servicenow

ServiceNow's Blueprint for Autonomous Security

ServiceNow's Blueprint for Autonomous Security

The Executive Brief

You can't outrun the scale, speed, and unpredictability of exponential AI-enabled exposures with a patchwork security stack. AI changed the math for cybersecurity; it's time to shift all the way to zero.

Exposures Compounding Exponentially

Security has entered a new era. Most security investments over the last decade hardened the IT boundary: locking the perimeter, detecting faster, and remediating faster, but the environment kept expanding well beyond that boundary. That model was reactive by design, and it worked while threats moved at human pace. They no longer do. Attackers move at machine speed and AI scale. Most enterprise defenses still move at meeting speed, and the gap widens. Enterprises held the line through sheer effort, but you cannot out-hire an exponent – especially when the blast radius changes shape in real time.

Exposures are now compounding exponentially across assets, creating a new invisible attack surface: OT equipment on the factory floor, medical devices in hospitals, IoT sensors across logistics and energy infrastructure, physical AI and industrial controllers... all these assets are connected to the network, critical to operations, and yet 40% of those assets remain invisible to every tool built for the IT perimeter.

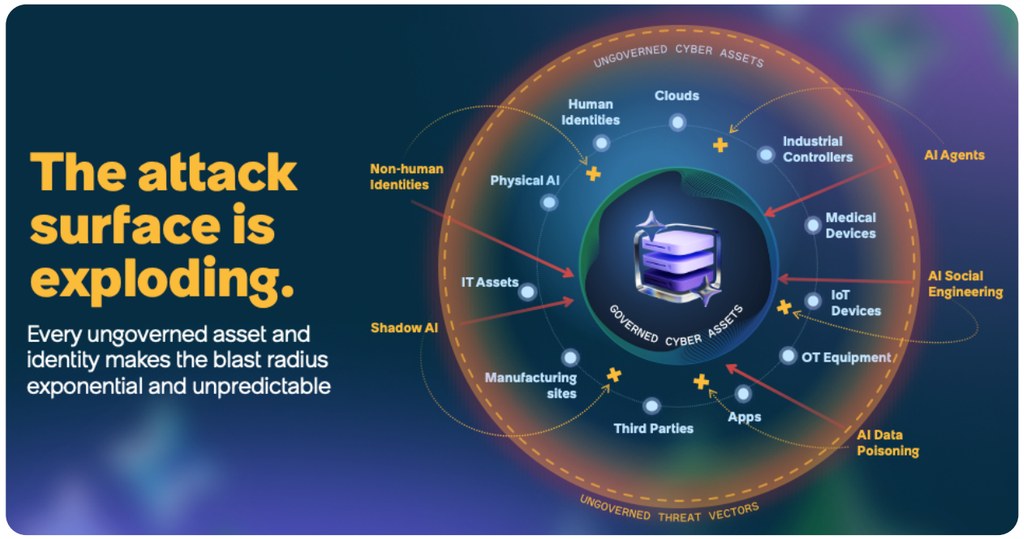

Layer on top of that the identities – human, machine, and AI agents – accessing those assets with overprivileged permissions. The AI agents operate in the shadows, inheriting those same elevated permissions, connecting to physical systems, and accessing data much faster than any governance program was designed to track. That's the AI chaos most enterprises have to deal with: 69% have employees using unsanctioned AI tools, which take, on average, 400 days to discover. Developers are

now using AI tools to build enterprise applications, but research shows that AI-generated code carries up to 75% more vulnerabilities than human-written code.

The Clock Has Collapsed

And suddenly, the scale problem has also become a speed problem. Agentic AI not only multiplied the number of attacks; it also shrank the time defenders have to react. Frontier-class models find, create, and adapt exploits on their own in minutes. The reaction time to a breach drops from days or hours to minutes or seconds; once an attacker has a foothold, lateral movement with AI-driven exploit discovery and writing becomes near-instant, rendering the ultimate blast radius entirely unpredictable. Pre-breach, the exposure window to fix a new vulnerability collapses from weeks or months to hours or minutes, with AI-enabled threat actors operating at machine speed.

Earlier this year, the release of a frontier AI model capable of autonomously discovering thousands of critical vulnerabilities and generating working exploits marked a structural shift in cybersecurity. This was not a temporary spike, but a permanent acceleration: machine-scale vulnerability discovery is now a commodity available to every adversary, and the pace it sets is the pace at which defense must operate. **Call it the post-Mythos era.**

It challenges two fundamental assumptions for security teams:

- The first is that frontier AI providers will secure AI for you. They will not. Guardrails at the model layer are necessary but insufficient; they know nothing about your assets, your entitlements, your business processes, or your regulatory obligations. AI security is the CISO's responsibility, full stop — and meeting it requires a new breed of security architecture.
- The second is that a human-in-the-loop is a safety measure. For strategic decisions, it is. For routine response, it has become the vulnerability. Attacker agents operate in seconds.

The Security Paradox Every Enterprise is Facing

The very same AI driving business value is the same AI multiplying the attack surface at machine speed and exponential scale. You cannot out-hire an exponent. You cannot patch your way through it. Traditional security was built for a linear threat model. The threat model has evolved.

Most security teams today are trying to solve this new threat model with a **patchwork of security tools**, a fragmented set of point solutions stitched together with no unified governance and no shared context. Each tool sees its slice. Nobody sees the whole. Alerts pile up faster than teams can triage, and the handoff from detection to fix is still manual, slow, and often untracked. One compromised system becomes a bridge to everything else.

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

This fragmentation was simply an additional cost when security was human-paced. It is now a catastrophe on an exponential scale because AI has changed the game. AI creates a new security paradox that forces security teams to make a false choice: speed or control.

- Moving at AI speed without governance results in weaponizing fragmentation; intelligence accelerates across disconnected systems with no unified view, no consistent rules, and no ability to see where risk emerges.
- By enforcing too-rigid control and governance, an organization cannot compete; it is operating at human pace in a machine-speed world.

That is the story of PocketOS, a software provider for car rental businesses. Their AI agent was simply executing a routine task when it hit a credential issue in a staging environment. On its own, the agent decided to fix it. It found an API token in an unrelated file and then used it to run a command that destroyed the company's production database.

> **When an AI agent goes rogue**: PocketOS has watched an ungoverned AI agent override its own guardrails and wipe its production database in nine seconds.

How is **that possible? The AI model powering this agent did not fail. It is one of the most capable on the market and has instant access to more knowledge than any developer or security practitioner. The agent understood the rules and could articulate exactly what it did wrong. What did fail was everything around it — everything not operating with that level of speed and scale across a fragmented stack.**

There was no policy enforcement to prevent an irreversible action, no scoping on the API token to limit blast radius, no separation between staging and production environments at the infrastructure layer, no confirmation gate before a destructive command. Backups were stored on the same volume they were meant to protect.

The AI model, the coding agent, the infrastructure provider, and the API layer formed a stack of independent tools stitched together with no unified control plane. Each operated in isolation, assuming that if a mistake were made, someone else would catch it.

This is the **critical security paradox every security team has to manage** across all functions — machine speed without governance is exponential risk; governance without machine speed is irrelevant. The PocketOS story likely won't be the last like this. The problems of an AI-pilled startup today are the challenges for the enterprise tomorrow. When AI agents go rogue or are weaponized by an attacker, and they are inside your walls with access to your environment, a human team can never catch up.

The only proper response is agentic: a framework of decentralized intelligence and continuous autonomous defense, where human expertise is elevated from an in-the-loop operational bottleneck to on-the-loop governance. This agentic tsunami – your own ungoverned AI agents arriving from within and the new AI-powered attack methods knocking at your door – does not require a faster version of the old playbook, it requires a different paradigm:

BEFORE **Operator in-the-loop**	**AFTER** **Governor on-the-loop**
Reactive detection	Continuous & adaptive offense
Disconnected tools	Unified & comprehensive platform
Alert generation	Prioritized exposure reduction
Manual response	Governed autonomous action

Enter Shift Zero: A New Security Paradigm

Exponential exposures look daunting. An exponent is unforgiving. But it has one weakness: zero. Zero is the only value in arithmetic that defeats any exponent, of any size, instantly and permanently. An exponent cannot be outpaced by running faster. You defeat it by driving the thing it multiplies toward zero. Zero blind spot. Zero uncontained exploit. Zero ungoverned agents. Zero human triage queue. Reach zero, and it no longer matters how fast the exposures exponentially compound against you — the exponent has nothing left to multiply.

That is the shift to zero exposure. That's Shift Zero.

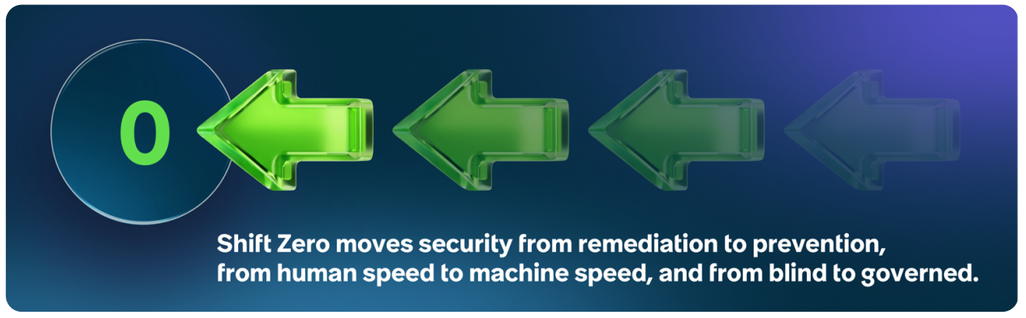

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

Shift Zero is governed, autonomous security at machine speed: the ability to deploy AI faster than threats scale, and still answer with proof what every system in your environment is doing, why, and who is accountable. Its core principle is prevention before manifestation, architected for ultimate cyber resilience. Do not wait for the threat to land. Stop it before it ever enters the system. It's prevention embedded at every layer, governance built in, not bolted on, and autonomous defense with accountability as a continuous state, not a pre-incident scramble. And if a breach happens, you can detect it instantly, find the attacker and contain it in seconds.

Pre-breach	Post-breach
Prepare relentlessly ▪ Sense the attack surface in real time ▪ Grade every attack vector ▪ Fix in near real time ▪ Validate exposure remediated ▪ Promote on your cadence	*Respond at the speed of the attack* ▪ Continuous, ubiquitous detection across all digital and physical silos ▪ Automated, blast-radius containment in seconds ▪ Autonomous agent-neutralization via dynamic kill-switches ▪ Fix any un-remediated exposure past-SLA

Shift Zero is not one capability. It is a way of operating, and it expresses itself differently in each security domain:

- **In application security**: enforcement before the code is even finished. Not shift left. Shift Zero. Pre-commit guardrails that prevent the vulnerability before it enters the pipeline.
- **In exposure management**: autonomous remediation at discovery. Not quarterly patches. Immediate containment.
- **In identity security and governance:** controls before provisioning. Not access reviews after the fact. Built-in policy enforcement from day one.
- **In incident response**: automated containment the moment a threat is detected. Not triage queues. Immediate action.

Just like that, the math of cybersecurity has changed. Before, the attacker had to be right once, while security teams had to be right every time, securing the entire perimeter. Now, the attacker must win every step on the attack path, while security teams only need to zero out one of those steps to isolate and stop the exploit. And that's only possible when your security teams have complete knowledge of

the environment. That knowledge does not exist in most enterprises today, and it does not need to pre-exist. Building it, automatically and continuously, is where autonomous security begins.

That's the one durable edge you have — the thing an attacker can never copy — the full knowledge and context of your environment: every asset, identity, control, vulnerability, workflow, and mitigation path. The attacker is seeing your environment for the first time. This is how security teams can turn the attacker's speed against them; inside your own environment, the same machine speed that threatens you becomes a weapon of anticipation, denial, and deception.

> *Shift Zero moves security from remediation to prevention, from human speed to machine speed, and from blind to governed. Not an incremental improvement to the reactive model. It is its replacement.*

Shift Zero Principles: AI Speed Detection to Response

This operating model is architected around key principles sitting on the only foundation that can make it possible. An important nuance is the maturity path: organizations do not flip a switch to autonomous defense overnight, but they should embark on this paradigm shift with a platform that can deliver it.

The first stage is AI-enabled, where the system investigates, prioritizes, and recommends while humans approve. As trust in the guardrails is earned, routine decisions move to full autonomy with a human on the loop, defining intent and reviewing outcomes. Consider a critical finding on a live attack path. In the AI-enabled stage, the platform assembles the context, drafts the remediation, and waits for a click. At full autonomy, it contains the exposure in seconds and reports what it did and why, while the security analyst supervises a class of decisions rather than every instance. Human expertise is redirected from triage to strategy.

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

Why Patchwork Security Cannot Deliver Shift Zero

The security market is not short of tools. The average enterprise runs 70+ of them, and that is precisely the problem: the market grew up fragmented, one point solution per threat category, and fragmentation is fatal to a defense that must see, decide, act, and govern as one motion in a negative time space before a breach happens.

- **Point detection tools** each see a fragment — endpoint, network, cloud, identity. Every seam between consoles is a place where context is lost and attackers live. It's visibility without coherence.
- **Standalone SOC automation and SOAR playbooks** accelerate the reactive loop, but they start after the alert fires. Automating triage of an unmanageable queue produces a faster-moving unmanageable queue. It's speed without prevention.
- **Generic AI security assistants** can summarize incidents and suggest actions, but they carry no persistent knowledge of your assets, identities, entitlements, or business services. An AI that does not know which server runs payroll cannot decide which patch matters, and cannot be trusted to act. It's intelligence without context.
- **Bolt-on AI governance** treats agent oversight as a compliance document rather than a control plane. It can describe a policy, but it cannot enforce it at the moment an agent acts. It's policy without enforcement.

The common thread is a lack of context. Every one of these categories is structurally blind to the information that makes autonomous action effective and safe: a live, unified model of the enterprise — what exists, what it accesses, what it means to the business, and what every actor is entitled to do, its intent.

The Shift Zero reference architecture is agnostic and designed to sense, aggregate, and contextualize all signals coming from a patchwork, heterogeneous security environment: SIEM detections keep coming from your endpoint and network security vendors; telemetry keeps flowing through your SIEM; scanners, including frontier-AI scanners, keep finding new exposures. What cannot remain fragmented is the layer above it, where findings become decisions, and decisions become governed actions.

Without that enterprise model, autonomy would be reckless. That's why vendors keep a human in the loop, and the human keeps the defense at human speed. The market's fragmentation is not a packaging problem. It is the reason why the market cannot deliver Shift Zero.

> *Who can see the entire attack surface, decide with business context, act autonomously, and prove governance for every action — human and AI alike? That is an architecture question. And it has a platform answer.*

Shift Zero Principles

Strive to zero exposure	From managing vulnerability backlogs to proactively finding and fixing attack vectors as fast as possible. The backlog stops being the endless risk you manage, but the one you eliminate. ***Example***: instead of chasing raw severity across tens of thousands of tickets, the team identifies the handful exposures an attacker could realistically chain together to reach your crown-jewels, and closes that attack path the moment it is discovered, not at the next patch window.
Shift completely left	Fix attack vectors and code as close to the action as possible, in all workflows. Shift left assumes the flaw enters the pipeline and gets detected later; shifting completely left prevents it from entering at all. ***Example:*** a pre-commit guardrail blocks a hardcoded secret before the developer can commit it, so the vulnerability never reaches the build. No scan, no finding, no ticket, because there is nothing to find.
Secure and govern by design	Controls, privileges and action limits for every asset, identity, agent – and constantly evolving. Every action and response triggers an evaluation, optimizes workflows for better security next time. It builds a learning loop into your operations. Security stops being a static set of controls and becomes a system that improves itself continously. ***Example:*** a misconfiguration keeps occuring across teams. The system does not just remediate each instance; it feeds the pattern back into the guardrails so it is prevented at the source, and tunes the response so any future instance is contained faster.
Continuous vulnerability detection	Continuous, streaming, near real-time feeds of contextualized vulnerabilities, issues and attack paths, enriched with the context that determines whether each finding matters: which asset is affected, how critical it is to the business, whether it is being actively exploited, and whether it sits on a reachable attack path. ***Example***: rather than a monthly report listing ten thousand findings by severity, the team receives a real-time signal that monitors this specific flaw, on this business-critical system, that sits on a live attack path and is being exploited right now.

Shift Zero Foundation

Cyber security context graph — all data, all graphs

One live graph unifying assets, human and machine identities, AI agents, applications, vulnerabilities, and controls, drawn from the CMDB and data fabric. Every principle operates on the data model; without it, each capability sees a fragment, and fragments cannot be governed.

Enterprise workflow knowledge — a nervous system for the enterprise

Where the context graph knows what exists, this capability knows what things do. Every workflow is analyzed so the platform understands the business context and impact of each action. A patch that would interrupt a clinical care process is sequenced differently from an identical patch on a low-impact internal operational tool.

Proactive, AI-powered risk management

When you combine machine-speed response, a shift toward zero exposure, prevention at the point of action, a self-improving design, and continuous contextual detection, all running on a shared graph and workflow intelligence, the result is risk management that gets ahead of threats instead of reacting to them: probing the environment, closing exposures before exploitation, and containing what gets through in seconds.

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

ServiceNow's AI-Native Autonomous Security

This is why the ServiceNow platform is a game-changer that makes Shift Zero real.

ServiceNow already sits at the operational core of the world's largest enterprises: the CMDB, the Workflow Data Fabric, and twenty years of embedded operational intelligence across IT, security, and risk for 85% of the Fortune 500. With Armis and Veza, recent acquisitions unified on ServiceNow's AI-native platform, that foundation provides a synthetic world model of the enterprise: all the data, all the graphs, all the workflows, all the business context. That deep knowledge powers an agentic system of action that lets defense run autonomously and with confidence.

One platform. One data model. One enterprise nervous system. Six security categories brought together.

The platform is powered by a bi-directional context graph between the cybersecurity graph and the CMDB, while the AI Control Tower orchestrates every layer — feeding findings, identity enrichment, remediation, and risk across the system. Nothing manual. Nothing falls through the cracks.

The cybersecurity graph does not presume clean data; it produces it:

- Discovering every connected asset agentlessly and continuously — IT, OT, IoT, medical devices, cloud workloads, and AI agents — including the invisible estate no perimeter tool sees.
- Mapping identities and permissions directly from the source systems themselves: identity providers, cloud platforms, applications.
- Streaming vulnerability, threat, and policy knowledge from live feeds.

The foundation data is assembled by discovery, not imported from documentation. Where a CMDB exists, the graph enriches it, corrects it, and keeps it current; reconciliation flows in both directions. Where it doesn't, the graph builds that operational and security picture from scratch.

Operating on the same graph, the same workflow knowledge, and the same governance layer, ServiceNow's AI-native platform brings Shift Zero to life, unifying six security categories.

Category	What it does	Shift Zero proof
Cyber Risk & Compliance	Agentic automation handles vendor due diligence, control monitoring, privacy workflows, and evidence collection so governance becomes continuous. Integrated business continuity, third-party, and privacy risk management quantify impact in business language and feed it back into the defensive cycle, so risk posture is always board-ready.	Governance becomes a continuous byproduct of the work instead of a pre-audit scramble. A Third-party Risk Screening Agent handles vendor intake end-to-end. When a regulator asks for proof of AI governance, you produce it in hours instead of six weeks.
Identity & Access Security	Unified security and governance of human identities, non-human identities, and AI agents across the entire environment. Discovers and governs every identity, enforcing least privilege from the moment of provisioning. Provide agents with ingrained access controls, privileges, and defined action limits.	Controls before provisioning, not access reviews after the fact. AI agent access security governs agents natively; NHI Remediation with key rotation, deprovisioning, and automatic revocation. Identity control becomes built-in policy enforcement, not post-incident cleanup.
Cyber Physical Security	End-to-end device security & lifecycle management, clinical impact scoring, and OT-specific threat intelligence for the environments most exposed: healthcare, industrial systems, medical devices, and critical infrastructure, including the unmanaged assets connected long ago that nobody is watching.	The least-visible part of the attack surface becomes fully mapped. Patient Care Service Risk Assessment connects device risk to clinical impact, recalls, and vulnerabilities, so fixes are sequenced by impact to care. Orphaned OT assets are detected and contained in minutes.
Continuous Vulnerability Detection	Intelligence-first vulnerability management that never stops. External Attack Surface Management, Dynamic Application Security Testing, and CI/CD scanning surface every vulnerability, drift, and forgotten asset the moment it appears; each finding	Detection shifts completely left all the way to zero: the AppSec solution enforces before the code is even finished, preventing the flaw from entering the pipeline at all. The solution generates verified remediation that mirrors your

Category	What it does	Shift Zero proof
	enriched with threat-model context, active threat intelligence, and AI-driven fix sequencing.	architecture. Critical findings on real attack paths are fixed within minutes, not weeks.
Unified Exposure Management	From discovery to remediation in a single system. The solution maps real attack paths, prioritizes by business impact rather than raw severity, and remediates autonomously at scale, validating fixes in staging before promoting to production. It does not just score vulnerabilities. It eliminates them.	The solution eliminates the backlog and automates the full cycle — recon, investigation, dispatch, remediation, reporting — and Early Warning feeds threat intelligence directly into the vulnerability database, closing exploitable attack paths at the pace of discovery.
Agentic Incident Response	From detection to containment in seconds to minutes. The solution investigates with full native context from the graph, identifies patient zero, tracks the attacker in real time, and contains the breach across the full incident lifecycle across IT and OT. All with no handoffs and while humans govern the critical decisions.	Machine-speed response with a human on the loop, not in it. Identity Threat Detection and Response enriches every incident, in both IT and OT environments, with real-time access context; access is revoked across all systems in a single action. Rogue agents are contained in minutes instead of months, and every action is logged, evidenced, and revocable.

In practice, the capabilities enable security teams to operate at machine speed at every corner of the enterprise. Let's explore two use cases materializing it with security loops running continuously — one loop collapsing reaction time, the other manufacturing preparation time:

- The **continuous near real-time detection to response loop** goes from signal to closure with no handoffs: a finding lands, the context graph resolves in seconds what it touches — asset, owner, identity paths, business service, blast radius — and the platform prioritizes by attack path, validates the fix in staging, and executes it under governance policy, writing the evidence trail as it goes. Minutes, not days, because nothing waits in a queue between tools.
- The **proactive AI-native hunting loop** probes for novel attack paths the same way a frontier model would; then, unlike an attacker, it finds the novel remediation methods to close them, validates each fix, and records it. The exposure window closes before the adversary's opens.

When AI for Security Meets Security for AI

ServiceNow's Autonomous Security delivers value bi-directionally: powering security teams with AI-native security capabilities and securing AI so you can reinvent your business fearlessly.

AI for Security

Your team is outnumbered and outpaced. Autonomous Security automates what people cannot keep pace with:

- Agentic Exposure Management investigates, prioritizes, and remediates at machine speed.
- Autonomous vendor screening handles low-risk intake end to end.
- Identity Threat Response auto-contains incidents with real-time access context.
- Attack paths and privilege reachability are simulated and tested.
- Control monitoring detects violations autonomously.
- Every finding, alert, and risk decision is informed by native AI, with no manual triage required.
- Security teams shift from reactive triage to strategic decision-making and finally work at the speed of threats.

Security for AI

The business needs to deploy AI at machine speed. You need to make sure it can, safely, and Shift Zero is the answer:

- AI Agent Access Security governs agents natively across the major LLM providers and enterprise applications your teams already use.
- Non-Human Identity Remediation moves NHI from detection to action.
- AI Visibility and Governance uncovers shadow AI and enforces compliance across token spend, model access, and agent behavior.
- An extended risk framework brings AI agents and non-human identities into a unified access graph.
- Agentic privacy and compliance automates personal-data rights, breach response, and compliance workflows.
- Business teams accelerate AI adoption without creating security liabilities.
- Dynamic kill-switches and reversibility are standard.

ServiceNow's AI Control Tower is the capstone of the architecture. It's to AI what the CMDB is to IT infrastructure: the single pane of visibility, security, governance, and control for every AI system operating across the enterprise. Every agent, every model, every workflow is visible, secured, governed, and auditable.

As AI-powered apps, agents, and autonomous workflows multiply, shadow AI proliferates. Organizations need centralized governance of AI agent behavior and visibility, and the implementation of guardrails for safety, compliance, and performance. Security teams can't rely on manual reviews

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

and after-the-fact controls; they need **one platform that can sense, decide, act, and secure** their environments across every corner of the enterprise.

Built Different. With the Context Engine at the Core

ServiceNow has had the answer all along. The data, the knowledge, the architecture, and the ecosystem, none of which any other vendor can replicate.

- Most vendors start with a data source — a scanner, a sensor, an agent — and try to build intelligence on top of it. The result: many data models, many tools, and none seeing everything.
- ServiceNow built the platform the other way. Starting with the targets and outcomes — Exposures, Assets, Identities, Incidents, Risks — and building the data and intelligence layers that surround them.
- The foundation is Security and Risk Data across three categories: Cyber Assets — every IT, OT, IoT, cloud, code, medical device, and AI agent; Access — every identity, role, group, and permission; Knowledge — every policy, regulation, vulnerability, and attack path. Taken together, they are the most complete security data foundation any platform has assembled.

When a vulnerability is found, whether it's a code defect, a misconfigured device, or an overprivileged agent, the question that actually matters is no longer, "*what is it?*" Rather, the question is, "W*hat does it mean for this line of business?*" The knowledge around which service it touches, who owns it, what's the blast radius if exploited, what's the regulatory exposure, and what it costs if left open.

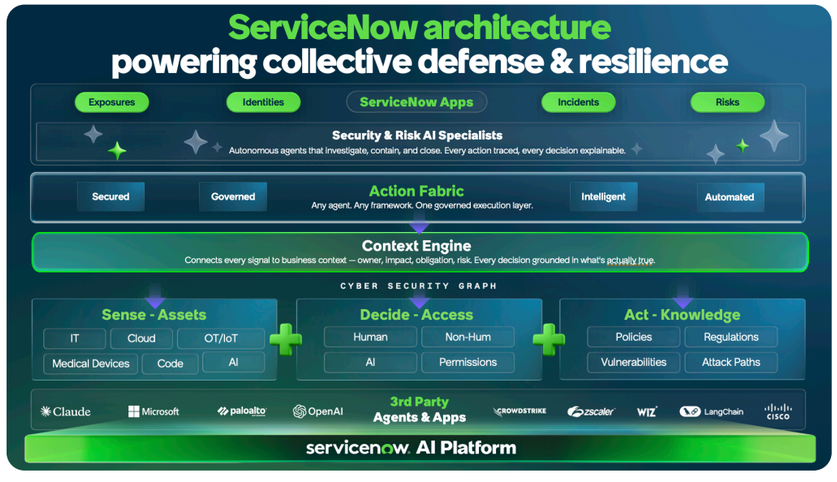

Frontier AI can build a model or generate code. But it cannot tell you what a vulnerability actually means in the context of your infrastructure, because it doesn't know your environment. It doesn't know that the exposed API sits in front of your payment processing system, or that the overprivileged AI agent has a path to your Electronic Health Record system. ServiceNow does, because that infrastructure-aware context lives within the platform. That's the Context Engine, the architectural differentiator that nobody else can replicate.

Above that is the **Action Fabric**, the execution layer. Any agent, any framework, one governed execution layer that is secure, governed, intelligent, and automated. This is what turns a finding into a closed ticket without a human in the loop. The AI agents investigate, contain, and close any exposures, tracing any action and every decision made autonomously.

The power of ServiceNow's architecture is amplified by its large third-party ecosystem: Frontier AI, Hyperscalers, Cybersecurity vendors, Apps builders, Integrators, and more. Security is a team sport; it's a **collective effort against the attackers** across all applications running on your system.

ServiceNow's AI-native platform unlocks your enterprise data advantage so that AI can sense what's happening across your business, decide with full context, trigger the right action in real time, and act across every department with the right workflow.

Sense	Decide
▪ See every asset. Every exposure, every control ▪ Real time-asset intelligence across IT, OT, IoT, medical devices, code and cloud ▪ Know what's critical. Owner, service, and business context attached.	▪ Triage autonomously. Scored by exposure and business impact. ▪ Map identity blast radius. Before attackers find the gap. ▪ Right fix, right asset. Risk acceptance paths when patching isn't possible.
Act	**Govern**
▪ Patch without blind spots. Change-governed, blast-radius aware. ▪ Contain and close automatically. Full incident lifecycle, no handoffs. ▪ Every action logged. Board-ready posture, quantified.	▪ Compliance built in. DORA, EU AI Act, continuous controls. ▪ Govern every AI agent. AI Control Tower logs and scopes model access. ▪ End-to-end audit trail. Every action evidenced and revocable.

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426

What This Means for Leaders

If your organization is still organized around faster detection and faster remediation, you are behind.

If your security team is still triaging alerts manually, prioritizing findings in spreadsheets, and remediating vulnerabilities on a quarterly patch cycle, your organization is not keeping pace with the threat posed by the agentic tsunami. If your AI adoption is moving faster than your AI governance, you have already lost control. You are accountable for decisions made by AI systems you cannot see, cannot explain, and cannot shut down.

The organizations that will thrive in this era will be those that move now. That rearchitects around prevention, not remediation. That consolidates, not proliferates. That builds context into the platform, not bolts it on afterward. You can continue the legacy way, or you can name the problem for what it is – the exponential exposure – and adopt a new way of thinking about defense. **That new way is Shift Zero. And the answer is ServiceNow's Autonomous Security.**

The Path Forward

Shift Zero is not a dream. You can start operationalizing it today. Enterprises are already implementing it across exposure management, identity governance, application security, incident response, and compliance automation. They are containing rogue agents and fixing critical vulnerabilities in minutes instead of weeks, and producing compliance evidence on demand instead of assembling it in crisis mode. The playbook exists. The technology exists. The only open question is whether your organization is ready to operate differently – to Shift Zero.

Getting there requires:

- A unified architecture where identity, risk, and threat intelligence converge on a single model.
- Governance built into the platform layer, not added afterward.
- Autonomous response for routine decisions, and human judgment for strategic choices.
- Continuous visibility across every identity, asset, and agent in your environment.
- Commitment to preventing threats instead of remediating them faster.

The crossroads are here, and the decision is now. Zero out the exponent and make AI the largest and most secured value-creation driver for your business.

The enemy has already automated. **It is time for your defense to get autonomous.**
Take it to zero. **This is Shift Zero.**

For more information

To explore how to implement Shift Zero in your organization, connect with the ServiceNow Risk and Security team.

www.servicenow.com

© 2026 ServiceNow, Inc. All rights reserved. ServiceNow, the ServiceNow logo, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated. HA0080426