




Securing agentic identity

A strategic imperative for the AI-driven enterprise

September 2026

Table of contents

Foreword	03
Executive summary	04
Identity as the foundation for bounded autonomy	06
The target architecture	08
Assurance and regulatory alignment	10
Technology evolution trends	13
A phased roadmap for action	14
Measures that reveal real progress	15
Future state	16
Recommendations	16
Conclusion	17
Connect with us	19

Foreword

India stands on the threshold of a transformation that will be remembered as a defining chapter in our digital journey. For the past two decades, Indian enterprises have digitalised processes. Today, they are moving far beyond digitalisation towards autonomous execution. Across banking, manufacturing, healthcare, telecommunications, energy, retail, public services and Global Capability Centres (GCCs), organisations are rapidly deploying agentic AI systems capable of making decisions, initiating actions, collaborating with other systems and delivering business outcomes with minimal human intervention.

This is not simply another technology cycle; it is the emergence of a new digital workforce. India's AI ambitions are accelerating at an unprecedented pace. Through the India AI Mission, the government has committed over INR10,000 crore to build a national AI ecosystem, democratise compute, develop indigenous foundation models and establish a Safe & Trusted AI framework. India's AI infrastructure has already expanded beyond 34,000 GPUs, creating the foundation for large-scale AI innovation and adoption. As organisations embrace this opportunity, a critical reality emerges.

Every AI agent is an identity. Every identity demands trust. Every trusted identity demands governance. For decades, cybersecurity has been designed primarily around the protection of human users. Enterprises are entering an era where thousands, and eventually millions, of machine identities will act autonomously across applications, cloud platforms, business processes, supply chains and customer ecosystems.

These AI agents will approve transactions, negotiate workflows, access confidential data, orchestrate business processes, interact with customers, generate software and make decisions and execute actions at machine speed.

The question is no longer whether organisations will deploy agentic AI. It is whether they can do so securely. India has one of the world's largest digital economies and is home to over 1,800 GCCs, a thriving start-up ecosystem, globally significant financial institutions and world-leading Digital Public Infrastructure (DPI). AI is expected to become a major force multiplier for India's economic growth and competitiveness.

If we secure agentic identity today, we will protect our organisations, strengthen citizen trust, help secure India's digital economy and establish India as a global leader in responsible and trustworthy AI. The age of autonomous business has arrived. Let us ensure it is built on a foundation of trust.



Gaurav Shukla
Partner and Leader - Cyber
Deloitte South Asia



Nitin Varma
SVP & Managing Director,
Saviynt India & SAARC

Executive summary

Agentic AI promises to redefine enterprise productivity by enabling autonomous digital workers that can reason, decide and act across business processes. The next phase of enterprise AI is moving from providing insights to independently driving actions.

Agentic AI can interpret intent, select tools, call APIs, retrieve sensitive data, initiate transactions, write code and delegate work to other agents. This creates extraordinary economic potential.

As organisations increasingly deploy AI agents to perform tasks once reserved for humans, competitive advantage will depend on the adoption of emerging AI capabilities and the ability to govern and scale these new identities securely.

Agentic AI is not merely another technology upgrade; it is the dawn of a new digital workforce.

However, a critical gap lies beneath the AI transformation. Traditional identity systems were built around human users, applications and workloads with consistent behaviour. Agentic systems challenge those assumptions. Their permissions may be dynamically determined at runtime based on the task and context. Their authority may originate with a user, a business process or another agent. Their actions may cross multiple clouds, models, tools and data domains before producing an outcome. Enterprises must therefore understand each agent's purpose, provenance, delegation, context and behaviour.

Today's enterprises must secure an ecosystem increasingly dominated by such agentic identities, machine identities, APIs and automated workloads. As these non-human identities multiply and operate autonomously across digital environments, they introduce new risks that demand stronger visibility, governance and accountability to protect business-critical assets.

Cost of inaction

Legacy Identity and Access Management (IAM) frameworks were never designed to govern autonomous systems that operate at machine speed, make independent decisions and interact across complex digital ecosystems. As AI-driven identities proliferate, organisations face heightened security exposure from rapidly escalating privileges, growing accountability and audit gaps due to limited visibility and increasing regulatory and board-level scrutiny over AI governance, transparency and risk management. What was once a security concern has now become a strategic business and compliance imperative.

The strategic imperative

Securing AI and non-human identities is no longer solely a cybersecurity challenge; it is a business-critical priority. As enterprises accelerate the adoption of agentic AI, they must ensure that innovation is matched by strong identity governance, visibility and accountability. The organisations that will lead in the AI era are those that govern AI identities with the same discipline, oversight and control applied to human identities. With increasing regulatory expectations around access governance, auditability and AI accountability under frameworks such as GDPR, SOX, HIPAA, NIS2 and

DORA, robust identity security has become a strategic enabler of trust, compliance and sustainable growth.

Agentic identity must become a first-class security domain. Every production agent should have a unique identity, a named owner, a declared purpose, bounded authority, verifiable provenance and a complete record of its actions. Identity must connect the initiating intent to the executing workload and the resulting business outcome. That connection is the foundation of trust.

Identity types in the agentic enterprise

A human identity is trusted to perform a role. A workload identity is trusted to execute a function. An agentic identity is trusted to exercise judgement and act. That distinction changes everything. For a human identity, the enterprise asks whether the person is authenticated and entitled. For a workload identity, the enterprise asks whether the service is genuine and authorised to reach the resource. For each agent action, the enterprise must assess whether it is authorised for the intended purpose, principal, tool and data, while accounting for the prevailing risk conditions. This is the progression from entitlement to access to authority. Access determines whether an identity may cross a boundary. Authority determines what the identity may do, how far its permissions extend, whether it may delegate to another identity, when human intervention is required and what evidence must be preserved. Agentic identity is therefore not simply another category of non-human identity. It is the control mechanism that binds intent, delegation, action and accountability in an autonomous enterprise.

Emerging security challenges

Many organisations struggle to detect unusual AI activity, make real-time access decisions and maintain continuous compliance monitoring. These gaps are compounded by agentic identities with excessive permissions, service accounts without clear ownership and credentials, such as API keys and OAuth tokens, that provide access to systems but are rarely reviewed or rotated. Without stronger lifecycle governance, these unmanaged identities and credentials create a rapidly expanding security and compliance blind spot.

The five failure modes that matter most

Together, these weaknesses give rise to five critical failure modes that organisations must address as agentic systems scale.

Invisible agency: Agents operate without a complete inventory, clear ownership, or a reliable link to the applications, models, tools, credentials and data they use.

Unbounded delegation: Agents receive broad standing access, delegate authority to another agent or combine individually acceptable permissions in ways that create unacceptable outcomes.

Credential without context:

Credentials are treated as sufficient proof of authorisation, even when user intent, purpose, workload identity, device posture, transaction value or environmental risk has changed.

Behavioural drift: Agents' actions move away from their approved purpose because of prompt manipulation, model changes, compromised tools, poisoned context or evolving workflows.

Broken accountability: Logs capture isolated API calls but cannot reconstruct the chain from human or system intent through delegation, reasoning, tool execution and business impact.

Structural identity challenges in the agentic enterprise

Human-paced controls cannot govern machine-speed authority. An agentic identity can create identities, acquire privileges, invoke tools and move across cloud and data environments long before a conventional access review begins. Permissions granted at deployment can drift from runtime purpose, inherited bias can influence autonomous decisions, unclear ownership can obscure accountability and one compromised credential can open a path for rapid, largely invisible

lateral movement. These are not isolated control failures; they expose a structural limitation in traditional IAM. The agentic enterprise requires identity security that operates continuously, evaluates authority at the moment of action, limits delegation, detects behavioural drift and preserves an unbroken chain of accountability from originating intent to business outcome.

Identity as the foundation for bounded autonomy

The enterprise security perimeter is no longer defined by networks or devices. Network boundaries cannot express the full meaning of an agent's action, and application controls rarely see the entire delegation chain. The enterprise security perimeter is increasingly defined by who or what can access critical systems, data and decision-making processes. Identity is the only control plane capable of linking the actor, the delegator, the

workload, the requested capability, the target resource and the surrounding risk context.

As AI agents, machine identities and automated workloads operate across cloud and SaaS environments, attackers are shifting their focus to compromised credentials, excessive privileges and unmanaged access pathways. This makes identity both the primary attack

surface and the most effective point of control. In the AI era, organisations that cannot continuously discover, govern and contain every identity will struggle to maintain security, accountability and trust. Identity security is therefore no longer a supporting capability; it is the control plane for cyber resilience, regulatory governance and enterprise risk.

A mature agentic identity control plane should deliver the following six integrated capabilities:

Discover: A mature agentic identity control plane must continuously identify every agent, service principal, credential, model, tool, endpoint and dependency across the enterprise.

Register and attest: Each AI agent must receive a unique identity only after its origin, workload integrity, accountable owner, declared purpose, operating environment and approved capabilities have been verified.

Govern: Ownership, risk classification, access certification, segregation of duties, exception management and timely retirement must be enforced throughout the agentic identity's lifecycle.

Delegate and authorise: The control plane must preserve who or what the agentic identity represents, constrain the scope and duration of delegated authority, evaluate every consequential action against intent, purpose, policy, data sensitivity and current risk and prevent uncontrolled onward delegation.

Observe and respond: Identity, policy, tool, data, execution and outcome telemetry must be continuously correlated to detect behavioural drift, interrupt unsafe activity, revoke compromised authority and trigger human intervention when required.

Prove: Finally, the control plane must preserve tamper-evident evidence that reconstructs the complete chain of accountability, from originating intent and delegated authority to policy decisions, agentic identity execution and resulting business impact.

Design principles for secure agentic identity

The following design principles translate these capabilities into practical controls for governing agentic identities.



Verify identity before granting access:

No agent should access enterprise tools or data without a unique, verifiable identity and an accountable owner.



Align permissions with purpose:

Base permissions on the agent's declared business purpose and approved capabilities, rather than assigning broad roles for convenience.



Define delegation explicitly: Record the initiating principal, scope and duration of delegated authority, along with any restrictions on further delegation.



Limit the duration of access:

Use temporary credentials, just-in-time access, proof of possession and automated credential rotation. Eliminate embedded secrets and minimise standing privileges.



Authorise every significant action:

Evaluate high-impact actions against current intent, resource sensitivity, transaction value, operating environment, behavioural risk and policy.



Apply human oversight based on risk:

Require approval for irreversible, high-value, safety-sensitive actions and policy exceptions, while allowing routine, low-risk activity to proceed without unnecessary review.



Trace every action from intent to outcome:

Connect user or system intent, agent identity, delegated authority, workload, model and tool activity, policy decisions and resulting outcomes.



Revoke access without delay:

Enable the immediate suspension of agents, revocation of delegated authority, rotation of credentials, disabling of tools and termination of related sessions.



Block actions when verification fails:

Prevent consequential actions when identity, delegation, policy or origin cannot be verified.

A lifecycle model from intent to retirement

Trust in an AI agent cannot begin at authentication or end at deployment. It must be designed into the agent's purpose, built into its software supply chain, verified through registration and attestation, constrained through narrow delegation, enforced at every consequential action and continuously tested through behavioural and outcome monitoring. Every material change must renew that trust, and retirement must conclusively revoke authority and eliminate residual access. This creates an unbroken evidence chain from design intent to runtime action and, ultimately, secure decommissioning.

The target architecture

The target state is not a new identity silo for AI. It is an orchestration layer that connects existing identity governance, access management, privileged access, workload identity, policy decision, security analytics and AI governance capabilities.

This architecture should include the following components:

- **A central agent registry** that records owner, purpose, risk tier, model, tools, data domains, environments, dependencies and lifecycle state.
- **A trust service** that establishes workload identity through platform-native credentials, certificates, attestation or other cryptographic mechanisms.
- **A delegation service** that carries the initiating principal, approved purpose, scope, audience, duration and delegation depth.
- **A policy decision and enforcement layer** that evaluates action-level context and can allow, deny, constrain, step up or require human approval.
- **A credential service** that issues short-lived, audience-bound, proof-of-possession credentials and rotates or revokes them automatically.
- **An evidence fabric** that correlates identity, prompts or instructions where permissible, model and tool activity, data access, policy decisions, approvals and outcomes.
- **An incident-control channel** that can disable an agent, isolate a tool, revoke authority and prevent cascading delegation.

Design-time identity layers for agentic identity

To secure and govern AI identities end to end, organisations need the following three connected layers that provide visibility, governance and runtime control.

Find it using identity security posture management.

Continuously discover every AI agent and non-human identity, assess security posture, score risk, detect shadow AI and identify privilege drift or anomalous behaviour before exposure becomes an incident.

Govern it with identity governance and administration.

Establish clear ownership and purpose for every identity, manage registration and lifecycle, certify access, enforce least privilege and maintain a complete audit trail and provenance chain.

Enforce it with runtime authorisations

Apply intent-aware authorisation and real-time policy enforcement to every action. Dynamically allow, restrict or block access, or escalate high-risk actions for human approval, while eliminating unnecessary standing permissions.

The objective is to discover every identity, govern its access and verify every significant action at runtime.

Run-time identity layers for agentic identity

A secure AI agent request requires an end-to-end chain of trust connecting the user who initiates the request, the authority given to the agent and the

system that performs the action.

Standard authentication protocols, such as OAuth and OIDC, verify the user and record consent. Token exchange or JSON Web Tokens (JWTs) allow the agent to receive limited authority for a specific task, rather than inheriting the user's full access. Workload identity standards, such as SPIFFE or WIMSE, give the agent a unique and verifiable identity. Proof-of-possession controls ensure that credentials can be used only by the agent to which they were issued. Together, these layers make each request authenticated, appropriately authorised and traceable, while reducing the risk of credential theft, reuse or misuse.

Modern Identity providers (IdPs) capabilities



Modern identity security requires continuous, context-aware control across the entire session lifecycle. Continuous access evaluation reassesses risk in real time, while fine-grained authorisation applies attribute, relationship and policy-based decisions to every request. Context-rich tokens carry the intelligence needed for adaptive access, and federated machine identities establish trusted workload authentication across cloud environments. Together with end-to-end session visibility, these capabilities create a unified security model in which access is not simply granted once, but continuously verified, adjusted and revoked as risk and context change.

Governance is a shared operating model

Identity security is reaching an inflection point. The rapid expansion of cloud, SaaS and hybrid environments, combined with autonomous agentic identities operating beyond traditional governance boundaries, is fundamentally reshaping the identity perimeter. No single function can secure agentic identity because no single function owns the entire risk.

These responsibilities must operate as one integrated governance model, not as disconnected control towers.

Accountability may be shared across the enterprise, but it must never be ambiguous. Every agent must have a named owner, every delegated authority must have a defined boundary and every material outcome must remain attributable to a responsible business decision.



Assurance and regulatory alignment

The regulatory case should be expressed with precision. Most laws and standards do not yet prescribe a complete “agentic identity” architecture. They do, however, create obligations or expectations relating to access control, accountability, human oversight, security, auditability, data protection, operational resilience and risk management. The practical requirement is to demonstrate that autonomous actions remain attributable, authorised, proportionate, reviewable and controllable.

Organisations should map agentic identity controls to the obligations that apply to their sector and geography, while tracking emerging guidance for AI-agent security and interoperability. Draft protocols and developing standards should inform architecture experiments, but they should not be represented as settled compliance requirements.

Global regulatory landscape for AI identity

Region	Regulation	Key requirements (AI identity and access)
EU	AI Act (Art. 9–15)	Human oversight, meaningful human control, detailed access logs, risk-tiered controls, consent controls for biometrics
	GDPR (Art. 22)	Data subject rights, portability and erasure support, consent withdrawal and secure cross-border identity safeguards
US	SOX	Executive-certified access controls, internal control coverage, segregation of duties and audit trails
	HIPAA	Minimum necessary access, encryption, access logging, patient consent, AI vendor compliance (BAA)
	PCI-DSS	Network segmentation, strong authentication, penetration testing and secure key management

AI identity management is no longer optional. It is a compliance imperative.

Industry-specific regulations and security frameworks

Industry	Framework/ regulation	Key AI identity and access requirements
Financial services	Basel III	Requires AI systems to implement operational risk control, AI identity governance, stress testing including identity failures
	MiFID II	Access controls for trading AI, kill switches, audit trails and regulatory reporting
Healthcare and Life Sciences	FDA 21 CFR Part 820	Requires AI medical devices to implement design controls, including identity management, lifecycle identity integration and risk management for identity hazards
	ICH GCP (Clinical Trials)	Data integrity, audit trails, e-signatures and investigator oversight via identity controls
	Health Insurance Portability and Accountability Act (HIPAA)	Mandates minimum necessary access principles for AI systems
Manufacturing and critical infrastructure	NIST CSF	AI asset inventory, access control, behaviour monitoring, incident response, recovery planning
	IEC 62443	Requires zone and conduit models to account for AI system communications, criticality-based controls and defence-in-depth access mechanisms

Industry-wide demand for AI-aware identity governance is accelerating.

Emerging AI governance standards and initiatives

Area	Framework / Initiative	Key focus
AI governance and ethics	ISO/IEC 23053	AI risk management including identity controls
	IEEE 2857	Privacy engineering and access control
	NIST AI RMF	Identity as core AI risk control
	W3C	AI identity and authentication specifications
Global harmonisation	GPAI	Identity management best practices
	OECD AI principles	Transparency and accountability via identity controls
	UN AI ethics	Human rights in AI identity management

Technology evolution trends

Critical technology developments are redefining how AI identities are created, authenticated, trusted and governed. As autonomous agentic identity becomes an active participant in enterprise ecosystems, identity architectures must evolve beyond static, human-centric controls towards models that are cryptographically resilient, federated, intelligent and capable of operating continuously at machine speed.

Quantum-safe AI identity

The emergence of quantum computing will require enterprises to reassess the cryptographic foundations protecting AI identities and communications. Post-quantum authentication, quantum-resistant certificate authorities, secure key-exchange mechanisms and quantum-safe digital signatures will become increasingly important for establishing long-term trust. Organisations must begin preparing for this transition now, particularly where agentic identity handles sensitive data, executes privileged actions or relies on credentials that must remain secure for many years.

Federated AI identity ecosystems

AI agentic identities will increasingly operate across organisational, cloud and platform boundaries, making isolated identity models unsustainable. Federated and decentralised identity architectures could enable organisations to authenticate agentic identity, verify their authority and preserve accountability across interconnected ecosystems. Technologies such as verifiable credentials, distributed ledgers, decentralised identifiers and self-sovereign identity models may help establish portable trust without surrendering control to a single identity provider.

AI-powered identity management

AI will create new identities and transform how they are governed. Machine learning can strengthen identity analytics, detect behavioural anomalies, optimise access policies and predict emerging risks before they become incidents. Over time, identity management will evolve from periodic, manually intensive administration towards autonomous lifecycle governance that continuously discovers identities, adjusts privileges, responds to risk and removes access when it is no longer required.

GenAI and autonomous agentic identity

Large Language Models (LLMs) and autonomous agentic identity are emerging as first-class enterprise principals capable of calling APIs, accessing data, executing code and initiating business actions. Unlike traditional service accounts, their permissions may change according to user intent, task context, delegated authority and runtime conditions. This shift demands identity controls that can continuously determine not only who an agentic identity is, but also on whose behalf it is acting, what it is permitted to do and whether its behaviour remains aligned with its intended purpose.

Multi-agent orchestration

When agents act together, risk no longer resides in a single identity. It emerges from the orchestration. Individually acceptable permissions can combine into excessive authority, delegated tasks can obscure accountability and one compromised agent can influence an entire chain of execution. Secure multi-agent orchestration must therefore verify every participant, constrain every delegation, evaluate cumulative

authority, isolate compromised agents and preserve a complete record from initiating intent to collective outcome. Agents may operate and collaborate autonomously, but any transfer of authority must be controlled and traceable.

Rapid cloud and multi-cloud adoption

Enterprise workloads increasingly span across cloud SaaS platforms and private-cloud environments simultaneously. As identities move across these trust boundaries, fragmented controls and incomplete inventories make it difficult to establish ownership, trace access and enforce consistent policies. The challenge is no longer simply managing identities within individual platforms; it is maintaining continuous visibility and governance across the entire multi-cloud identity fabric.

Data and model supply chain risks

Third-party models, datasets, plug-ins, tools and fine-tuned language models introduce inherited risks that may not be visible during deployment. These components can introduce unsafe behaviour, hidden dependencies, excessive access requirements, malicious instructions or compromised integrations into enterprise AI workflows. A trusted AI identity architecture must therefore govern not only the agentic identity itself, but also the models, data, tools, credentials and external services that shape its behaviour and authority.

A phased roadmap for action

Turning this strategic objective into operational reality requires a focused, phased approach:

01

In the first 30 days, know every critical agentic identity and contain immediate exposure.

02

By day 60, establish ownership, policy, baseline controls and bounded authority.

03

By day 90, enforce runtime controls, demonstrate that every consequential action is authorised, traceable and controllable and establish an audit-ready posture.

Measures that reveal real progress

Success should be measured not by the number of agents deployed, but by how securely and effectively they are governed. The following indicators reveal whether an organisation is making meaningful progress:

- 01  Percentage of production agents with a unique identity, named business owner, declared purpose and current risk tier.
- 02  Percentage of agent credentials that are short-lived, automatically rotated, audience-bound and protected against replay.
- 03  Percentage of high-impact actions subject to action-level policy and, where required, human approval.
- 04  Number and age of orphaned agents, unknown credentials, excessive privileges and unresolved ownership gaps.
- 05  Mean time to revoke an agent's authority and contain downstream sessions.
- 06  Percentage of agent actions that can be traced from initiating intent through delegation, execution and business outcome.
- 07  Rate of policy exceptions, behavioural anomalies, blocked actions and confirmed privilege drift.

Future state

The future of agentic identity security will be defined by autonomous, self-healing ecosystems that continuously detect risk, adapt access and remediate threats at machine speed. AI-driven analytics will identify abnormal behaviour before compromised agentic identity can cause harm, while dynamic policy engines will

adjust privileges in real time and automated certifications will prevent access creep as models evolve. A unified AI identity control plane will bring together identity discovery, risk assessment, lifecycle access management, policy governance, audit monitoring and advanced analytics

within a single architecture. As regulatory expectations continue to evolve, organisations that establish robust agentic identity security today will create resilient, accountable and auditable AI environments, laying the foundation for responsible innovation and sustainable business growth.

Recommendations

AI identities are dynamic, programmatic and capable of enterprise-scale impact within a single action. Governing them as conventional user accounts creates unacceptable exposure to credential sprawl, privilege abuse and compliance failure. Organisations must adopt an integrated approach combining AI-driven

risk analytics, zero-trust enforcement and automated lifecycle governance to secure AI deployments at machine speed and scale. As regulators demand demonstrable controls over AI pipelines, boards assume greater accountability, and attackers increasingly target machine credentials, agentic identity

security must become an immediate enterprise priority. Organisations that act now will build resilient, auditable and trusted AI environments that enable innovation without compromising control.

Conclusion

Securing agentic identity is more than a cybersecurity task; it is a foundation for trust. Agentic AI will not be secured by placing traditional controls around a new technology and hoping that existing assumptions hold. The enterprise must govern agents as actors: identifiable, owned, purpose-bound, minimally privileged, continuously evaluated and accountable for every consequential action.

The organisations that move first will not be those that give agents the most freedom. They will be those that create the clearest contract between autonomy and authority. Without it, enterprises cannot safely unlock the productivity gains of AI. By treating agentic identity as distinct, high-risk identities requiring continuous oversight, organisations can innovate with confidence. With an agentic identity control plane, innovation can scale without allowing control to dissolve.

The strategic objective is clear. Identify every agent, constrain its authority, verify every action, and preserve accountability from intent to outcome.

The next era will be shaped by enterprises that empower independent action while maintaining strong governance. The future of work is human-machine collaboration. The security strategy must reflect that reality.



01.0000:00
02.3000:00

About Deloitte

Deloitte is one of the world's largest and most diversified professional services organisations, providing assurance and advisory, tax, management consulting, and enterprise risk management services through more than 345,374 professionals in more than 150 countries. Our organization includes a unique portfolio of competencies integrated in one industry-leading organisation. Deloitte Touche Tohmatsu India LLP (DTTI LLP) also referred as Deloitte India is a member firm in India that provides non-audit consulting services. Our experienced professionals deliver seamless, consistent services wherever our clients operate.

In India, Deloitte is spread across 12 cities with over 12,000 professionals, who are proficient at delivering the right combination of local insight and international expertise to our clientele drawn from across industry segments. Deloitte is well-equipped to deliver solutions to the complex challenges faced by organisations across the public and private sectors. Our edge lies in our ability to draw upon a well- equipped global network and teaming this with customised services at a local office.

We have been consistently recognized as leaders by Gartner in Cybersecurity, the Data and Analytics space, as well for Public Cloud Infrastructure Managed and Professional Services and Oracle Cloud Application Services.

<https://www2.deloitte.com/in/en.html>

About Saviynt

Saviynt is the identity security platform built to secure the modern enterprise. With Zuma, its complete AI security platform, Saviynt helps organizations build, run, and scale AI with confidence. The world's leading brands trust Saviynt to reduce risk, improve visibility, and support compliance while enabling business agility.

For more information, visit www.saviynt.com.

Connect with us

Deloitte India

Sathish Gopalaiah

President, Consulting
Deloitte South Asia
sathishtg@deloitte.com

Gaurav Shukla

Partner and Leader - Cyber
Deloitte South Asia
shuklagaurav@deloitte.com

Anand Venkatraman

Partner
Deloitte India
anandv@deloitte.com

Ravindra Usmanpurkar

Partner
Deloitte India
uravindra@deloitte.com

Dhruv Khanna

Partner
Deloitte India
dhruvkhanna@deloitte.com

Saviynt

Nitin Varma

SVP & Managing Director
Saviynt India & SAARC
nitin.varma@saviynt.com

Snigdhodev Bhattacharya

Director Sales
Saviynt India & SAARC
s.bhattacharya@saviynt.com

Sanjeevi Kumar R

Director Sales (South & Sri Lanka)
Saviynt India & SAARC
sanjeevi.kumar@saviynt.com

Parna Das

Director Sales (West)
Saviynt India
parna.das@saviynt.com

Contributor

Deloitte India

Shashikant Shetty



Saviynt, the Saviynt logo, and other Saviynt marks are trademarks or registered trademarks of Saviynt, Inc. in the United States and other countries. Unauthorized reproduction, distribution, or transmission of the publication in any form is prohibited without prior written permission from Saviynt and Deloitte India, with standard exceptions for noncommercial uses and critical reviews under copyright law.

For more information, visit www.saviynt.com. | ©2026 Saviynt, Inc. All rights reserved

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited ("DTTL"), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.

© 2026 Deloitte Touche Tohmatsu India LLP. Member of Deloitte Touche Tohmatsu Limited