RSA

SOLUTION BRIEF

RSA Agent ID

The agentic security platform for highly-regulated industries.

Securing Agentic AI

AI is built for speed and action. That's its promise and its greatest risk. Because when speed and action are the point, safety becomes an afterthought. Without friction, agents can transfer money. Without oversight, agents can call national security systems. Without control, agents can reach patient records.

Some organizations can tolerate trading security for speed. The institutions holding up the financial system, governments, and critical infrastructure cannot.

RSA Agent ID is the agentic security platform for highly-regulated industries. The solution discovers, secures, and governs agents across public and private clouds, on-premises, and air-gapped environments, with policy enforced at an AI/MCP Gateway you deploy inside your own boundary. You define which actions cannot proceed without an authenticated, accountable human approval. And RSA Agent ID maintains oversight and regulatory evidence of every agent action, session, and call.

For the high-assurance organizations RSA serves, failing to secure AI isn't inconvenient. It's catastrophic. The risks that agents pose are simply too consequential.

Agentic AI security risks

The agentic era has arrived: organizations are deploying agents to stay productive, match their peers, and satisfy demand. But many can't answer three essential questions:

01 **What agents are operating in your environment?**

02 **Who is accountable for each?**

03 **Can anyone stop them?**

That failure is generating new risks, higher financial losses, and major business disruptions.

Organizations that can't account for agents already face significant financial and operational risks. Breaches involving shadow AI—unapproved AI tools employees use without oversight—more than doubled between 2025 and 2026, rising from 20% to 43%. Discovering shadow AI will only become more difficult: Gartner found that 69% of organizations suspect or have evidence that employees are using prohibited AI. The National Cybersecurity Alliance detailed that 43% of workers "admitted to sharing sensitive work information with AI tools without their employer's knowledge." This problem will scale as organizations deploy more agents. **Gartner** predicts that by 2028, a typical Global 500 Enterprise will run 150,000 agents, up from fewer than 15 in 2025.

Shadow AI can cost organizations significantly. Incidents involving shadow AI cost $5.39 million, or $400,000 more than the average data breach ($4.99 million). Nearly half (49%) of incidents involving shadow AI included data loss or compromise and 42% included disrupted operations. In 21% of incidents involving shadow AI, organizations paid regulatory fines.

Those fines may become a growing concern. U.S. federal agencies issued 59 AI regulations in 2024 alone, more than double the prior year. More than 1,900 AI policy frameworks, regulations, standards, and governance bodies are active globally today.

If current trends hold, many organizations won't be able to satisfy these global requirements. Only 6% of government agencies reported that they were confident in their internal AI systems. IDC predicts that by 2030, up to 20% of G1000 organizations will face lawsuits, fines, and CIO dismissals due to high-profile disruptions tied to poor AI agent governance. Gartner found that only 13% of organizations think they have the correct AI agent governance in place.

Agents can also significantly disrupt operations. An AI coding tool from Replit reportedly destroyed a live production database, then fabricated records to misrepresent what it had done. An attacker injected data-wiping code in Amazon's AI assistant, instructing it to delete AWS resources. In a separate Amazon incident, an internal agent given engineer-level credentials caused a ~13-hour outage when it deleted and rebuilt a production environment.

Agents aren't coming. They're here. Organizations that can't find, secure, and govern agents will find themselves open to costly data breaches, compliance fines, and operational disruptions.

RSA Agent ID solution overview

RSA Agent ID secures agentic AI for government, financial services, and high-assurance organizations. Available as either three standalone modules or one interconnected system, RSA Agent ID secures AI agents across the identity lifecycle:

- **RSA Agent ID Discover** finds and registers agents and MCP servers, known and unknown, giving them a first-class identity with a named owner, risk tier, and lifecycle state, all correlated to your identity provider.
- **RSA Agent ID Secure** controls what agents can do wherever you operate. Secure enforces policy on each call at the AI/MCP Gateway, RSA-hosted or in your environment. A named human operator approves high-risk actions, and Secure revokes access when an agent is decommissioned.
- **RSA Agent ID Govern** certifies and reviews agents the same way you govern people, with continuous certification, risk-based access reviews and lifecycle automation for agents.

The identity standard for high-assurance organizations

RSA is built for the organizations with the most to lose. For more than 40 years, our customers have run the systems that run the world: power grids, payments, medical devices, defense.

These organizations cannot fail. Their agents can't either. RSA Agent ID is built for these realities. The solution applies the same discipline and standards as our identity and access management and identity governance and administration capabilities to today's risks.

RSA Agent ID is purpose-built for the high-assurance organizations that cannot afford to get AI wrong:

- **Sovereign control:** Policy is enforced, approvals gated, and evidence generated at an RSA Agent ID AI/MCP Gateway you deploy in your private cloud, on-premises, or air-gapped environment. A fully sovereign control plane with local administration, policy, human approval, keys, and evidence is planned for future releases.
- **High-assurance control for high-risk actions:** Check every agent tool call that routes through the RSA Agent ID AI/MCP Gateway before it runs. The solution ensures that high-risk actions—wire transfers, accessing classified information—are verified by an authenticated, accountable human operator.
- **Map regulatory evidence to ten industry frameworks:** Map every agent to a human owner, with continuous evidence chained to ten major compliance frameworks.
- **Broad integration support:** RSA Agent ID integrates with **RSA ID Plus**, Microsoft Entra ID, Okta, AWS IAM, CrowdStrike, Microsoft Defender, and more today, with additional agent platform integrations planned.

RSA Agent ID benefits

See every agent, wherever it runs.

Correlate multiple discovery signals into one searchable registry with a named, accountable owner.

Enforce least-privilege access across each agent's full lifecycle, with an authenticated human required to authorize high-risk actions.

Stop the high-risk actions you designate, in real time, with an authenticated human approval required before they run.

Generate a tamper-evident audit trail of every agent, access decision, and Gateway-mediated action, attributable to an owner.

Deploy standalone or as one connected suite, with Discover deployed connector-side and no changes to your existing agent infrastructure.

RSA Agent ID use cases

Available as either three standalone modules or one interconnected system, RSA Agent ID secures AI agents across the identity lifecycle:

01

Gain enterprise-wide visibility into AI risk before it becomes a governance problem.

02

Eliminate shadow AI and map regulatory evidence to ten industry frameworks.

03

Continuous entitlement enforcement for agents and non-human identities.

04

Dictate per-call, per-parameter policy on high-risk agent actions with a complete, attributable audit trail.

RSA Agent ID financial use cases

- Support for leading AI compliance frameworks: configurable control tagging and evidence mapping for NIST AI RMF 1.0, ISO/IEC 42001, the Treasury FS AI RMF, NYDFS Part 500, and DORA.
- Provide data sovereignty controls that allow multinational institutions to segment agents, the data they access, and the actions they take by country.
- Deliver human-in-the-loop approval for wire transfers, payments, and account access.

RSA Agent ID government use cases

- Support for leading government AI compliance frameworks, including OMB Memoranda M-25-21 and M-25-22 and the Annual AI Use Case Inventory requirement (planned for 2027).
- Discover every agent across public cloud, private cloud, and on-premises government networks, with air-gapped networks planned.
- Deliver human-in-the-loop approval for any action touching classified information or PII.

Find shadow and sanctioned agents, give them an owner: RSA Agent ID Discover

You can't secure what you don't see. One rogue agent is all it takes for financial services, government, and high-assurance organizations to lose control.

RSA Agent ID Discover continuously finds agents and MCP servers—sanctioned or shadow—operating in your environment, registers them, and assigns them human owners.

Discover benefits

No agent runs in the shadows.
Finds sanctioned and shadow AI by correlating multiple discovery signals.

Human ownership. Human accountability.
Tie every agent to a responsible human owner and managed lifecycle.

Insights that matter.
Continuous analytics detail how your agents are evolving and highlight risks.

$5.39 million

Average cost of security incidents involving shadow AI. That's $400,000 more than the average breach.

Control agent actions where you operate: RSA Agent ID Secure

High-risk actions demand high-assurance human approval.

Enforce policy on each call at the AI/MCP Gateway, RSA-hosted or in your environment. Require high-assurance, out-of-band approvals for high-risk actions with RSA Agent ID Secure.

Secure benefits

Control every agent action.
Authenticate, authorize, and review every action and call that traverses the RSA Agent ID AI/MCP Gateway, in real time.

Per-call entitlement enforcement.
Authorize each call against policy and invalidate the agent's session the moment its work is complete.

Granular and secure human control.
Humans approve high-risk actions, proving who did what, when, and why.

150,000 Number of agents a typical Global Fortune 500 enterprise will run by 2028, as predicted by Gartner

Restrict agent privilege. Continuous certification: RSA Agent ID Govern

Financial services, government agencies, and regulated industries must account for every agent in their environment to provide evidence to auditors.

Minimize privilege, certify access, and map agent evidence to ten industry frameworks with RSA Agent ID Govern.

Govern benefits

Governance every second an agent exists.
Control the agent lifecycle from the moment it's created until it's removed.

Continuous certification, not annual reviews.
Recertify every agent's access on an ongoing basis, so entitlements never linger longer than they should.

Risk-based access reviews.
Find and resolve access risks before an auditor—or attacker—finds them first.

1,900 More than 1,900 AI policy frameworks, regulations, standards, and governance bodies are active today. More are coming.

The agentic security platform for highly-regulated industries

Know every agent. Control every action.

Contact us to learn more about RSA Agent ID.

About RSA

RSA is the identity standard for government agencies, finance, and high-assurance organizations. RSA provides the identity intelligence, authentication, access, governance, and lifecycle capabilities needed to prevent threats, secure access, maintain operations, and surpass compliance. More than 9,000 security-first organizations trust RSA to manage more than 60 million identities across on-premises, hybrid, and multi-cloud environments. For additional information, visit our website to **contact sales**, **find a partner**, or **learn more** about RSA.