PingIdentity

Executive Brief

# Are You AI Ready?

# AI Agents Are Already Here

Agentic AI is a new, novel class of non-human identity that brings transformative potential and real operational risk. It also introduces a new customer channel and a digital workforce. As AI agents begin to act on behalf of users and systems, they challenge traditional identity, security, and governance models. If you want to lead the agentic economy, you'll need to drive agent-enabled innovation while enforcing strong identity, delegated access, data protection, and human oversight.

This brief provides an executive framework, a pragmatic readiness checklist, and a 90-day action plan to help enterprises operationalize AI readiness at scale.

## Why Identity for AI Should Be a Top Priority Today

### AI Agents Are Redefining How Businesses Operate

AI agents are becoming a new channel for customer engagement and revenue via agentic commerce and customer service assistants, as well as a new digital "workforce." Both require enterprise-grade identity, authorization, lifecycle, and governance to operate safely.

### AI Agents Are a Distinct Identity Type

AI agents are non-deterministic, capable of planning and acting, and often operate on behalf of humans. They can't be treated like traditional apps or service accounts, and they can't be treated like traditional human users.

### AI Agents Introduce a New Layer of Risk

Security risks include goal hijacking, over-privileged access, cognitive or memory corruption, deception, and resource exhaustion. Mitigation requires least-privilege design, auditability, human-in-the-loop controls, and continuous monitoring.

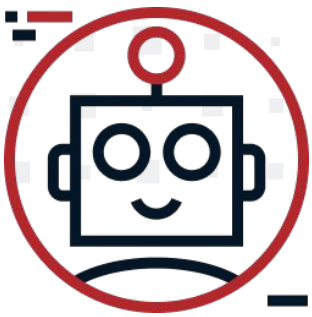

## Among Organizations Adopting AI Agents...

- **66%** report increased productivity
- **57%** report cost savings
- **55%** report faster decision-making
- **54%** report improved customer experience

*- PwC AI Agent Survey 2025*

*Source: https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html*

# What Is AI Readiness?

Being AI-ready is no longer aspirational—it's operational. For the C-suite, this means having the governance, security, and infrastructure to safely scale AI agents across customer and workforce environments.

It's not about the number of pilots launched. It's about whether your enterprise can support autonomous agents without compromising trust, compliance, or control.

To be truly ready, your organization must:

- **Enable agent-driven innovation** across customer journeys and internal workflows.
- **Enforce Zero Trust identity and access** for AI agents through delegated authorization and scoped, short-lived tokens.
- **Harmonize and protect data**, ensuring agents operate on accurate, policy-aligned information.
- **Prepare your workforce** to operate securely with AI, including oversight, approvals, and risk recognition.

## Start With the Basics: Know Your Agents

Securing agentic AI begins with understanding what types of agents your organization is enabling. Each behaves differently and demands specific identity and access strategies.

**Personal Agents (BYO)**
These are external agents—like ChatGPT or Gemini—used by customers or employees. Enterprises must detect and register these agents dynamically (e.g., via DCR), delegate access securely, and invoke human approvals for high-risk actions.

**Digital Assistants (Customer-Facing or Workforce-Facing)**
These enterprise-managed agents operate inside your trust boundary, either supporting customer interactions or assisting internal teams. In both cases, digital assistants must have registered identities and operate under tightly scoped permissions—even when acting on behalf of a user. To ensure consistent, secure access aligned to each assistant's role and context, you need centralized policy enforcement, token exchange, and service-to-service credentials.

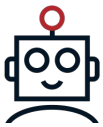

**Digital Workers**
These fully autonomous agents operate within your trust boundary. Like employees, they need their own credentials and lifecycle management—but with tighter controls, such as just-in-time (JIT) access and continuous monitoring.

# Your Framework for AI Readiness

Now that you know what your organization needs to be AI-ready, here's a proven framework for establishing a foundation to secure AI agents and unlock their potential without losing control.

## The 6 Pillars of Agentic AI Security

### Identity for AI

Agents should be first-class identities in their own right. **Classify** personal agents, digital assistants, and digital workers based on their autonomy and trust boundary position. **Govern** their full lifecycle with assigned custodians, registration, and regular reviews. **Detect** and tag agent activity, including GUI-based agents, to differentiate and monitor their behavior from human users.

### Delegation, Not Impersonation

Never share human credentials with agents. They must act through **authenticated delegation**, and **least privilege** should be the default. Use OAuth extensions like DCR, PAR, RAR, and token exchange to **define scope, consent, and context**. Apply **just-in-time** and **short-lived tokens** to reduce exposure, and always maintain a clear audit trail to distinguish agent and human actions.

### Human-in-the-Loop (HITL)

For sensitive or high-risk actions, enforce **explicit human approvals** using out-of-band flows like CIBA. These **approvals should be verifiable**—via biometrics or strong MFA—and logged to preserve accountability and support regulatory compliance.

### Interoperability and Control

Prepare for agents that interact vertically with tools (via **MCP**) and horizontally with other agents (via **A2A**). These protocols enable secure, policy-driven access to tools, data, and peer coordination, setting the foundation for scale and interoperability.

### Data Harmonization, Authorization, and Zero Trust

Ensure AI agents operate on **harmonized, policy-aligned data**. Use **relationship-based access control** (ReBAC) to enforce fine-grained authorization—especially for RAG workloads. Apply Zero Trust principles to govern access based on identity and context, not location or network perimeter.

### People and Operating Model

AI security is as much about people as it is about policy. Build a **role-based learning program** around privacy, delegation, and oversight. Anchor **governance** in the NIST AI Risk Management Framework (GOVERN, MAP, MEASURE, MANAGE) to ensure accountability, continuous improvement, and risk alignment.

# Check Your AI Readiness

As AI agents scale across the enterprise, they introduce new dimensions of risk, complexity, and accountability. Identity becomes the control point for enabling innovation without sacrificing security.

This readiness checklist outlines the foundational capabilities every organization must establish to safely scale agentic AI—spanning threat mitigation, data governance, delegated access, human oversight, and governance alignment.

Use this as a strategic gate before moving from pilot deployments to production environments.

## Threat Mitigation and Runtime Controls

| | |
|---|---|
| ☐ | Embed OWASP Top 10 defenses into your SDLC and runtime guardrails to address LLM and agent risks (prompt injection, insecure tool use, excessive autonomy, training data poisoning, and sensitive data leaks). |
| ☐ | Use authorization systems with external consistency to ensure ordered, enforceable content and access updates, especially across distributed systems. |
| ☐ | Enforce permissions-aware design for Retrieval-Augmented Generation (RAG) and other AI workloads. Filter documents and responses based on user entitlements to prevent data leakage. |

## Data Governance and Protection

| | |
|---|---|
| ☐ | Establish a cross-functional data governance committee to define shared business terms, data ownership roles, and change management processes across silos. |
| ☐ | Label sensitive data (PII, regulated, confidential) and enforce attribute-based or relationship-based access control (ReBAC) at ingestion and access layers. |
| ☐ | Enforce Zero Trust at every layer—API gateways, service meshes, and data pipelines—using short-lived tokens and explicit, auditable policies. |

## Human Readiness and Culture

| | |
|---|---|
| ☐ | Launch a role-based learning program aligned to NIST SP 800-50r1, covering agent-specific behaviors like human-in-the-loop approvals, secure delegation, and AI-driven risk recognition. |
| ☐ | Ensure employees understand how to work safely with AI agents, especially when they operate on behalf of humans or access sensitive systems. |
| ☐ | Go beyond completion rates—track improvements in approval quality, reduction in privilege creep, and adherence to delegated workflows. |

## Governance Alignment with NIST AI RMF

| | |
|---|---|
| ☐ | **GOVERN**: Define AI-specific roles, risk tolerances, and accountability structures. |
| ☐ | **MAP**: Inventory AI agent use cases, associated data flows, and regulatory constraints. |
| ☐ | **MEASURE**: Continuously assess AI performance, bias, security, and the effectiveness of HITL and least privilege enforcement. |
| ☐ | **MANAGE**: Remediate gaps, iterate controls, and tie improvements to measurable changes in risk posture. |

# From Pilot to Policy: A 90-Day Plan to Scale AI Security

Scaling agentic AI securely doesn't require boiling the ocean. But the sequence of events does need to be intentional.

This 90-day plan outlines a pragmatic path from initial pilots to institutional controls. Focus on quick wins early, then layer in scalable patterns and governance. This helps you move from experimentation to enterprise readiness without compromising trust, compliance, or agility.

| Phase | Use Case Activation | Security Controls | Data Governance | Workforce Enablement |
|---|---|---|---|---|
| **Days 0–30**<br><br>Establish the Foundation | Identify 3 priority agent use cases, one each (personal agent, digital assistant, and digital worker)<br><br>Map current identity, authorization, and data flows | Pilot permissions-aware AI for one RAG scenario, with pre/post-filtering by entitlements, short-lived credentials, and audit logging<br><br>Launch one HITL approval flow for high-risk actions using out-of-band approvals | — | — |
| **Days 31–60**<br><br>Expand Controls and Coverage | Extend secure delegation and JIT tokens to two additional AI use cases | Classify and register all agents<br><br>Tag agent sessions for enforcement and analytics | Begin sprint to define shared data terms and initial labeling of sensitive data | — |
| **Days 61–90**<br><br>Operationalize and Govern | — | Apply Zero Trust policies to service-to-service agent calls and data access | Finalize governance model and oversight for scaled agent operations | Launch learning program (SP 800-50r1) focused on HITL, delegation, and AI risk awareness |

# Your AI Readiness Scorecard

Use this scorecard to assess your organization's readiness to scale agentic AI.

Rate each capability from 1 (not started) to 3 (fully implemented) to pinpoint where further investment is needed before moving from pilot projects to enterprise deployment.

| Score<br>1 = Planned<br>2 = In Progress<br>3 =Ready | | | Capability | Description |
|---|---|---|---|---|
| 1 ☐ | 2 ☐ | 3 ☐ | Agent Identity Management | AI agents are registered, classified by type, and governed with lifecycle controls and custodianship. |
| 1 ☐ | 2 ☐ | 3 ☐ | Delegated Access | Agents use scoped delegation (not shared credentials), and short-lived/JIT tokens are enforced. |
| 1 ☐ | 2 ☐ | 3 ☐ | Human Oversight | High-risk agent actions require human approval, with audit trails and periodic review. |
| 1 ☐ | 2 ☐ | 3 ☐ | Permissions-Aware Design | Agent and RAG workflows are filtered by entitlements to prevent data leakage. |
| 1 ☐ | 2 ☐ | 3 ☐ | Protocol Readiness | MCP and A2A standards are being adopted to ensure interoperability and policy enforcement. |
| 1 ☐ | 2 ☐ | 3 ☐ | Zero Trust Enforcement | Identity-based, least-privilege access controls are enforced across services and clouds. |
| 1 ☐ | 2 ☐ | 3 ☐ | Data Governance | Sensitive data is labeled, protected, and harmonized across AI pipelines. |
| 1 ☐ | 2 ☐ | 3 ☐ | Workforce Training | Cybersecurity and privacy education includes agentic behaviors, delegation, and HITL. |
| 1 ☐ | 2 ☐ | 3 ☐ | AI Governance Framework | Governance follows the NIST AI RMF (GOVERN, MAP, MEASURE, MANAGE). |

**Add Your Scores and Rate Your AI Readiness**

- **22–27**: You're well-positioned to scale agentic AI with confidence.
- **16–21**: Invest in targeted gaps before expanding use cases.
- **Below 16**: Focus on foundational readiness—scaling too soon increases risk exposure.

# Trust Ping Identity with Your Agentic AI Plans

Ping Identity provides a comprehensive foundation for securing AI agents across your enterprise. From authentication and authorization to lifecycle governance, threat detection, and delegated access, we enable organizations to manage digital workers, internal assistants, and external personal agents with confidence.

Our platform supports just-in-time access, agent classification, and policy enforcement, so that you can securely adopt AI at scale, without compromising control.

## Explore the Ultimate Guide to Identity for AI

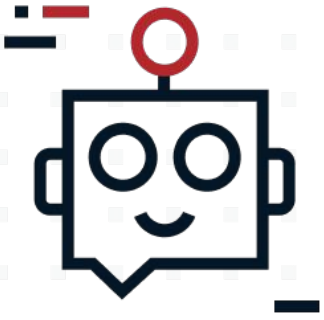

At Ping Identity, we believe in making digital experiences both secure and seamless for all users, without compromise. That's digital freedom. Ping enables enterprises to combine our best-in-class identity solutions with third-party services they already use to remove passwords, prevent fraud, enable Zero Trust, or anything in between. And they can do it all with a simple drag-and-drop canvas. That's why more than half of the Fortune 100 choose Ping Identity to protect every single digital interaction for their users, while making experiences frictionless. Learn more at www.pingidentity.com.

#4318 | 12.25 | v01