




IDENTITY SECURITY FOR THE AGENTIC ENTERPRISE

GOVERNING HUMAN, MACHINE AND AI IDENTITIES AT SCALE

SEPTEMBER 2026

TABLE OF CONTENTS

Executive Summary	3
The agentic shift and what it changes	4
Where enterprises stand today	5
The evidence regulators now demand	7
AI-ready identity requirements and SailPoint's role	8
Aligned for AI-ready identity	9
Operating model before platform	9
Running identity operations at scale	10
DXC delivery accelerators and implementation assets.....	11
Governing identity in practice: customer use cases	12
Measuring the return	13
A phased roadmap	14
Conclusion	14
Sources and references	15



Executive Summary

Identity is the control layer that determines who and what can access enterprise systems, data, applications, and critical business processes. As enterprises adopt cloud, automation, and AI, effective identity governance must extend beyond employees, partners, and customers to devices, workloads, service accounts, APIs, automated processes, and AI agents. This broader control plane enforces least privilege, automates access lifecycles, and creates auditable evidence of access decisions, enabling organizations to scale AI securely and meet regulatory obligations.

Identity is the enterprise control layer for AI. As organizations embed generative and agentic AI into core operations, the identities able to access data, modify records, and execute sensitive actions are growing faster than legacy governance frameworks can manage. Agentic AI—autonomous systems that reason, invoke tools, and act with limited human intervention—introduces identity principals that most programs were not designed to govern. Boards therefore need identity controls that make automated actions attributable, constrained, and reversible.

Industry evidence reinforces the urgency. Gartner identifies machine identity and access management as one of the least mature and fastest-growing areas of identity security and forecasts that, by 2028, approximately 25% of enterprise breaches will involve the misuse or compromise of AI agents (1). IDC forecasts that more than one billion AI agents will be deployed worldwide by 2029—approximately 40 times the 2025 level (2, 3). Together with DXC and SailPoint observations, these findings position identity as the control layer for AI, cloud, and digital transformation across human, machine, and autonomous identities.


For business and technology leaders, weak identity governance creates three material risks: unmanaged and overprivileged identities expand the attack surface; inadequate attribution and auditability weaken compliance with GDPR, NIS2, DORA, and the EU AI Act; and insufficient governance delays AI approvals and deployment, allowing better-prepared competitors to move faster.

This paper sets out what AI-ready identity governance requires:

- a unified inventory and ownership model
- spanning human, machine, workload and AI identities
- least privilege and scoped delegated authority for agents
- automated identity and credential lifecycles
- automated actions that are auditable by design

It then sets out how that capability is delivered in practice through DXC's **Digital Identity** advisory and governance operating model, SailPoint's AI-driven identity security platform and DXC OASIS-enabled managed services that run identity operations at machine scale.

The outcomes are measurable, and they should be managed as such rather than asserted. Organizations that establish identity as the foundation of AI governance will be positioned to move faster on AI with a defensible control model behind it.



The agentic shift and what it changes

Artificial intelligence has moved through three phases in enterprise use. Machine learning produced predictions inside defined boundaries. Generative AI produced content and responses at scale. Agentic AI acts: autonomous systems that plan, decide, invoke external tools and pursue goals with minimal human intervention.

That shift matters for identity because agents combine large language models with curated knowledge sources and external integrations to execute multi-step workflows across enterprise systems, accessing data, modifying records, invoking APIs and completing business processes on behalf of users and organizations. **The unit of control moves from the user session to the delegated task.**

In traditional environments, a human authenticates, receives role-based access, and has their actions logged. In agentic environments, the agent becomes an identity principal in its own right. **Governance must therefore cover the agent, the human, service, or organization it represents, the accounts and credentials it uses, the authority delegated to it, and the data, applications, and actions permitted under that authority. Because each control point operates at machine speed, governance requires task-scoped privilege, automated credential lifecycles, and action-level logging that records what the agent did and under whose authority.**

Most organizations are deploying agents into customer service, finance, supply chain and research faster than they are establishing any of the three governance controls.

Market trends: Gartner's strategic technology trends research indicates that agentic AI is expected to become embedded in a significant share of enterprise software by 2028, with autonomous decision-making becoming part of day-to-day work. This reinforces the need for identity governance models that can control not only human users, but also AI agents acting under delegated authority. Gartner predicts that, "by 2028, at least 15% of day-to-day work decisions will be taken autonomously through agentic AI, up from 0% in 2024." (4).

Leadership implication: Agent deployments become reversible and auditable only where each agent is registered as a distinct identity with a named business owner, delegated authority is scoped to the specific task rather than to the principal's full entitlement set, and action-level logging is a condition of production approval. Where those three are absent, the enterprise accumulates actions it cannot attribute, delegation it cannot bound and liability it cannot allocate.

DXC delivery experience: From DXC's delivery experience, organizations often achieve the quickest progress by strengthening ownership before expanding AI technology deployment. DXC identity engagements frequently uncover service accounts and agents that can run in production without named owners, defined lifecycles or clear approval records. Establishing an ownership register helps create accountability, strengthen governance and deliver early, measurable risk reduction.



"Discovering and governing non-human and agentic identities makes visible the fast-growing attack surface. Avoiding manual identity administration and establishing digital trust are needed if we are to scale AI securely."

Martin Reilly,
DXC Digital Identity

Where enterprises stand today

Organizations deploying agentic AI at scale run up against the same four constraints, and they compound each other. They are a gap in governance rather than a tooling defect, which is why acquiring a platform without closing them tends to reinforce the gap and channel resources to execute the governance deficit faster.

1. Discovery. Non-human identities proliferate across cloud platforms, DevOps pipelines, containerized environments and AI infrastructure. Service accounts, API keys, OAuth tokens, certificates and hardcoded credentials are embedded across code repositories, configuration systems, container registries and legacy applications.

Many organizations cannot state which identities hold access to critical systems, and most have no mechanism to discover identities or credentials that exist outside their primary directories or credential repositories.

2. Ownership. Where a service account is compromised or an agent's permissions are misused, responsibility is frequently impossible to establish. Teams often do not know who created an identity, who approved its permissions or who should remediate it. That gap delays incident response directly, because containment waits on a decision no one is clearly authorized to make. The absence of formally assigned owners and delegated backups further increases risk, as non-human identities frequently outlive projects, personnel changes or organizational restructuring. Without clear ownership delegation, critical remediation actions can be delayed when the primary owner is unavailable.

3. Standing privilege and credential sprawl. Non-human identities are commonly over-privileged, accumulate entitlements as new requests are granted and old ones are never revoked, and often share credentials across development, test and production environments. They are created readily and decommissioned rarely, and their credentials are issued without rotation. Secrets remain embedded in source code, configuration files, container images and deployment pipelines, frequently with no centralized vault or enforcement to route them through one. The result is a large population of identities holding excessive privileged access, each one threatening a broad blast radius for a single compromise.

4. Audit record. Regulatory and operational requirements demand a record of who did what, when and why. For agents a record with this information is harder to produce:

- what action was taken
- what data was reached
- and whether the behaviors were intended or anomalous.

Without action-level logging, these questions cannot be answered inside the window of scrutiny an auditor or regulator allows.



Market evidence: SailPoint's Machine Identity Crisis research (5), a global survey of more than 320 identity and access experts, security professionals and executives. It found that 69% of organizations manage more machine identities than human identities, with nearly half managing ten times as many, and that 72% consider machine identities harder to manage than human ones. SailPoint's agent research indicates the gap is widening rather than closing: 97% of AI agents have access to sensitive data, while only 21% of organizations report high confidence in their ability to manage AI agent security risk. Gartner's machine IAM research describes machine identity management as one of the least mature areas of most identity programs, despite being among the fastest growing in scale and exposure.

Leadership implication: Leadership should treat these constraints as one operating-model gap rather than four technical defects. The sequence is straightforward: discover and inventory identities, assign ownership, right-size privilege, automate lifecycles, and close logging gaps. Progress can be measured through the percentage of non-human identities with a named owner and the percentage of credentials under managed rotation.

Customer observation, anonymized: Discovery exercises frequently reveal a materially larger population of service accounts and API credentials than the organization's own inventory records, with a meaningful proportion holding standing privileged access and no identified owner. This gap is typically the first constraint on scaling automation safely.



"Most of our customers are surprised by just how many non-human identities are discovered. 80% of our customers find x10 more than they had estimated before the discovery engagement and for the rest it is nearer x50."

Patrick Hartmann

**DXC Global Identity
Governance Technical Lead.**

The evidence regulators now demand

Four frameworks converge on the same operational requirement: organizations must be able to show which identity authorized an automated action, under what authority, and with what oversight. **GDPR** requires traceability from automated decisions to the initiating identity; **NIS2** requires visibility, ownership, and lifecycle governance for identities with access to critical systems; **DORA** requires least privilege, recertification of access, and credential rotation evidence for critical access; and the **EU AI Act** requires scoped delegated authority, traceable decisions, and defined human intervention points for high-risk AI systems. A shared policy and audit layer can generate the evidence needed across all four compliance frameworks.

Leadership implication: Compliance failure in this area is rarely a policy gap. It is an evidence gap. Organizations frequently have the right policies but cannot reconstruct the identity audit trail behind an automated decision within the window a regulator or auditor allows. The practical response is to:

- map each obligation to a named identity control and a named evidence source
- extend records of processing and access to cover non-human and agent identities.



Audit evidence needs to be a continuously generated output of the identity platform rather than a periodic collection exercise. Where that holds, audit cycles shorten and regulated AI use cases can launch without a bespoke compliance program for each one.



AI-ready identity requirements and SailPoint's role

AI-ready identity requires five platform capabilities:

1. Continuous discovery of non-human identities.
2. Automated governance at machine scale.
3. Risk-informed access decisions.
4. Lifecycle automation across provisioning, credential rotation, certification, and decommissioning;
5. Continuous generation of audit evidence.

SailPoint addresses these requirements across human, machine, workload, and AI identities, enabling governance through a common policy and audit layer. The key architectural decision is to treat non-human and agent identities as first-class identity principals rather than as exceptions managed outside the identity platform.

AI-ready architecture treats machines, workloads, containers and agents as core identity entities, and in practice that means four things:

- Governance for all identity types consolidates onto a common policy and audit layer rather than fragmenting across systems.
- Credentials are retrieved from a centralized vault at runtime rather than embedded in code, so rotation no longer requires application redevelopment.
- Agents authenticate as distinct principals under authority scoped to their function, with actions logged for traceability.
- Zero Trust continuous authorization extends to non-human principals, so authority is reassessed as risk context changes rather than granted once at authentication.

Architectures built this way integrate with security operations for identity threat detection, apply consistently across multi-cloud and on-premises estates, and expose identity services through APIs. That combination allows governance rigor to be tuned to risk and business criticality rather than applied uniformly across every identity.

Leadership implication:

Platform capability alone does not produce governance. Automation applied to an immature process encodes the existing gaps and simply executes them faster, which is why discovery and ownership must precede large scale automation. Architecture decisions taken now also determine how much of the AI portfolio can be governed later, because architectures that treat machine and agent identity as special cases fragment policy enforcement and force manual workarounds that do not scale with agent volume.

The order that holds is to:

- establish the capability requirements
- assess current tooling against them
- prioritize discovery and evidence generation, which typically deliver the earliest measurable risk reduction.

Aligned for AI-ready identity

AI-ready identity governance requires both platform capability and sustained delivery. SailPoint provides a common policy and audit layer across human, machine, and agent identities, supported by machine-learning-based discovery and access decisioning. DXC provides the operating model, integration across hybrid and legacy estates, managed services for recurring obligations, and **DXC OASIS** as the measurement and operational intelligence layer. Together, they provide one accountable path from AI ambition to governed deployment, with a consistent control model and evidence base from strategy through operations.

Operating model before platform


Identity outcomes depend more on the operating model than on product selection. Before platform configuration begins, the governance model must assign ownership for every identity type, define decision rights for privileged and sensitive access, establish escalation and exception paths, and fix accountability for compliance evidence and incident response. Human judgment should remain mandatory for high-risk access, sensitive data, and policy exceptions, while automated intelligence supports decisions and records the intervention as audit evidence. Designing these controls for non-human actors from the outset prevents governance from degrading after implementation.

Leadership implication:

Platform selection and delivery capability should be evaluated as a single decision rather than two sequential ones. Organizations that procure the platform first and the operating capability afterwards commonly find that the governance model they needed was never designed, and that retrofitting it costs more than the platform did.

The question to put to any prospective partner is

- **who will own the operating model after go live**
- **and how its effectiveness will be measured.**



Running identity operations at scale

The constraint in most identity programs is operational capacity, not platform capability. Where the identity population grows faster than the team governing it, the manual operating model caps the rate at which AI can be adopted safely.

Fail to scale and:

- Access reviews accumulate.
- Credential rotation slips.
- Compliance evidence goes stale

DXC addresses the operational-capacity constraint through two complementary capabilities. DXC OASIS integrates identity, cloud, DevOps, security, and compliance data to prioritize risk, orchestrate lifecycle processes, monitor compliance, generate audit evidence, and trigger policy-based containment. DXC managed identity services run recurring activities such as access reviews, credential lifecycle management, anomaly triage, and evidence collection against governance-focused service levels. Together with SailPoint, they reduce operating cost per identity, improve detection and response, and allow governance to scale without proportional headcount.

Leadership implication:

Manual operations produce backlog, and backlog becomes risk, as overdue certifications, unrotated credentials and unresolved anomalies accumulate into unmanaged exposure.

To avoid risk:

- Automate the recurring, high volume work first, being discovery, provisioning and deprovisioning
- Credential rotation and evidence collection
- Reserve skilled analyst time for governance strategy, exception handling and policy development
- Decide deliberately which identity operations to retain in house and which to place with a managed service.

DXC delivery accelerators and implementation assets

Repeatability is what separates a delivered identity program from a bespoke one. Our practice runs on a set of reusable assets, refined across engagements, that shorten time to value and reduce the variance between one implementation and the next. Each asset exists because a client engagement demonstrated that building it from scratch was the most expensive part of the program.

Identity maturity assessment framework. A structured evaluation of the current identity estate against agentic AI requirements, covering discovery coverage, ownership completeness, privilege posture, lifecycle automation and evidence generation, is required. The output is a scored baseline with a prioritized gap list, and it doubles as the measurement baseline described in the following section.

Reference architectures. These are pre-defined patterns for governing human, machine, workload and agent identities on SailPoint, covering centralized secrets retrieval at runtime, scoped delegated authority for agents, continuous authorization for non-human principals and integration with security operations. Reference architectures remove the costliest element of an identity program, which is deciding the pattern while simultaneously building it.

Governance operating model templates. These provide proven ownership matrices, decision rights, escalation and exception paths, certification cycles, and accountability definitions for compliance evidence and incident response. Tailored to each client, they turn the operating model described in this paper into implementation-ready artifacts that organizations can adopt in weeks rather than develop over several quarters.

Integration playbooks. These are documented approaches for connecting SailPoint to the systems where identity governance usually stalls: legacy applications without native connectors, ERP and core banking platforms, DevOps toolchains, container registries, operational technology environments and multi-cloud identity providers. The playbooks carry the tested approach and the known constraints, so integration effort is estimated from experience rather than assumed.

Non-human identity discovery accelerators. These reusable assets and best practice methods surface service accounts, API keys, OAuth tokens, certificates and embedded credentials across directories, cloud platforms, code repositories, configuration systems and deployment pipelines. Discovery produces the inventory on which every later phase depends, and incomplete discovery is the most common reason later phases underdeliver.

Agent onboarding pattern. A standard route to production for AI agents is registration as a distinct identity, assignment of a named business owner, scoping of delegated authority to the task, configuration of action-level logging and definition of the human intervention points appropriate to the risk tier. Standardizing this route is what allows the tenth agent to be approved considerably faster than the first.



DXC OASIS operational intelligence. The operational layer described earlier, OASIS aggregates identity, cloud, DevOps, security and compliance data into a single view, orchestrating lifecycle processes across heterogeneous systems and generating timestamped audit evidence continuously rather than on request.

Managed service runbooks. These documented operating procedures for certification administration, credential rotation, anomaly triage, containment and evidence collection define service levels against governance outcomes rather than ticket throughput. Runbooks are what make service quality independent of which individuals are assigned to an engagement.

These assets are maintained as a practice investment rather than as engagement artifacts. They are updated as SailPoint's platform develops, as regulatory expectations sharpen and as agent architectures mature, which means a client engaging today inherits the accumulated experience of prior programs rather than the state of the practice when its own contract was signed.

The effect on time to value is measurable. Assessments are completed in weeks rather than months. Architecture decisions are ratified rather than researched. Integration estimates hold because they derive from comparable work. And the transition into managed operations is a defined handover rather than a negotiation, because the runbooks were written before the platform was configured.

Governing identity in practice: customer use cases

Four scenarios account for most demand: governing AI agents that access enterprise data; securing service accounts and API credentials; automating machine-identity access reviews; and supporting regulated AI deployment. Each uses the same control pattern: register the identity, assign ownership, scope delegated authority, automate the lifecycle, and log actions for audit and incident response. Once these controls exist, additional use cases can be onboarded faster and at lower marginal cost; treating each scenario as a separate project instead rebuilds and repays for the same controls.

Leadership implication:

Buyers should ask prospective partners to evidence their assets rather than describe their experience. Reference architectures, operating model templates, discovery tooling and runbooks are inspectable in a procurement process; delivery experience is not. Where a provider cannot show the artifacts, the client is funding their creation and carrying the risk of their absence.

Measuring the return

Identity investment is defended either on results or on activity, and only the first survives a budget cycle. Each benefit of AI-ready identity governance should therefore be paired with a baseline, a metric and a named executive owner.

Four categories carry most of the value:


1. **Risk reduction:** Complete discovery, enforced least privilege and proactive anomaly detection lower both the likelihood and the impact of identity-based breaches, with the scale of improvement depending on the maturity of the starting position.
2. **Compliance readiness:** Evidence is current rather than gathered under audit pressure.
3. **Operational efficiency:** This improves as automated rotation, certification and evidence collection eliminate routine administrative effort, enabling stronger governance without a proportional increase in team size.
4. **Accelerated AI adoption is where the greatest value is realized:**
Organizations with confidence in their identity control framework can approve and scale agentic AI use cases rapidly, while those lacking that confidence delay deployment and surrender first-mover advantage to competitors.

Two further effects compound over time. Human users meet less friction, through single sign-on, passwordless authentication and provisioning that grants appropriate access promptly. And leadership gains a defensible position on responsible AI adoption, which carries weight with customers, regulators and prospective talent.

Suggested measures: Track the percentage of non-human identities assigned to a named owner; measure the percentage of credentials rotated through a managed process; monitor the mean time to detect and contain identity-based incidents; report the percentage of access certifications completed on time; measure the time required to produce audit evidence; and track the elapsed time from approval of an AI use case to its production deployment.

Leadership implication:

Where no baseline is captured before implementation, improvement cannot be demonstrated and funding becomes difficult to sustain through the phases in which benefits compound. Capture the baseline before platform implementation begins, and report a small, stable metric set to the risk or audit committee on a quarterly cycle.



A phased roadmap

A phased roadmap builds control before scale.		
Phase 1	Months 0–3	Assesses the identity estate, establishes governance, and produces a costed roadmap with named owners.
Phase 2	Months 3–9	Implements the core governance platform, begins non-human identity discovery, centralizes sensitive credentials, and establishes baseline lifecycle and compliance controls.
Phase 3	Months 9–15	Extends lifecycle automation to non-human and agent identities and introduces continuous compliance monitoring and OASIS operational intelligence.
Phase 4	Months 15	Scales risk-aware governance, anomaly detection, and managed operations. Executive alignment, change management, cross-functional ownership, and stable outcome metrics support every phase.

Next step: Organizations at any stage of this journey can benefit from a structured assessment of current identity maturity against agentic AI requirements. DXC's **Digital Identity** advisory team works with leadership teams to establish that baseline and to produce a phased roadmap aligned to the organization's own AI adoption timeline.

Conclusion

Agentic AI is not introducing a new identity problem. Instead, it is highlighting a long-standing governance gap. Non-human identities already represent the largest identity population in most organizations. As AI agents increasingly act, transact and make decisions on behalf of the enterprise, these identities require the same rigor, control, and accountability as their human counterparts.

The choice facing leadership is therefore about pace. Organizations that cannot evidence who or what authorized an automated action will slow their own AI programs. Whether through their own caution or due to the scrutiny of regulators and auditors. Organizations that establish identity as their foundational control layer can approve new use cases faster, and each new use case inherits controls that already exist.

DXC and SailPoint provide that path in combination: SailPoint's AI-driven identity security platform, DXC's Digital Identity advisory and governance operating model, and **DXC OASIS-enabled** managed services to operate and maintain identity operations at scale. DXC welcomes the opportunity to discuss how identity can be established as the foundation of an organization's AI governance.



Sources and references

The evidence cited in this paper is drawn from publicly available research and published government regulations.

1. 'Gartner predicts that by 2028, 25% of organisational breaches will be traced back to AI agent abuse, underscoring how quickly this risk surface is expanding.' Gartner projection on AI agent abuse as a contributing factor in enterprise breaches by 2028, as reported in public industry coverage of Gartner research '[Why AI is forcing a reset of the identity stack](#) — Computer Weekly Security Think Tank, authored by Ted Ernst, Senior Director Analyst at Gartner (27 Apr 2026). [<https://www.computerweekly.com/opinion/Why-AI-is-forcing-a-reset-of-the-identity-stack>]
2. IDC FutureScape 2026 Predictions Reveal the Rise of Agentic AI and a Turning Point in Enterprise Transformation. 'IDC forecasts that by 2030, 45% of organizations will orchestrate AI agents at scale, embedding them across business functions.' Source: 'IDC FutureScape 2026 Predictions Reveal the Rise of Agentic AI — Business Wire official IDC press release (23 Oct 2026). [<https://www.businesswire.com/news/home/20251023490057/en/Agentic>]
3. 'IDC projects that the number of actively deployed AI agents will exceed 1 billion worldwide by 2029, approximately 40 times more than in 2025.'. IDC - Agent Adoption: The IT Industry's Next Great Inflection Point, December 2025, IDC blog, by Rick Villars. [<https://www.idc.com/resource-center/blog/agent-adoption-the-it-industrys-next-great-inflection-point/>]
4. 'Gartner predicts that, by 2028, at least 15% of day-to-day work decisions will be taken autonomously through agentic AI, up from 0% in 2024.' Source: Agentic AI is the top strategic technology trend for 2025 — ZDNET (29 Oct 2024). [<https://www.zdnet.com/article/agentic-ai-is-the-top-strategic-technology-trend-for-2025/>]
5. SailPoint, Machine Identity Crisis: The Challenges of Manual Processes and Hidden Risks, November 2024. A global survey of more than 320 identity and access experts, security professionals and executives. [<https://www.sailpoint.com/identity-library/machine-identity-security>]
6. SailPoint research on AI agent access to sensitive data and organizational confidence in managing AI agent security risk, 2025. '97% of all AI agents have access to sensitive data and only 21% of organizations are highly confident in their ability to manage AI agent security risks.'. [<https://www.sailpoint.com/press-releases/sailpoint-identity-security-solution>]
7. GDPR - Regulation (EU) 2016/679, General Data Protection Regulation. Regulation - 2016/679 - EN - gdpr - EUR-Lex [<https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>]
8. NIS 2 -Directive (EU) 2022/2555, Network and Information Security Directive 2 (NIS2). [<https://op.europa.eu/en/publication-detail/-/publication/9b84d482-85bd-11ed-9887-01aa75ed71a1/>]
9. DORA - Regulation (EU) 2022/2554, Digital Operational Resilience Act (DORA). [<https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng>]
10. AI Act - Regulation (EU) 2024/1689, Artificial Intelligence Act. [<https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng/>]

Further Reading:

- Gartner's strategic technology trends research indicates that agentic AI is expected to become embedded in a significant share of enterprise software by 2028, with autonomous decision-making becoming part of day-to-day work. This reinforces the need for identity governance models that can control not only human users, but also AI agents acting under delegated authority. [<https://www.gartner.com/en/documents/5850847>]
- Gartner, Innovation Insight: Improve Security with Machine IAM, March 2025. 'Inadequate security for machine identities poses significant security risks... IAM leaders should use this research to improve IAM controls for machines.' [<https://www.gartner.com/en/documents/6254251>]
- Gartner, Magic Quadrant for Privileged Access Management, October 2025, which extends evaluation criteria to privileged access for machines and cloud environments. [<https://www.gartner.com/en/documents/7051198>]
- KuppingerCole Analysts, Leadership Compass: Identity and Access Governance, Nitish Deshpande, Senior Analyst. [<https://www.kuppingercole.com/reprints/14bc3afd95f384b2d3e5198b0ad847b6>]



Patrick Hartmann

DXC Identity Governance
Technical Lead.



Martin Reilly

DXC Digital Identity Offerings.



Sugandha Sinha

DXC Intelligent Cyber Offerings.

Learn more at
DXC Cybersecurity



DXC Technology (NYSE: DXC) is a leading enterprise technology and innovation partner delivering software, services, and solutions to global enterprises and public sector organizations — helping them harness AI to drive outcomes at a time of exponential change with speed. With deep expertise in Managed Infrastructure Services, Application Modernization, and Industry-Specific Software Solutions, DXC modernizes, secures, and operates some of the world's most complex technology estates. Learn more on dxc.com.

[LinkedIn](#) | [Instagram](#) | [TikTok](#) | [YouTube](#)