



WHITEPAPER

The Application of Agentic AI for IT Help Desk Services

24/7

September 2026

 **evolver**

Executive Summary

Information technology (IT) help desks have become the operational front line of the modern enterprise, and they are under unprecedented strain. IT service ticket volumes are rising faster than budgets and staffing, driven by hybrid work, accelerating cloud migration, expanding application portfolios, and a steady tempo of cybersecurity events.

At the same time, leadership across both the U.S. Federal government and Fortune 500 organizations expects measurable improvement in user experience, first-contact resolution, and total cost of ownership across the IT Service Management (ITSM) lifecycle.

Traditional ITSM tooling, however capable, cannot close this gap on its own: analysts spend a disproportionate share of their time on triage, categorization, knowledge search, and routine ticket enrichment rather than on judgment-intensive work.

This whitepaper argues that agentic artificial intelligence (AI), deployed natively inside an ITSM platform such as ServiceNow, materially improves IT help desk outcomes when it is governed, orchestrated, and held inside the organization's security boundary.

Agentic AI differs from earlier generative assistants because it does not merely summarize or draft text; it plans and executes multi-step tasks against enterprise systems with human oversight. Industry IT analysts project that this capability will reshape service operations: Gartner forecasts that agentic AI will autonomously resolve eighty percent of common service issues by 2029 while reducing operational costs by roughly thirty percent.¹

Key takeaways for senior leaders:

1. **Govern agentic AI as a controlled capability.** Adopt the National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF) and the NIST Cyber Security Framework (CSF) 2.0 as the governance backbone, with model inventories, lifecycle monitoring, and evidencable documentation.²
2. **Keep data inside the boundary.** For U.S. federal and high-security environments, public AI services raise data-sovereignty and compliance concerns; the orchestration layer must run on-premises, hybrid, or air-gapped so sensitive data never leaves the ServiceNow environment.
3. **Avoid single-vendor AI lock-in.** Multi-model routing directs each task to the best-fit model for accuracy, latency, or cost, rather than paying a fixed per-module premium for one provider's models.
4. **Measure value before scaling.** Gartner cautions that more than forty percent of agentic AI projects may be canceled by the end of 2027 owing to unclear value or inadequate risk controls; bounded 60-to-120-day pilots with baselined metrics are the disciplined path forward.³

1. Introduction: The IT Help Desk Under Pressure

The IT help desk is where enterprise technology meets the people who depend on it. When it works well, it is nearly invisible; when it does not, productivity, mission readiness, and user trust all degrade. Three structural pressures now converge on help desk operations. First, demand for IT help desk services are growing: Larger hybrid workforces, continued cloud migration, and expanding application portfolios generate more service tickets across more channels. Second, supply is constrained as skilled IT analysts are scarce and expensive, while budgets rarely grow in step with demand. Third, expectations are rising. Leaders want faster resolution, higher first-contact resolution, and lower cost per ticket, all at once.

Conventional ITSM platforms organize this work through structured workflows, service-level agreements (SLAs), and a configuration management database (CMDB), but the cognitive labor still falls on people. IT analysts read each ticket, infer intent, classify and route it, search knowledge bases, gather context from multiple systems, and compose responses. Much of this is deterministic and repetitive, precisely the kind of work that drives queue backlogs, inconsistent quality, and analyst fatigue. The opportunity is to shift analysts toward decisions that require human judgment while assigning routine tasks to governed AI.

It is useful to think of the difference between earlier AI assistants and agentic AI as the difference between a calculator and an accountant. A calculator returns an answer to a defined problem; an accountant understands the goal, gathers the necessary information, and works toward a finished result. Agentic AI operates more like the accountant, carrying out multiple steps toward resolution under supervision. Gartner projects that by 2029 roughly seventy percent of enterprises will deploy agentic AI to operate IT infrastructure, up from fewer than five percent in 2025, shifting operators from manual responders to supervisors of AI-driven operations.⁴

2. What Agentic AI Changes in ITSM

From Generative Assistance to Autonomous Action

Generative AI assistants improved drafting and summarization, but they remained passive: a human still had to act on every suggestion. Agentic AI introduces a new operating model in which AI systems can act autonomously to complete tasks, navigating systems and taking steps on a user's behalf rather than only producing text.⁵ In ITSM, an agent can classify and prioritize incidents, gather relevant context, execute approved remediation steps, and document the resolution, with human approval for consequential actions.

The IT Help Desk Tiers, Reimagined

Agentic AI maps cleanly onto the established support tier model. The table below summarizes how AI agents augment each tier while preserving human accountability for high-impact decisions.

Tier	Traditional Model	With Agentic AI (ASSIST)
Tier 0	Static Frequently Asked Questions (FAQ) and IT portal self-service; low deflection.	Conversational virtual agent resolves common requests and password, access, and how-to issues autonomously.
Tier 1	Manual triage, categorization, routing, and basic troubleshooting.	AI agents auto-triage, classify, prioritize, and route; draft first responses for analyst approval.
Tier 2	IT Analyst gathers context across systems; searches knowledge manually.	Agents enrich service tickets, correlate incidents, and surface root-cause hypotheses and relevant knowledge instantly.
Tier 3/4	Engineering deep dive; tribal knowledge; slow documentation.	Agents draft incident timelines and post-incident knowledge; engineers focus on novel problem-solving.

Table 1. Help desk support tiers augmented by governed agentic AI.

Why IT & AI Governance Is the Precondition, Not the Afterthought

Autonomy without governance is a liability. Embedding AI into operational tools expands the attack surface through risks such as data poisoning, prompt injection, and model manipulation. NIST's AI Risk Management Framework places the GOVERN function at the center of trustworthy AI, stressing defined roles, training, monitoring, and documentation throughout the lifecycle. ⁶ NIST CSF 2.0 elevates GOVERN as the organizing function for cybersecurity outcomes. ⁷ The practical implication for help desk modernization is that AI oversight should be integrated into existing cyber governance rather than bolted on as a separate committee that never reaches the workflow.

This caution is grounded in evidence. Gartner found in 2025 that only fifteen percent of IT application leaders were considering, piloting, or deploying fully autonomous AI agents, in addition to the rising probability of canceled AI projects. ⁸ The lesson is not to avoid agentic AI but to deploy it where value is clear, controls are strong, and outcomes are measured against a baseline.

3. The Three Levers: Quality, Speed, and Cost

Agentic AI improves IT help desk performance along three reinforcing dimensions: quality, speed, and cost. Each lever is meaningful on its own; together they change the economics of IT support.

Enhance IT Service Quality

Quality in IT service is consistency under pressure; the same issue resolved in the same correct way regardless of which IT analyst, shift, or channel handles it. Agentic AI can improve consistency by summarizing context, correlating incidents, and applying repeatable reasoning to tickets. Analysts can review, accept, or override its recommendations, preserving human accountability. Vendor-reported outcomes in the ITSM market illustrate the magnitude: enterprises adopting agentic ITSM platforms have reported auto-resolution rates above seventy-five percent and double-digit gains in employee satisfaction. ⁹

Accelerate the Speed to Response

Speed is where automation compounds. Automating triage, classification, routing, and enrichment can reduce these tasks from minutes to seconds, lowering mean time to resolution (MTTR) and improving first-contact resolution. Selective automation converts repetitive work into standardized playbooks for enrichment, containment, and evidence packaging, an approach consistent with NIST's incident-handling guidance, which emphasizes preparation, repeatable processes, and learning loops. ¹⁰ The objective is not full automation but the reallocation of human attention to the decisions that require judgment.

Reduce Operational Cost

Agentic AI can reduce costs in two ways. First, deflection and automation reduce the analyst hours consumed per ticket, lowering labor cost, the dominant expense in any help desk. Second, multi-model orchestration reduces AI spend itself by routing each task to the most cost-effective model that meets the accuracy and latency requirement, rather than paying a flat premium for a single provider's models on every transaction. This matters because the prevailing commercial model for platform-native AI is expensive. ServiceNow's Now Assist is sold as a separately licensed add-on, per product line, and generally requires customers to be on a higher-priced Professional or Enterprise tier before the AI capability can be purchased at all. ¹¹ Industry analyses estimate the add-on at roughly fifty to one hundred dollars or more per fulfiller per month, which for a 500-fulfiller organization can add several hundred thousand dollars annually on top of existing platform cost. ¹²

4. Case Studies

ASSIST, Evolver's agentic AI capability for ServiceNow powered by L2H.ai, has been deployed and placed into operational use at two large-scale federal IT organizations, one within the U.S. Department of War (DoW) and one within the U.S. Intelligence Community. These high-security environments provide demanding proving grounds for agentic AI in IT help desk operations.

The Challenge

The two federal IT organizations needed to increase service capacity while operating within strict security and compliance requirements. Their solution needed to integrate with existing ServiceNow environments, keep sensitive data inside the security boundary, and avoid dependence on a single AI provider.

The Solution: ASSIST

The Evolver ASSIST solution was rapidly installed, integrated, tested, and deployed via L2H.ai's AI orchestration platform for ServiceNow. The solution embedded vendor-agnostic, multi-model AI agents directly within the agency's ServiceNow instance and automated triage, summarization, root-cause analysis, and knowledge-article generation across ITSM and IT Operations Management (ITOM). Critically, the solution was workflow-native and non-disruptive: no data left the ServiceNow boundary, and existing automation, integrations, and CMDB investments were preserved while AI reasoning was added across every workflow.




Figure 1. Evolver ASSIST architecture: L2H.ai Airwall multi-model orchestration embedded inside the agency ServiceNow security boundary.

How Airwall Outperformed the Native AI Add-On

The distinction that mattered most at both the selected DoW & IC IT organizations was architectural. ServiceNow's native AI capability, Now Assist, is licensed per product line as an optional add-on and is typically gated behind a higher subscription tier, so enabling AI across modules multiplies cost module by module.¹³ The L2H.ai Airwall layer took a different approach. Airwall avoided dependence on a single provider by dynamically routing each task to the best-fit approved model, including OpenAI, Anthropic, Gemini, Llama, Mistral, or an agency model. Routing could be optimized for accuracy, latency, or cost while remaining inside the security boundary under unified guardrails and audit logging.

This produced three advantages over the native add-on. First, higher quality, because the orchestration layer could select the strongest model for a given task type rather than accepting one model for all tasks. Second, lower cost, because inexpensive models handled high-volume routine work while premium models were reserved for complex reasoning, avoiding a blanket per-fulfiller AI premium. Third, stronger sovereignty and governance, because every inference, routing decision, and policy enforcement was logged inside the agency environment for full auditability, aligned to FedRAMP, FISMA, and NIST SP 800-53.

Why multi-model routing wins, by analogy

Multi-model routing works like triage: routine tasks go to lower-cost resources while complex tasks are routed to more capable models. This improves resource efficiency while reducing dependence on a single provider.

Measured Outcomes

Against a baselined legacy help desk, the ASSIST pilot delivered improvements across every core service metric. The figures below are representative of the gains targeted and typically observed in the ASSIST engagement model and are consistent with the fifty-to-seventy-percent faster resolution range Evolver specifies for the offering.


Metric	Baseline	With ASSIST	Change
Mean Time to Resolve (MTTR)	100 (index)	38 (index)	- 62%
First-Contact Resolution	100 (index)	152 (index)	+ 52%
Tier-0/1 Deflection Rate	100 (index)	245 (index)	+ 145%
Cost per Ticket	100 (index)	52 (index)	- 48%

Table 2. Representative ASSIST pilot outcomes versus baseline (indexed to 100).

The Cost Picture Over Three Years

The illustrative three-year model compares costs for roughly 600 fulfillers under three approaches: the status quo, ServiceNow with Now Assist, and Evolver ASSIST with L2H.ai routing. The per-module path carries the highest total cost because the AI premium applies broadly and grows with every module enabled, while ASSIST reduces both labor hours and AI spend.

Figure 2. Illustrative Example: three-year cumulative cost of ownership. Figures are illustrative and directional, intended to show relative trajectory rather than a specific contract quotation.



5. Implementation: A Disciplined, Phased Approach

Evolver delivers ASSIST through a bounded, measurable, four-phase model that proves value inside the boundary before scaling. A 60-to-120-day pilot covering a single ServiceNow module or bureau is the standard entry point.

Phase	Timeline	Key Activities
1. Discovery & Baseline	Days 0–30	Review the agency ServiceNow instance, ticket volumes, and integrations; baseline MTTR, first-contact resolution, deflection, and cost per ticket; prioritize use cases; define data-sovereignty, hosting, and model-governance policies.
2. Airwall Deployment	Days 30–60	Deploy the Airwall orchestration layer inside the boundary in the selected hosting model; configure multi-model routing; activate triage, summarization, RCA, and knowledge agents; enforce guardrails, audit logging, and role-based access.
3. Workflow Integration & Pilot	Days 60–120	Embed AI reasoning into Flow Designer and ITSM/ITOM processes; tune Tier 0–2 workflows with agency knowledge; run a side-by-side pilot against baseline; deliver analyst training and change management.
4. Sustainment & Scale	Ongoing	Operate 24/7 Tier 0–4 help desk with continuous AI-assisted resolution; conduct quarterly governance reviews with executive dashboards; optimize routing as new models are approved; expand across modules and bureaus.

Table 3. The ASSIST four-phase delivery model.

Governance and Compliance Alignment

Throughout all phases, ASSIST maintains full auditability, policy enforcement, and model governance aligned to federal compliance frameworks, including FedRAMP, FISMA, and NIST SP 800–53, within a Zero Trust posture. This satisfies the dual-governance requirement, governing cybersecurity risk and AI risk together, that NIST CSF 2.0 and the NIST AI RMF jointly establish.¹⁴

Leadership Recommendations

Agentic AI can improve resolution time, first-contact resolution, deflection, cost per ticket, and service consistency. Achieving those gains requires strong governance, measurement, and security controls. The evidence from the DoW and IC example engagements are that the decisive factors are architectural: keep data inside the boundary, orchestrate across multiple models rather than locking into one provider's per-module premium, and govern AI as rigorously as any other operational control.

For senior leaders in both federal agencies and Fortune 500 enterprises, the path forward is straightforward and low-risk:

- **Start with a bounded pilot.** Baseline current performance, then run a 60-to-120-day pilot on one module or bureau with explicit success criteria.
- **Insist on boundary control and multi-model orchestration.** Require that sensitive data never leaves the environment and that the architecture avoids single-vendor AI lock-in.
- **Govern AI as part of cyber governance.** Align to NIST AI RMF and CSF 2.0 with model inventories, monitoring, and auditable evidence.¹⁵
- **Measure relentlessly.** Track MTTR, first-contact resolution, deflection, cost per ticket, and model spend on executive dashboards, and expand only where value is proven.

The Evolver ASSIST technology solution puts these recommendations into practice by combining Evolver's Enterprise IT Help Desk Services with L2H.ai Airwall orchestration on ServiceNow. The solution is designed to improve service quality, accelerate response and resolution,

and reduce operating costs while keeping sensitive data inside the required security boundary. For organizations modernizing IT support under pressure, agentic AI is no longer experimental; with the right architecture and governance, it is a proven, measurable advantage.

About the Authors



Gregory A. Garrett is the Chief Operating Officer and Chief Innovation Officer for Evolver. He leads all of Evolver's technology thought leadership and technology business units, which provide advanced IT, AI, cybersecurity, electronic security, and eDiscovery services to federal agencies and Fortune 500 companies.

With over 25 years of executive experience and a background as a CIO, CTO, CISO, and COO, he has managed more than \$40 billion in global tech contracts and served in key roles across both large and mid-sized firms. A retired U.S. Air Force Colonel and accomplished author, Garrett has also contributed extensively to industry thought leadership through 24 books, more than 150 articles, expert testimony, and over 250 speaking engagements.



David Hayes serves as CIO of Commercial IT at Evolver, where he leads the design, deployment, and operation of large-scale information management for commercial customers. Bringing decades' worth of deep expertise in data analytics, performance, scalability, and continuity of operations across petabyte-scale environments, David has driven the adoption of advanced technologies that streamline information workflows and help clients reduce expenses. His leadership spans multiple industries, including financial services, aviation, telecommunications, and automotive, reinforcing Evolver's reputation for secure, innovative data solutions.



Dr. Brian McElyea is an accomplished cybersecurity executive recognized for leading large-scale digital transformation initiatives and building high-performing Cyber Centers of Excellence that drive mission-aligned security outcomes. With over two decades of leadership across healthcare, defense, and federal sectors, he has guided global organizations through the complexities of risk management, compliance, and modernization with measurable results.

Works Cited

Aisera. "[Aisera Named a Visionary in 2025 Gartner Magic Quadrant for AI Applications in IT Service Management.](#)" Aisera, 5 Sept. 2025. Accessed 2 June 2026.⁹

eesel AI. "[A Complete Guide to ServiceNow AI Pricing in 2026.](#)" eesel AI, 2026. Accessed 2 June 2026.¹¹

eesel AI. "[ServiceNow License Types and Costs 2026: ITSM, CSM, ITOM Pricing Explained.](#)" eesel AI, 2026. Accessed 2 June 2026.¹³

Gartner. "[Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029.](#)" Gartner, 5 Mar. 2025. Accessed 2 June 2026.^{1, 5}

Gartner. "[Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027.](#)" Gartner, 25 June 2025. Accessed 2 June 2026.³

Gartner. "[Gartner Survey Finds Just 15% of IT Application Leaders Are Considering, Piloting, or Deploying Fully Autonomous AI Agents.](#)" Gartner, 30 Sept. 2025. Accessed 2 June 2026.⁸

National Institute of Standards and Technology. [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#). U.S. Department of Commerce, Jan. 2023. Accessed 2 June 2026.^{2, 6, 14, 15}

National Institute of Standards and Technology. [Cybersecurity Framework \(CSF\) 2.0](#). NIST CSWP 29, 26 Feb. 2024. Accessed 2 June 2026.^{2, 7, 14, 15}

National Institute of Standards and Technology. [Incident Response Recommendations and Considerations for Cybersecurity Risk Management](#). NIST Special Publication 800-61r3, Apr. 2025. Accessed 2 June 2026.¹⁰

PagerDuty. "[Gartner Report: Predicts 2026 — AI Agents Will Transform IT Infrastructure and Operations.](#)" PagerDuty, 2026. Accessed 2 June 2026.⁴

Redress Compliance. "[ServiceNow Now Assist AI Pricing: Enterprise Cost Guide 2026.](#)" Redress Compliance, 3 Apr. 2026. Accessed 2 June 2026.¹²