requirements rather than strategic mandate.

“We’re going to see fewer and fewer whole-scale shifts of applications to the cloud in favor of load rebalancing.”

Manju Naglapur

Senior Vice President and General Manager,
Cloud, Applications & Infrastructure Solutions,
Unisys

Strategies for success

- Establish workload placement criteria and an approval forum (e.g., cost, latency, risk).
- Strengthen private cloud capabilities (e.g., autoscaling, observability, self-service) if and where necessary.
- Use common platform tooling and policies across on-premises and cloud.
- Track repatriation opportunities where unit economics favor owned capacity.
- Explicitly budget for data movement and the cost of sovereign storage.



The wholesale move to the cloud is done. Now we optimize — the loads, the response times, the cost. All of this is dynamic now."

**Head of Digital and Innovation,
U.S. Biopharmaceuticals Company**

Looking ahead

These ten insights reveal a technology landscape that's maturing fast. AI moves from experimentation to execution, cloud strategies shift from migration to optimization and security programs balance prevention with recovery.

These changes create real opportunities for leaders who act decisively. Three factors will determine organizational success through 2026.

Focused execution over grand ambitions

Success comes from targeted deployments rather than sweeping transformations. Organizations must identify high-ROI AI applications — knowledge assistants, coding agents, service operations tools — and deploy them with clear metrics. This means building pipelines of focused projects, investing in quality improvements alongside cost reduction and balancing workload placement across private, public and sovereign clouds. Winners will compound small victories into meaningful advantage.

Prepared resilience

As technology becomes more central to operations, stronger defenses become essential. Organizations must establish post-quantum cryptography roadmaps, implement recovery-speed metrics and rehearse AI-enhanced breach scenarios. This includes maintaining immutable backups, segmenting networks and meeting evolving data sovereignty requirements. Resilience affects customer trust, insurance costs and regulatory relationships.

Thoughtful workforce evolution

AI will reshape roles rather than eliminate them. Organizations must redirect productivity gains toward quality, growth and modernization while managing change transparently. This means reskilling at scale, redesigning early-career programs around AI-assisted workflows and pairing automation with upskilling. Developers, analysts and service agents will work alongside AI tools — success requires preparing people for that partnership.

The path forward demands both urgency and precision. Move too slowly and competitors capture the advantages of packaged AI and optimized cloud economics. Rush without preparation and risk failed deployments or workforce disruption.

Organizations that succeed will stay clear about their technology priorities while adapting as tools mature — rapid learning cycles, governance that enables speed and measuring what matters: quality improvements, recovery speed, customer experience. The shifts we've outlined will determine competitive positioning for years ahead.



unisys.com

© 2026 Unisys Corporation. All rights reserved.

Unisys and other Unisys products and services mentioned herein, as well as their respective logos, are trademarks or registered trademarks of Unisys Corporation. The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date in the document properties. This content is provided for general information purposes only and should only be construed as guidance. Decisions should not be based solely on the information presented and should not replace consultation with professional advisors. All brands and products referenced in this document are acknowledged to be trademarks or registered trademarks of their respective owners. Nothing within this document may be reproduced or transmitted without citing "Top IT Insights for 2026: Navigating the Future of Technology and Business, Unisys Corporation." Unisys disclaims any responsibility for losses incurred by a reader as a result of the reader's decisions made based on any information included herein.

