servicenow

ServiceNow's Blueprint for Agentic Business

The Technical Foundation

servicenow

ServiceNow's Blueprint for Agentic Business

The Technical Foundation

servicenow

Table of Contents

Purpose of This Document

As AI intelligence trends toward commoditization, the durable competitive advantage shifts to platforms that can apply it with the right context, governance, and execution infrastructure where enterprise work actually happens. This paper provides a technical assessment of that market shift and explains why ServiceNow's architecture is uniquely positioned to lead it.

1. The Problem: Architectural Fragmentation Meets the AI Moment

The modern enterprise runs on an average of 367+ applications across the employee experience alone, each with its own data model, security perimeter, API surface, and governance logic. Beneath the application layer sit thousands of connected assets (IoT, OT, cloud infrastructure, medical devices) that most organizations cannot fully inventory, let alone govern.

This is a platform problem, not a management problem. Every department optimized for its own domain: Salesforce for sales, Workday for HR, Zendesk for support. The result is deep but disconnected systems with no shared operational model. Data is siloed. Workflows terminate at application boundaries. Security policies are enforced inconsistently across integration points.

Over the past 12–18 months, markets have reacted dramatically to every major AI breakthrough. DeepSeek alone wiped out nearly $1 trillion in tech value. Fears around AI agents, digital workers, and vibe coding triggered another ~$1 trillion selloff across software stocks in a single week, with the sector down nearly 30%+ from recent highs.

The market narrative is clear: if AI can write code, automate tasks, and spin up agents, what happens to enterprise software? But the market is confusing functional replication with enterprise readiness. An AI agent can generate a ticketing form. It cannot replicate decades of business process modeling, security architecture, regulatory compliance, deep system integrations, and operational resilience.

The first wave of enterprise AI adoption confirmed this. Despite record investment, enterprise AI maturity declined 20% year over year. The models aren't the problem. The foundations are. Vendors bolted AI assistants onto existing applications as sidecars, producing shallow intelligence layered over disconnected process. The models can reason. They cannot execute across systems with governance, context, and audit-grade accountability.

The core architectural issue: the current generation of AI solutions lacks access to unified operational data, sits outside the flow of work rather than inside it, falls short of enterprise security and compliance requirements, and forces organizations to redesign processes around the AI instead of augmenting what already works.

The market draws a provocative conclusion: "AI will eat software." Some feature-specific software, yes. Workflows with rich business context? No.

AI doesn't replace enterprise orchestration. It depends on it.
Without workflows, AI is just expensive advice.

What enterprises need is Data, AI, Workflows, and Security operating on a single platform with a shared operational model, unified governance and context, and the ability to execute across every domain. ServiceNow's AI Platform architecture addresses this by unifying Data, AI, Workflows, and Security on a single platform with a shared operational model and consistent governance. The sections that follow detail the technical foundations of that architecture and why it creates structural advantage.

2. The Key Insight: Applied AI vs. Raw Intelligence

The competitive landscape is defined by a fundamental asymmetry: frontier AI models can be built, improved, and even commoditized within training cycles measured in months. Enterprise operational context (the workflows, integrations, data relationships, and institutional knowledge required to apply AI reliably) compounds over years and cannot be accelerated with compute alone.

The closer you get to the operational core of an enterprise, the more complexity, governance, reliability, and context matter. That's precisely where ServiceNow has a structural advantage:

- 80B+ workflows, 6.5T transactions executed annually (growing at ~25% YoY) across IT, HR, customer service, security, and business functions
- 20+ years of embedded workflow intelligence at the operational core of mission-critical operations
- AI embedded in real-time business processes, operating within the flow of work
- Enterprise visibility with AI Control Tower (AICT), positioned as the neutral party between hyperscalers and AI startups
- Flexibility, choice, and ecosystem openness: any cloud, any hyperscaler, any LLM, deployable on-prem, private cloud, or public cloud, enabling data sovereignty and residency compliance

ServiceNow already sits at the operational core of these enterprises. 85% of the Fortune 500. 98% renewal rate for 20+ consecutive quarters. This isn't a commercial metric — it's deep operational embedding. The platform where work already happens is the platform best positioned to make that work autonomous.

A useful analogy: a GPS helps an individual optimize a route. Powerful but localized. ServiceNow operates like air traffic control, coordinating thousands of moving parts, enforcing safety and policy

constraints, routing work across teams and systems, and maintaining a real-time operational view. One assists an individual. The other manages the system.

This distinction captures the core strategic positioning: AI that assists individuals versus AI that runs enterprises. But the point isn't that individual AI assistance lacks value. It's that enterprise AI must go beyond it. When autonomous workflows handle the routine, people gain the autonomy to do the work they actually signed up for. That's the real promise: putting AI to work *for people*.

The Economic Reality: Intelligence Commoditizes. The AI Control Tower Compounds.

The cost of intelligence has dropped by an order of magnitude in the last three years. Gemini Flash, a highly capable model for many applications, is about four times cheaper than GPT 5.2. Claude Sonnet commands a premium for coding prowess, but the benchmarks show Gemini 3.1 Pro and GPT 5.2 Codex are closing fast. The gap will narrow. So will the cost differential.

Intelligence is already cheap. The durable value accrues in the thing that can safely *execute* in the enterprise: the governed platform that knows who you are, what you're allowed to do, and how to do it across systems, with audit-grade proof.

As model intelligence continues improving and becomes broadly available, application features become cheaper to create and easier to copy. Enterprises don't pay durable multiples for features that can be reproduced quickly. They pay for the institutional AI Control Tower that enables safe execution:

1. **Governance OS:** Permissions, workflow enforcement, audit trails.
   Any agent must run through it.
2. **Best domain agents (on-platform):** The model isn't special. The platform's operational data and execution context are.
3. **Data flywheel:** Every action, exception, resolution, and failure becomes compounding operational history that improves automation, risk detection, and workflow design.

This produces a counterintuitive but critical insight: AI agents need the platform more than humans do. Humans have intuition about boundaries. They know not to peek at colleague compensation, to get approvals before payroll changes, to watch deadlines. A powerful agent can do far more than a human, but that amplifies risk unless the platform supplies the guardrails. The more capable the agent becomes, the more it depends on identity resolution, entitlements, workflow constraints, integration governance, audit evidence, and change management.

The Open Platform Advantage: Any AI, Any Data, Any Cloud

The AI Control Tower only works as the enterprise orchestration layer if it is genuinely neutral. A control tower locked to one hyperscaler or one model provider is just another silo.

ServiceNow is architected for openness at every layer.

Any AI model: enterprises can leverage any LLM provider (NVIDIA, OpenAI, Google, Anthropic, ServiceNow's own) grounded in their enterprise context and governance. AI Agent Fabric provides connectivity for agents to work together across tools, teams, and vendors. Any data: Workflow Data Fabric and Zero-Copy Connectors federate live queries across customer data warehouses, data lakes, and transactional systems without replication or vendor lock-in. Any cloud: deployable on-prem, private cloud, or public cloud across any hyperscaler, with configurable data residency and sovereignty controls.

Hyperscalers want to consolidate workloads. Data platforms want to centralize data. Model providers want to standardize on their APIs. Each optimizes for its own layer. ServiceNow optimizes for the enterprise, orchestrating across all of them. That neutrality is what makes the AI Control Tower trusted.

"ServiceNow is destined to be the best platform, the operating system of enterprise AI agents."

Jensen Huang, CEO, NVIDIA

3. The AI Native Architecture for Sense, Decide, Act, Govern

True autonomous workflows at enterprise scale are not powered by a model alone. They require a tightly integrated stack of data, execution infrastructure, intelligence, and governance operating together. ServiceNow is uniquely positioned because we already own and operate this stack at the center of enterprise operations.

Four interconnected capabilities combine to power the AI Control Tower, enabling autonomous workflows that sense context, decide the right action, execute work across the enterprise, and govern every step. Humans define the boundaries. The platform enforces them. Autonomous workflows operate within those boundaries at speed and scale no human team could match.

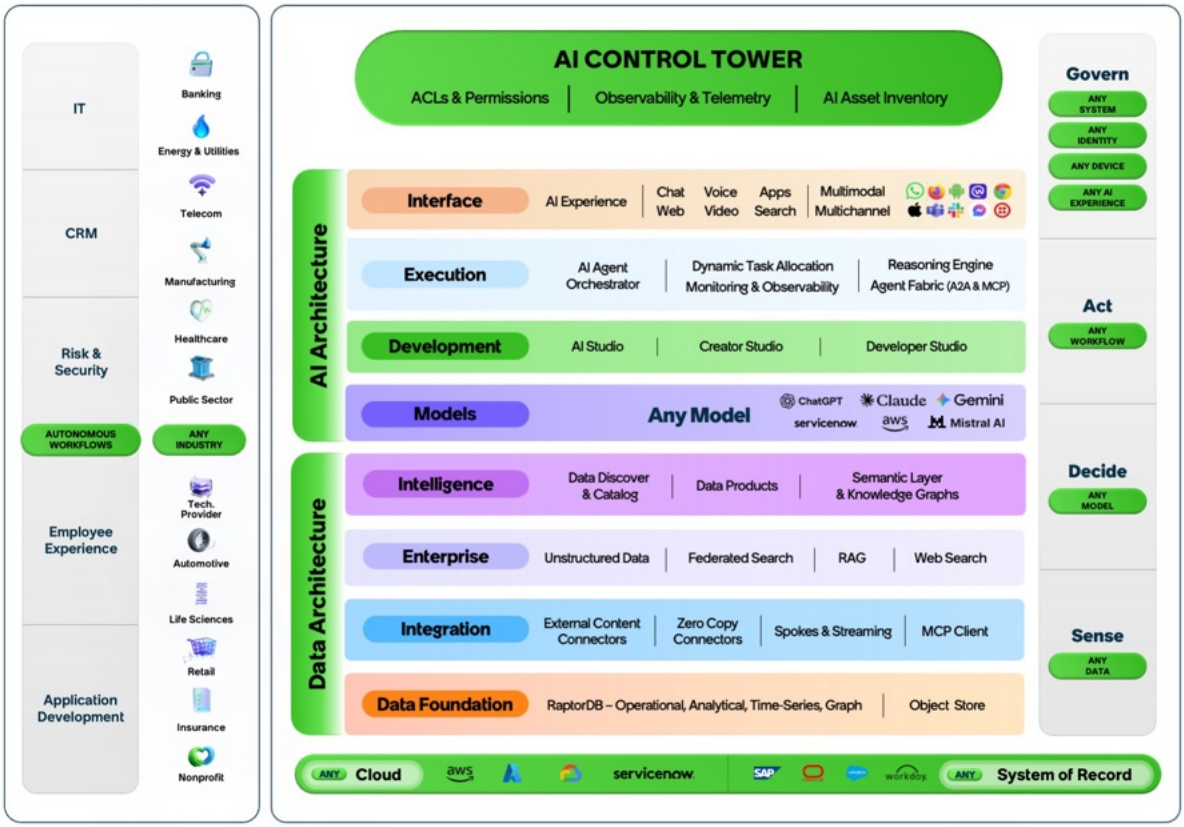

3.1 SENSE — Unlock Your Enterprise Data Advantage

Most LLMs give you models trained on the publicly available data. ServiceNow gives you enterprise context.

ServiceNow has deep, read-write agency into the digital fabric of organizations, providing a continuously updated model of assets, services, dependencies, ownership, and operational state.

Data Models: The Foundation of Enterprise Operations

ServiceNow's architecture is built on the Common Service Data Model (CSDM), a unified relational schema connecting digital assets, business services, people, locations, contracts, and issues in real time. This creates a live operational model of enterprise digital infrastructure and relationships rather than static data snapshots.

The CMDB (Configuration Management Database) alone represents an extraordinary investment. It maps every digital asset to the business services it supports, enabling impact analysis when changes occur. When a network switch fails, ServiceNow can automatically identify which business services are

affected, notify the right teams, trigger failover procedures, and create incident records. The data model encodes these relationships in real time.

RaptorDB Pro: Real-Time Analytics on Operational Data

RaptorDB provides an HTAP (Hybrid Transactional/Analytical Processing) database engine that runs transactional and analytical queries on the same live dataset. The read-replica architecture scales analytical workloads near-linearly without impacting transactional performance.

Other SaaS platforms either limit analytical queries per fixed tenant size or extract data to separate analytical stores. Both introduce latency and operational/analytical inconsistency. RaptorDB eliminates this: analytics run directly on operational data with zero lag.

Native graph and time-series capabilities in the same engine allow AI Agents to query transactional records, run analytics, and traverse relationships in a single system without data movement or replication delays. Real-time analytics on live operational data at scale. No ETL, no secondary stores, no data drift.

This is the speed and scale AI agents need to reason and act on billions of records in real time.

Workflow Data Fabric: Real-Time Federation Over Operational Systems

A federated data layer that queries and combines data across enterprise systems without replication or centralization. Supports point-to-point APIs, streaming ingestion, virtualized access, and real-time federated queries.
With Zero-Copy Connector (ZCC), ServiceNow federates live queries across both ServiceNow-resident data and customer data warehouses, data lakes, and transactional systems. This provides a unified view of the customer's entire data estate for operational decisioning without data movement. No other SaaS vendor supports this.

Traditional Data Platforms are processing engines requiring storage-layer access (Iceberg/Delta tables). This creates two problems:

a. **Logic duplication.** Business rules and security must be replicated in the data platform. This is extremely likely to be incorrect.
b. **Stale data.** Sources expose replicated snapshots, not live operational state.
With ServiceNow's federated architecture, queries execute against live source systems with native business logic and security intact. No duplication, no replication lag, no logic drift.

Data Catalog, Lineage, and Semantic Layer (powered by data.world)

Provides structured understanding of enterprise data for AI agents through cross-domain discoverability, provenance tracking, semantic business mapping, and relationship modeling.

ServiceNow embeds data.world's catalog natively within the platform where workflows execute. It discovers all data assets across the customer's entire data estate (platform-resident and federated) with relationships, lineage, provenance, and governance unified where operational decisions happen. No other SaaS vendor integrates catalog capabilities directly into their operational platform.

Traditional Data Platforms catalog replicated data in their warehouses. ServiceNow catalogs live operational data in real-time, where business decisions execute and actions are taken. Trusted, governed data context flows directly to AI agents operating on live systems. No catalog-to-execution gap, no stale metadata, no external tooling required.

The Sense capability extends to connected assets: IoT, OT, cloud, medical devices, identities, AI agents. Continuous discovery that knows what exists and how it's connected. Vendors without ServiceNow's history in enterprise IT cannot replicate this.

Vertical and Industry Depth

Enterprise context is not generic. How a hospital triages a critical system failure is fundamentally different from how a bank responds to a compliance exception or how a manufacturer handles an unplanned production stoppage. The workflows, regulatory constraints, and data models are distinct. Getting them wrong has real consequences.

ServiceNow's advantage extends into deep vertical expertise refined through thousands of production deployments:

- **Healthcare and Life Sciences:** Clinical workflows, medical device management, HIPAA-governed data handling, and patient safety protocols. AI coworkers operate within care-specific compliance boundaries that generic platforms cannot encode.
- **Financial Services:** Trade operations, regulatory reporting, SOX controls, and multi-jurisdictional compliance. The RSU example in this paper illustrates the cross-system complexity that is routine in this industry.
- **Government and Public Sector:** FedRAMP-authorized deployments, classified environment support, and citizen service workflows with strict data residency requirements.
- **Telecommunications:** Network operations, field service orchestration, and customer lifecycle management at scale, where IT and OT converge and a single outage can affect millions of subscribers.

- **Manufacturing:** Production line monitoring, supply chain coordination, and quality management workflows that bridge IT systems with operational technology on the factory floor.

This depth compounds over time. Each deployment adds industry-specific workflow patterns, exception handling, and regulatory knowledge into the platform's operational intelligence. Competitors must recreate this from scratch in every industry they enter. ServiceNow's AI Agents inherit it at deployment.

IT-OT Convergence and Cyber-Physical Asset Intelligence

Enterprise AI cannot sense what it cannot see. Traditional IT platforms discover and manage digital assets: servers, applications, endpoints. But the modern enterprise also runs on billions of operational technology assets, IoT devices, medical equipment, industrial controllers, and building systems that sit outside conventional IT visibility. Most organizations cannot fully inventory what's connected, let alone assess what's exposed.

Upon closing, Armis will further extend ServiceNow's Sense capability into this cyber-physical domain with agentless discovery and continuous monitoring across every connected asset type — managed, unmanaged, IoT, OT, medical, and industrial devices — feeding real-time asset intelligence directly into the CMDB and Context Graph. AI Agents will be able to sense across the full operational environment, not just the IT layer.

When paired with ServiceNow's business-context CMDB, which maps assets to the services, processes, and teams they support, Armis will create a continuously updated map of the entire enterprise environment. For industries where IT-OT convergence is accelerating (manufacturing, healthcare, energy, telecommunications), this delivers a unified asset model that no other platform provides.

Global Infrastructure and Data Sovereignty

Enterprise AI cannot operate where it cannot comply. Data locality expectations have intensified globally, and regulatory frameworks now require data, compute, and AI inference to reside within specific national boundaries. Beyond locality, many jurisdictions impose requirements around who can access data, how it is operated, and how infrastructure is owned and managed — such as the US Government's IL4 and IL5 requirements, or Australia's IRAP framework.

ServiceNow's global network of data centers and hyperscaler deployments enables customers to run within their own regulatory jurisdiction across dozens of countries today. This is a structural differentiator. AI-native startups and task-level agent frameworks lack sovereign infrastructure entirely. Hyperscalers offer regional deployment but cannot provide the unified workflow execution,

governance, and data model that ServiceNow delivers within each jurisdiction. Model providers offer intelligence — they do not offer compliant, governed execution on sovereign soil.

For autonomous workflows at enterprise scale, this is a prerequisite. AI Agents that sense, decide, and act on enterprise data must do so within the data residency and regulatory boundaries their organizations are subject to. Enterprise agentic solutions must operate where the enterprise operates, adhere to the regulations the enterprise operates within, and meet those needs today, not tomorrow.

ServiceNow does.

3.2 DECIDE — Ground AI in Business Reality

To enable autonomous workflows, AI models need to reason with business accountability and explainability. By grounding any AI model provider in enterprise context, rules, and knowledge, ServiceNow enables decisions that are aligned, predictable, and auditable.

ServiceNow focuses on making models safe for your business: aligned with enterprise policy, permissions, and operational reality.

Contextual AI Grounding

CMDB, Context Graph, and workflow history give AI real-time awareness of who is involved, what assets are affected, what policies apply, what happened before, and what should happen next.

- **Real-time Enterprise Knowledge Graph** connects people, roles, assets, services, and workflows into a unified model that AI can dynamically query for ownership, dependencies, and operational state.
- **Semantic Layer across systems** integrates CMDB, workflow data, and enterprise systems to ground AI decisions in real-time operational context rather than isolated data.
- **Continuous operational learning** is automatically updated through system activity, creating a compounding institutional knowledge layer that improves over time.
- **Historical context and continuity** stores past resolutions and execution paths, enabling AI to maintain continuity across long-running processes and resume with full context.
- **Pattern-based intelligence** learns from patterns, exceptions, and outcomes at scale, allowing AI to reference what worked, what failed, and what's typical.
- **Enhanced reasoning and impact analysis** links data to real enterprise relationships for better next-best-action selection and impact analysis based on operational dependencies.

Multi-Agent Orchestration

While Claude Cowork and OpenClaw are opening up multi-agent workflows on individual machines, ServiceNow already applies this model at enterprise scale, grounded in real processes, governance, security, and execution infrastructure.

ServiceNow has been operating with coordinated, workflow-driven multi-agent systems in production since March 2025. ServiceNow was also the first enterprise software vendor to add production support for MCP & A2A (in May 2025).
This is a modular agentic architecture designed for enterprise complexity, not task-level automation scaled up.

3.3 ACT — Execute Autonomously Across the Enterprise

Everyone else delivers AI that thinks. ServiceNow delivers AI that thinks and acts without compromising trust. Platform capabilities like Agent Orchestrator, Agent Studio, and Agentic Playbooks execute work end-to-end, from autonomous IT helpdesks to updating CRM records based on customer signals. The difference between AI assistance and autonomous execution.

The Workflow Engine: Operational Intelligence Built Over 20 Years

ServiceNow orchestrates enterprise operations through deterministic, policy-driven workflows that execute reliably across IT, HR, customer service, security, finance, and legal domains. These workflows coordinate human decisions and automated actions while enforcing governance, compliance, and security at every step.

Enterprise execution requires consistency and control that AI which generates answers probabilistically alone cannot provide. ServiceNow bridges this gap by combining AI reasoning with deterministic workflow execution, ensuring actions follow defined policies, maintain audit trails, and produce reliable outcomes even as AI explores creative solutions.

Enterprise AI requires both the probabilistic intelligence of AI combined with the deterministic controls of workflows.

Unified Operational Data Model

All workflows execute against the same relational schema (CSDM). An incident record references real configuration items, business services, ownership chains, and policy constraints in real-time. Workflows inherit this context natively.

When a network switch fails, the platform automatically identifies affected business services, notifies correct teams, triggers failover procedures, and creates incident records. The data model encodes these relationships as first-class objects, not external lookups.

Competitors operate workflows as isolated automation layers sitting on top of disconnected systems. ServiceNow workflows execute within the operational model where state, relationships, and policy already exist.

Governance as Native Execution Architecture

ServiceNow's workflow engine treats governance (approvals, SLAs, compliance gates, audit trails, rollback) as native execution primitives. Flow Designer, Process Automation Designer, and Integration Hub execute within the same governed runtime. RPA scripts, external system calls, and AI actions inherit platform-level security and policy enforcement automatically. Competitors layer governance on top of workflow engines, requiring separate security configuration at each integration point. This creates governance drift: different automation layers enforce slightly different versions of the same business rules.

ServiceNow enforces policy at the execution layer. Every workflow step runs through the same policy engine and ACL framework. AI agents can operate autonomously at scale because policy violations prevent execution before they happen. Competitors must choose between agent autonomy and enterprise control. ServiceNow provides both.

Workflows as Executable Agent Tools

ServiceNow AI Agents invoke existing workflows as high-level operations rather than low-level APIs. A workflow like "provision user" already encodes the full execution path: create AD account, update HRIS record, route manager approval if above threshold, configure access based on role, send notifications, log audit trail.

Competitors expose system APIs, forcing agents to orchestrate these steps on every invocation. The agent must infer approval hierarchies, system update sequences, and rollback procedures. This probabilistic orchestration introduces variance: agents hallucinate outdated thresholds, omit critical steps under token limits, or invent sequences that violate compliance. Hard-coded prompts miss dynamic organizational state and context.

ServiceNow workflows are pre-built enterprise operations. Twenty years of production deployments have encoded common patterns (user provisioning, approval routing, incident resolution, change management) as tested workflows that integrate with CMDB, HRIS, and asset data. Competitors force customers to build these operations from scratch against raw APIs. ServiceNow agents inherit decades of encoded operational knowledge as proven, context-aware execution primitives.

servicenow

The Autonomous Workforce

ServiceNow's Autonomous Workforce extends the human workforce with AI coworkers that operate alongside people within existing workflows. These AI coworkers handle routine, high-volume work so teams can focus on higher-value decisions and outcomes.

They operate with:

- **Role-Scoped Permissions (ACL Engine):** Same role hierarchy governing 80B+ workflows. RBAC-enforced boundaries. Agents cannot self-escalate.
- **Platform-Persisted State (RaptorDB):** Context and decision history persist in the transactional data layer, not model context windows. No session resets.
- **Enterprise Grounding (CMDB + Context Graph):** Live queries against asset relationships, ownership, policies, and resolution patterns at execution time.
- **Cross-System Orchestration (Integration Hub):** 500+ certified spokes and MID Servers with bi-directional, policy-aware connectivity.
- **Audit and Transactional Integrity:** Immutable audit trails, transactional rollback, inline SLA and compliance enforcement.

Developer Enablement: Build Agent + App Engine

While most AI tools help developers generate code, ServiceNow AI helps them generate enterprise-ready applications. Because the platform is grounded in existing workflows, data models, permissions, and integrations, what gets built is immediately deployable within the enterprise. Developers move from idea to production faster, building on top of decades of business process capital rather than recreating security, governance, and orchestration from scratch.

When an out-of-the-box workflow doesn't exist, App Engine and Build Agent let anyone build new workflows with AI, inside the guardrails of enterprise security, governance, and policy. Innovation moves fast without creating risk.

servicenow

What this looks like in practice, ServiceNow vs. building from scratch:

Dimension	Without ServiceNow (From Scratch / Vibe Coding)	With ServiceNow
Time to Production	Fast to prototype, slow to harden; months–years to enterprise readiness	Rapid path from idea to production on an enterprise-ready platform
Security, Identity & Compliance	Must build RBAC, SSO, audit trails, and regulatory controls from scratch; high risk of gaps	Built-in roles, permissions, auditability, and enterprise-grade compliance by default
Infrastructure, Scale & Resilience	Must design for hosting, performance, failover, monitoring, and recovery	Proven infrastructure, scalability, and operational resilience already in place
Data, Integrations & Context	Requires stitching together APIs, systems, and data models across the enterprise	Deep integrations, enterprise data models, and system context already embedded
Workflows & Cross-Department Processes	Must custom-build approvals, SLAs, orchestration, and coordination across teams	Native workflow engine spanning IT, HR, Finance, and operations out of the box
Cost & Risk to Finish	Easy to start, expensive and risky to productionize and maintain at scale	Build on existing process capital, reducing risk and accelerating time to value

How It All Comes Together: The Autonomous IT Resolution

Consider a common enterprise scenario: an employee can't connect to VPN and submits an IT ticket saying, "*I can't access internal apps.*"

What a standalone LLM does: It can suggest steps: restart the VPN client, check network connectivity, reinstall software, verify credentials. It might ask smart follow-up questions. But it's operating without real context. It doesn't know who the employee is, what device they're using, what access they should have, whether there was a recent security policy change, or whether a certificate just expired. Most importantly, it cannot take action. It can guide, but it cannot resolve.

What ServiceNow does through Sense, Decide, Act, Govern:

Sense: When the ticket is created, an AI worker embedded inside the ServiceNow workflow immediately understands the full context. It knows the employee's role, device, location, access profile, and recent change history. It can see that the employee's VPN certificate expired overnight after a scheduled security rotation. It checks device telemetry, confirms there's no broader outage, and determines the issue is isolated.

Decide: The AI determines the right remediation path based on policy, context, and historical patterns.

Act: Within policy boundaries, it triggers a certificate refresh, pushes the updated configuration to the employee's device, revalidates access permissions, and tests connectivity. Once the connection is restored, it verifies the employee can access internal applications again.

Govern: The resolution is documented. The audit trail is complete. Policy compliance is verified. The pattern is logged for future learning.

The entire process happens in minutes. No human intervention. Complete accountability.

Stand-alone LLM	ServiceNow
Ticket → Advice → User follows steps → Maybe resolved	Ticket → Sense (user, device, access, history) → Decide (policy, context) → Act (remediate, verify) → Govern (audit, learn) → Resolved

Beyond IT: Why Regulated, Cross-System Workflows Expose the Deeper Gap

The VPN example illustrates autonomous IT resolution. But the real test comes when money, regulated data, and multiple systems of record are involved.

Consider a real scenario: an employee checks their brokerage account on RSU vesting day and the share count looks wrong. Not wildly off, just enough to trigger concern. *Where are my shares?*

What doesn't work: A standalone LLM explains RSU vesting, walks through withholding logic, pulls policy summaries, but can't answer the actual question: "Why is my distribution wrong, and can you fix it?" Meanwhile, a Cowork-style desktop agent notices something suspicious (a withholding election that appears to have flipped) but can't verify why because it lacks access to the enterprise system of record. So it guesses. It blasts multiple teams in parallel, starts drafting messages containing personal stock transaction details, and creates noise, risk, and embarrassment fast. The LLM can explain but can't act. The desktop agent can act but can't govern. Neither can resolve.

What ServiceNow does through Sense, Decide, Act, Govern:

Sense: The system resolves the employee's identity in the enterprise sense. Which employment entity, equity plan version, tax jurisdiction, payroll provider, stock plan administrator, and brokerage are involved. Those facts live as structured, relational data across HR, payroll, finance, tax, and equity administration, connected by relationships that change over time and are full of edge cases. A lived-in data model built from thousands of real-world scenarios. Commodity intelligence can read and reason. It can't produce clean, governed, current enterprise truth.

Decide: The system enforces permissions *below* the model. If the AI tries to compare the employee to similar cases (a reasonable analytical approach in isolation), enterprise rules prevent it. Compensation data is governed by field-level access controls, role-based visibility, need-to-know boundaries, and jurisdiction constraints that differ by workflow, role, and circumstance. The system literally cannot show or do what it's not allowed to.

Act: "Where are my shares?" becomes a governed workflow spanning multiple vendors: HRIS, payroll, stock plan admin, brokerage, and the general ledger. The platform investigates root cause, validates corrections, coordinates across vendors, enforces segregation of duties, and tracks SLA compliance. Each API call executes as the right identity, under the right policy, for the right purpose, with the right audit trail. Calling an API is easy. Governed, federated execution across vendor boundaries is the hard part.

Govern: The enterprise must produce correct shares *and* evidence: who initiated the case, what changed where and when, who approved what, which systems were touched, SLA compliance, exception handling. For regulated workflows, the audit trail isn't documentation. It's the product. The issue turned out to be a cross-entity transfer that changed the employee's tax profile and created a temporary mismatch between systems. The AI could explain how vesting works. The platform could reconcile records across vendors safely, with approvals, and with a trace of what changed where.

Standalone LLM	Local Desktop Agent	ServiceNow
Explains RSU logic → User still stuck	Guesses → Blasts teams → Creates risk	Sense (identity, entity, plan) → Decide (policy, permissions) → Act (federated execution across vendors) → Govern (audit trail, approvals, evidence) → Resolved

The VPN example shows what autonomous IT looks like. The RSU example shows why the platform advantage deepens as workflows cross systems, involve regulated data, and require federated execution with governance. The more complex the work, the more the enterprise needs the AI Control Tower.

From Architecture to Outcomes: The Five Autonomous Solution Areas

Sense, Decide, Act, and Govern are the platform architecture. The outcomes they enable map to five solution areas that ServiceNow delivers today, each representing autonomous workflows applied to a major enterprise domain:

Autonomous IT shifts from reactive firefighting to AI that anticipates, resolves, and secures issues before they disrupt the business. **Autonomous CRM** reimagines customer experience from lead to resolution with AI agents that sell, serve, and support across front, middle, and back office. **Autonomous Employee Experience** eliminates the manual grind of fragmented systems, so

employees spend less time navigating and more time delivering value. **Autonomous App Development** turns application creation from a months-long IT project into a days-long creative act — low-code and AI-assisted development inside platform security, governance, and compliance, where every app inherits enterprise-grade trust from day one. **Autonomous Risk & Security** unites SecOps and GRC for the AI era, using AI to continuously assess risk, remediate vulnerabilities, and apply governance in real time.

Each is powered by the same Sense/Decide/Act/Govern architecture. The technical depth in this paper explains *why* they work. The solution areas are *where* they deliver.

3.4 GOVERN — Trust at Enterprise Scale

Sense, Decide, and Act are only as powerful as the governance layer that holds them together. AI Control Tower is that layer: the integration point for the entire architecture, not a compliance function bolted on after the fact.

Autonomy cannot exist without strong governance. ServiceNow provides enterprise-grade controls that make AI safe and deployable in high-risk, regulated environments. The human role shifts from executing work to governing the system that executes it, defining the policies, permissions, and boundaries within which autonomous workflows operate.

AI Control Tower (AICT): The Capstone of the Architecture

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

As AI-powered apps, agents, and autonomous workflows multiply, shadow AI proliferates. Organizations need centralized governance for AI usage and behavior, visibility into how AI is acting across the enterprise, and guardrails for safety, compliance, and performance. Security teams can't rely on manual reviews and after-the-fact controls. They need the same power and speed the rest of the business is gaining.

AICT delivers this and positions ServiceNow as the neutral orchestration hub that manages any AI ecosystem (internal, partner, or third-party), making it central to how enterprises operationalize AI at scale. One platform orchestrating Sense, Decide, Act, and Govern across every corner of the enterprise.

Enterprise Security and Compliance

Key capabilities:

- Role-based access control
- Policy enforcement frameworks
- Full auditability of AI decisions and actions
- Compliance alignment
- Data segmentation and protection
- Human-in-the-loop oversight
- SOC 2 Type II certification with immutable audit trails for every action
- Multi-regulatory compliance architecture (GDPR, HIPAA, FedRAMP, SOX, CCPA) with configurable data residency and retention policies

This is the governance foundation autonomous workflows require.

Identity Governance for the Agentic Enterprise

Modern enterprises must manage a diverse array of identities: employees, partners, systems, applications, devices, and increasingly, autonomous AI agents. Traditional identity and access management tools were designed for human accounts. They struggle with machine identities, API keys, and autonomous agents making real-time decisions, creating blind spots around privilege sprawl and cross-system access paths.

As enterprises deploy hundreds of AI coworkers across departments and systems, identity governance becomes a critical control surface. Every AI coworker needs scoped permissions. Every action needs an auditable identity chain. Every access decision must enforce least privilege in real time, not through periodic human reviews.

As part of ServiceNow, Veza will bring its patented Access Graph into the AI Control Tower, mapping and analyzing access relationships across human, machine, and AI identities. With over 30 billion access permissions under management across Fortune 500 and global enterprises, Veza provides the trusted identity layer that AI coworkers need to act safely at scale. Enterprises will be able to centrally govern, monitor, and enforce AI access and actions across their entire ecosystem. No other major platform combines front-end AI automation with this depth of identity security.

Cyber Exposure Management Across the Full Attack Surface

Governing AI at enterprise scale requires knowing not just who has access, but what's connected and what's exposed. Upon closing, Armis will extend the AI Control Tower into proactive cybersecurity by connecting real-time asset discovery, threat intelligence, and risk prioritization with automated remediation and response workflows.

Rather than fragmented views in separate tools, Armis and ServiceNow will deliver a unified, end-to-end security exposure and operations stack that can see, decide, and act across the entire technology footprint: IT, OT, IoT, cloud, medical devices, and industrial systems. Exposure insights will automatically flow to the right teams, trigger remediation at scale, and deliver measurable, continuous reduction of enterprise risk.*

**The acquisitions of Armis and Veza are subject to customary regulatory approvals and closing conditions. References to combined capabilities in this document reflect anticipated benefits following the completion of each transaction.*

Mission-Critical Reliability and Global Scale

As AI moves from assistance to autonomous execution, reliability becomes as important as intelligence. ServiceNow's operational maturity provides the proven, enterprise-grade foundation that autonomous workflows demand.

99.99% Uptime Architecture: Enterprise SLA commitments with no more than 52.6 minutes of unplanned downtime per year. Active-active failover, geographically distributed data centers, sub-second degradation detection, and zero-downtime deployments. This is the reliability threshold where enterprises trust systems with autonomous operations.

Multi-Instance Customer Isolation: Dedicated environments with isolated databases per tenant for independent scaling, maintenance, and disaster recovery. Strong isolation ensures data security and performance stability when AI agents execute workflows for hundreds of millions of users globally.

Continuous Availability & Disaster Recovery: Automated failover with Recovery Time Objectives in minutes and Recovery Point Objectives in seconds. Daily backups with fast-restore SLOs and comprehensive continuity protections for long-running autonomous workflows.

Enterprise-Grade Performance & Scale: Horizontally scalable architecture with performance monitoring from database queries to CDN edge caching. Capacity planning systems predict and pre-scale for load spikes, ensuring always-on execution across regions and time zones.

Built-In Compliance & Audit Infrastructure: SOC 2 Type II with immutable audit trails. Multi-regulatory compliance (GDPR, HIPAA, FedRAMP, SOX, CCPA) with configurable data residency and retention policies.

This operational maturity is critical. As AI runs workflows, reliability, security, performance, and scale become as important as intelligence. ServiceNow starts with a proven, enterprise-grade foundation.

4. Why Competitors Fall Short

The AI Control Tower combines Sense, Decide, Act, and Govern into a unified platform. Every competitor delivers pieces. None delivers the whole.

4.1 Standalone LLMs: Intelligence Without Action

Large language models represent a major leap in reasoning and language understanding. By design, however, they are general-purpose intelligence engines rather than enterprise operating systems.

Standalone LLMs lack the foundational capabilities required to deliver autonomous workflows at enterprise scale:

- **No native enterprise context.** They do not understand an organization's workflows, roles, assets, dependencies, or process history without heavy customization. Even leading model providers increasingly acknowledge that the breakthrough isn't the LLM alone. It's the context, memory, and system integration wrapped around it that makes AI truly valuable.
- **No execution layer.** They can suggest actions but cannot securely and reliably orchestrate multi-step workflows across enterprise systems.
- **No native rollback or transactional safeguards.** No built-in mechanism to revert changes, restore prior state, or enforce controlled recovery the way workflow-driven enterprise platforms can.
- **No persistent operational memory.** They do not inherently retain institutional knowledge across cases, decisions, and outcomes.
- **No built-in governance model.** Enterprises require permission-aware actions, audit trails, policy enforcement, and compliance controls.
- **No system of record integration.** LLMs operate outside the systems where work actually happens.

In Sense/Decide/Act/Govern terms: LLMs can partially Decide. They cannot Sense enterprise context, Act within operational systems, or Govern outcomes.

4.2 Vibe Coding: Prototypes Without Process Capital

Vibe coding dramatically lowers the barrier to building software. It accelerates ideation, prototyping, and even early-stage automation. But enterprise value is created by decades of accumulated business process capital, not by software alone. In practice, vibe coding struggles to translate into durable, production-grade outcomes because it cannot easily recreate the deep operational foundations that businesses rely on.

Key gaps:

- **Business Process Capital Cannot Be Generated.** Enterprises run on layered workflows built over years: approvals, SLAs, exception handling, cross-functional coordination. Vibe coding can create new apps but cannot replicate the maturity and reliability of processes refined through real operational experience.
- **Data Model Depth.** Enterprise data models encode years of business logic, relationships, and edge cases. Generated applications typically start with simplified schemas that lack the richness needed to support mission-critical operations at scale.
- **Security & Compliance by Design.** Enterprise systems must enforce role-based access, auditability, data residency, and regulatory controls from day one. Vibe-coded solutions often treat these as afterthoughts, creating risk when moving beyond prototypes.
- **Resilience & Operational Hardening.** Production environments require uptime guarantees, failure handling, observability, and recovery mechanisms. Vibe coding makes it easy to create working software but much harder to create software that is durable under real-world load and failure conditions.
- **Integrations Across Systems.** Enterprises depend on deep, bi-directional integrations across dozens or hundreds of systems. Recreating stable, secure, and governed system-wide integration layers is complex and time-intensive.
- **Cross-Department Coordination.** Most enterprise work spans multiple departments, approval chains, and policy boundaries. Vibe coding typically optimizes for local use cases rather than the interconnected operating model of a large organization.
- **Economics of Scale.** Vibe coding makes the first 20% easy: the prototype, the interface, the early automation. The remaining 80% (hardening, integrating, governing, and maintaining) is where most of the cost and complexity lives. For most enterprises, rebuilding this foundation from scratch is prohibitive.
-

4.3 Data Platforms: Insight Without Execution

Data Platform companies play a critical role in the AI ecosystem by organizing, preparing, and modeling enterprise data. They excel at powering analytics, model development, and data-driven insights. Autonomous workflows at enterprise scale require more than data and models, however.

Key gaps:

- **Strong in data, limited in process.** They help organizations understand their data but do not manage operational workflows.
- **No system of action.** Insights generated must still be executed through other enterprise platforms.
- **Lack of embedded business context.** Data platforms structure information but do not inherently capture the operational relationships between people, processes, and decisions.
- **No execution-native governance layer.** Enterprises need AI that can act within permissioned, policy-aware workflows.
- **Distance from day-to-day operations.** Data platforms power intelligence but do not sit inside the flow of work where incidents, approvals, service requests, and operational tasks are resolved.

In Sense/Decide/Act/Govern terms: Data platforms Sense well. They partially Decide. They cannot Act or Govern.

4.4 Digital Workers & Agent Frameworks: Tasks Without Operations

Anthropic's digital worker concept and frameworks like OpenClaw represent an important evolution from chat assistants to task-level agents. These systems can automate knowledge work such as legal review, research, and content creation. They primarily operate at the task layer, however, rather than the enterprise execution layer.

Key limitations:

- **Limited process awareness.** They can complete tasks but are not deeply embedded in enterprise workflows where multi-step business processes are managed.
- **Shallow enterprise context.** They depend on external integrations to access data but do not natively understand relationships across people, systems, roles, and approvals.
- **No ownership of execution infrastructure.** They rely on other platforms to carry out actions across systems.
- **Fragmented governance and lacking security controls.** Enterprise-grade control requires consistent enforcement of permissions, roles, and compliance within the workflow itself. Systems like OpenClaw lack built-in role-based access control, policy enforcement, auditability, and permission-aware execution, making them difficult to safely deploy in regulated enterprise environments without significant additional security layering.
- **Task-centric rather than operations-centric.** They assist functions (legal, marketing, analytics) but do not manage end-to-end operational flows across departments.

OpenClaw has captured the market's imagination by proving that agentic AI can act, not just respond. The deeper irony is that the same design choices that make it powerful also make it risky. OpenClaw

introduces meaningful security risk because its agents can run commands, access local files, and interact with external systems. In many setups, the primary agent operates with host-level access, creating potential for remote command execution, data leakage, or privilege misuse. Sensitive credentials such as API keys and OAuth tokens are often stored locally in plaintext. Deployments exposed without strong authentication have already led to real-world exploits.

Leading voices have highlighted the contradiction: Cisco's AI security team called it "groundbreaking" but a security "nightmare." VentureBeat noted it proves agentic AI works while proving security models don't. Andrej Karpathy cautioned against running it on personal machines.

OpenClaw proves autonomous execution is real and simultaneously demonstrates why ungoverned execution isn't enterprise-ready.

Anthropic's Cowork explicitly lists risks showing it is not enterprise-ready, advising users to avoid granting access to sensitive files, limit browser extension access, monitor for prompt injection, and restrict internet access to trusted sites.

In Sense/Decide/Act/Govern terms: Digital workers and agent frameworks can Decide and partially Act on individual tasks. They cannot Sense across the enterprise, Act across operational systems at scale, or Govern with enterprise-grade controls.

5. Where the Landscape Is Heading

ServiceNow's position is strong. But the landscape isn't standing still. Model providers are expanding beyond intelligence into workflow tooling and enterprise context. Data platforms are extending into operational execution. Enterprise adoption patterns are still forming, and platform consolidation is not guaranteed. The structural advantages described in this paper are durable. But they compound only through relentless focus on deepening operational embedding, strengthening governance, and delivering faster than the market around us.

6. How We Deliver

ServiceNow is becoming an AI-native enterprise — AI embedded into every product, every feature, every interaction with the platform. Not AI bolted on. AI built in. That means AI-native engineering practices, AI-native pricing and packaging, multimodal user experiences, and a modernized tech stack that delivers innovation faster than ever. That shift is already underway.

1. Expand the Act Layer: Autonomous Workflows as the Applied AI Execution Layer
Be the system where AI-powered work actually happens, where it's executed rather than discussed or monitored. AI in chat is advice. AI in workflows with full enterprise context is action. Expand AI Workers and autonomous workflows into finance, legal, procurement, and supply chain. Every new workflow

domain adds context, strengthens the execution advantage, and makes applied AI more accurate and harder to displace.

2. Deepen the Sense Layer: Compound Accuracy and Trust
Contextual AI Grounding with CMDB, Context Graph, and workflow history is a compounding differentiator that improves with every execution. The Armis acquisition extends this into cyber-asset management, creating institutional memory that makes AI more accurate and trusted. Competitors layer AI on top. ServiceNow embeds it inside operational reality.

3. Make Govern a Defining Strength: AI Control Tower as the Standard
Position AI Control Tower as the standard for trusted enterprise AI. With Gartner predicting 40% of agentic AI ventures will fail by 2027 due to governance challenges, our governance-first architecture becomes a strategic differentiator. Enterprises will choose platforms that apply AI safely, not just powerfully.

4. Accelerate Innovation Velocity: Be AI-native
The shift from biannual to monthly releases is critical. Platform complexity creates slower iteration versus AI-native tools. This must be addressed through architectural modernization and increased R&D investment.

5. Create the best AI Front Door for the Enterprise
ServiceNow is building a single conversational AI interface that collapses 367+ applications into one consumer-grade front door — one place to ask questions, take action, and resolve issues across IT, HR, CRM, Finance, and beyond, grounded in enterprise context, permissions, and workflows.

7. The ServiceNow Advantage: Applied Intelligence in the Flow of Work

Intelligence is rapidly becoming a commodity accessible to anyone through an API. The differentiator isn't who has the smartest model. It's who can apply intelligence at the precise moment it matters, grounded in the right context, embedded in the flow of work, and capable of turning insight into action.

ServiceNow's unique IP is AI embedded inside the operational fabric of the enterprise, grounded in real workflows, assets, relationships, and history. We don't layer intelligence on top of work. We apply it inside the execution layer where work actually happens. Competitors can access frontier models. They cannot easily replicate our operational context, our execution position, or the institutional knowledge that compounds with every workflow.

If we anchor AI inside real workflows with deep operational context and strong governance, ServiceNow becomes the operating system for autonomous workflows at enterprise scale.

The more capable AI agents become, the more they need a platform that supplies identity resolution, entitlements, workflow constraints, integration governance, audit evidence, and change management. Intelligence will keep getting cheaper. Trusted execution will keep getting more valuable.

That's the AI Control Tower for Business Reinvention. That's where ServiceNow lives.