

Hype Cycle for ITSM, 2026

24 June 2026 - ID G00846493 - 95 min read

By: Chris Laske, Siddharth Shetty, Chris Matchett

Initiatives: Delivery of Functional Responsibilities

The drive to automate service management practices or even move to a zero-touch operation is taking center stage. In order to avoid missteps, use this information to cut through the hype and determine which innovations to build into a service management strategy.

Analysis

What You Need to Know

Heads of infrastructure and operations (I&O) are experiencing a steady barrage of vendor hype associated with rapid advancements in the capabilities of IT service management (ITSM) technologies. At the same time, they face increasing pressure from CIOs to integrate AI into their work practices to streamline service delivery and improve levels of quality. Heads of I&O must strike a balance between innovation and technical resilience.

The 2026 Hype Cycle for ITSM provides critical insights into the practices and technologies on the forefront of IT service management. New to this year's Hype Cycle and top of mind for many clients is AI agents. Generative AI continues to evolve and is rapidly moving toward the Slope of Enlightenment. Service response has been rebranded as service operations and moved from the Innovation Trigger to The Peak of Inflated Expectations, given increased interest along with technology alignment.

To provide more in-depth analysis of AI capabilities in ITSM, a parallel Hype Cycle with this specific focus has again been published in 2026: [Hype Cycle for AI in ITSM, 2026](#)

The Hype Cycle

The 2026 Hype Cycle for ITSM has significant updates:

- Digital workplace is a critical adjacent practice area for ITSM, since stability of endpoints can lead to significantly fewer service desk contacts. Autonomous endpoint management jumped from the Innovation Trigger to the Peak of Inflated Expectations, while digital end-user experience (DEX) tools and digital employee experience progressed through the Trough of Disillusionment.
- Technology change automation in support of capabilities such as standard changes has matured and is becoming more broadly adopted, advancing from the Trough of Disillusionment to the Slope of Enlightenment.
- Along with the newly added AI agents, numerous innovations associated with AI that highlight the advancements in artificial intelligence are rapidly advancing on the Hype Cycle in recognition of the hype and prominence of these technologies in the marketplace.
- Reduced hype on topics in the service management space has precipitated the departure of innovations, including virtual support agents, collaborative support hub, and observability, while innovations such AI-focused problem management were eliminated due to their inclusion in a more integrated set of AI-powered IT agent advisory features.


Automation continues to generate hype in I&O and service management. Infrastructure automation, policy as code, automated incident response, and technology change automation (TCA) all remain, with infrastructure automation and automated incident response holding relatively steady in the Slope of Enlightenment.

Technology reliability, bolstered by AI, is top of mind for many Heads of I&O as they look to proactively reduce the number and impact of incidents. Agent-Native I&O has been retitled as Agent-Centric I&O as a pattern for organizing infrastructure and IT operations designed to exploit agentic AI in making autonomous and semi-autonomous decisions. Event intelligence remains near the Slope of Enlightenment as more data is generated and demands for reliability are increasing. Autonomous endpoint management is accelerating endpoint patching and streamlining configuration management, while protecting the digital employee experience.

Multiagent systems have reached the Peak of Inflated Expectations as agentic AI, potentially across disparate platforms, is attracting significant market attention. Configuration management databases (CMDBs) are still a focal point for ITSM, driving the importance of the technology dependency hub. More vendors are working to align DEX, event intelligence, and ITSM capabilities to optimize technology support, encouraging clients to consider how service operations teams can be established. Security is always a critical concern as teams collaborate to minimize risk, with threat exposure management remaining on the Peak of Inflated Expectations.

Figure 1: Hype Cycle for ITSM, 2026

Hype Cycle for ITSM, 2026



Gartner

The Priority Matrix

Immediate considerations:

- Generative AI has become pervasive in the ITSM vendor landscape. Heads of I&O must determine how to leverage this technology to optimize service desk performance.
- Given the increasing complexity of technology environments and the ongoing need to reduce business impacting incidents, event intelligence should be a key point of focus.

- With broad adoption of agile methodologies and rapid deployments, organizations need to take advantage of technology change automation to avoid being a bottleneck to the delivery of value.

Short-term considerations:

- IT knowledge generation has become increasingly vital as organizations work to deflect more tickets from their service desks. Heads of I&O must work with their teams and embrace AI to increase the quality and quantity of knowledge content to improve results.
- Heads of I&O are under increasing pressure to streamline service management practices. AI-powered agent advice is a burgeoning capability that can provide intelligent recommendations to agents, improving service levels and reducing mean time to repair.
- More and more ITSM platforms are being delivered with AI agents as well as AI agent builders. Heads of I&O must evaluate the capabilities of AI agents to determine how they can potentially add value to service management practices.

Long-term considerations:

- Technology dependency hubs are a growing area of need and interest as organizations see the increasing value in reconciling conflicting dependency data across discovery, observability, and cloud platforms. The growing number of tools in these categories, along with the need to expedite incident response and limit risk associated with business-impacting changes, makes this an increasingly important governance capability.
- AI applications in ITSM have become a driving force in service management. Yet, many AI capabilities are still relatively immature and emerging. Heads of I&O must closely manage AI adoption. Companion research, including the [Hype Cycle for AI in ITSM, 2026](#) as well as the [Magic Quadrant for AI Applications in IT Service Management](#) and [Critical Capabilities for AI Applications in IT Service Management](#) can provide valuable guidance in this regard.

Table 1: Priority Matrix for ITSM, 2026

(Enlarged table in Appendix)

Benefit ↓	Years to Mainstream Adoption			
	Less Than 2 Years ↓	2 to 5 Years ↓	5 to 10 Years ↓	More Than 10 Years ↓
Transformational	Digital Employee Experience Generative AI	Infrastructure Platform Engineering Multiagent Systems Service Operations Site Reliability Engineering	Agent-Centric I&O Threat Exposure Management	
High	Automated Incident Response Autonomous Endpoint Management Technology Change Automation XLA	AI Agents AI-Powered IT Agent Advisory DEX Tools Infrastructure Automation Policy as Code	AI Applications in IT Service Management Enterprise Service Management Technology Dependency Hub	Natural Language Case Extraction
Moderate	Event Intelligence Solutions Software Asset Management Tools	IT Knowledge Generation	Swarming Support	
Low		Operations Assistant		

Source: Gartner (June 2025)

On the Rise

Natural Language Case Extraction

Analysis By: Chris Matchett

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Definition:

Natural language case extraction uses natural language technologies and LLMs to automatically recognize incidents and requests from organic conversations and create associated tickets in a case management system (e.g., an ITSM platform or a CRM customer engagement center).

Why This Is Important

Natural language case extraction (NLCE) is an essential capability for a functional collaborative support hub. It scans collaboration channels and support forums for conversations, data, knowledge and metrics for support activities that have not gone through a case management system and pulls them into the platform. NLCE agentically automates the slow, labor-intensive work previously done by IT service desk agents, who would manually search for incidents to add into the ITSM platform.

Business Impact

- Provides visibility of support effort and impact outside of formal contact centers.
- Provides advance warning of service issues that employees are discussing.
- Automatically creates incident tickets without requiring a service desk contact.
- Improves knowledge generation and maintenance by finding undocumented solutions used in peer support.
- Enables reporting of time spent by IT and business experts on the collaborative support hub and measures the impact of peer IT support.

Drivers

- NLCE is technically feasible using generative AI large language models (LLMs), which are well-suited to extracting context and data from unstructured content such as conversations.
- Vendors in AI applications for ITSM have plans for this capability within their roadmaps.
- NLCE appeals to employees and customers who prefer peer support and discussion channels.
- Organizations that implement a collaborative support hub for IT seek NLCE solutions to connect the hub to the case management system, such as their ITSM platforms.

Obstacles

- Software vendors don't have functional implementations of NLCE, despite the technical feasibility, resulting in its low availability.
- Organizations must build their own AI application with integrations to their case management system(s) to be able to use this functionality today.
- Employees may resist or mistrust AI features that monitor conversations.

User Recommendations

- Add a conversational agent to communal channels as a participant to offer assistance alongside formal support channels and human moderators.
- Enable virtual support agents to mark conversations as complete once successfully resolved in order to manage the clutter of common and repeat issues that have known solutions. This allows human expert responders and community moderators to focus on the other conversations that a conversational AI agent is unable to deal with.
- Use NLCE on conversations that are instantiated by swarming, as well as on more permanent collaborative support hub channels.
- Ask your case management platform vendor about these capabilities and how to activate them. The roadmaps are prioritized on customer demand.

Gartner Recommended Reading

[Innovation Insight: Collaborative Support Hubs](#)

[How to Operate and Manage a Collaborative Support Hub](#)

[Critical Capabilities for AI Applications in IT Service Management](#)

Technology Dependency Hub

Analysis By: Ankita Hundal, Roger Williams

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Definition:

A technology dependency hub (TDH) reconciles conflicting dependency data across discovery, observability and cloud platforms, establishing a consistent basis for operational decisions. It does not replace existing tools — it governs how their data is trusted and used during change, incident and automation activities.

Why This Is Important

The issue is no longer lack of visibility; it is conflicting interpretations of the same environment. Teams act on different versions of dependency data, leading to contradictory changes, delayed incident response and unsafe automation. TDH establishes a consistent basis for decisions.

Business Impact

TDH improves change success rates, reduces incident triage time and prevents automation from amplifying bad assumptions by aligning dependency views across teams. Without governance, it introduces overhead — if data ownership and precedence are unclear, it becomes another inconsistent layer.

Drivers

- Enterprises operate multiple tools that represent dependencies differently, forcing teams to make decisions based on conflicting data rather than a shared understanding.
- Dynamic cloud and container environments outpace static models, making dependency data inconsistent across systems rather than simply outdated.
- DevOps and platform teams require dependency-aware systems for release and self-service workflows, but existing models cannot support decision making at their speed.
- Automation and AIOps initiatives fail to scale when dependency data cannot be trusted, exposing that the issue lies in data alignment, not intelligence.
- Regulatory requirements such as DORA and NIS2 demand defensible and current dependency understanding, increasing pressure to reconcile data across systems.
- Cloud migration and modernization efforts stall when dependency conflicts create uncertainty around impact and risk.
- Repeated change failures and prolonged incident resolution have exposed dependency misalignment as an operational risk, not a tooling gap.
- Legacy rationalization efforts fail when organizations cannot determine which systems are actually connected and in use.

Obstacles

- Dependency data lacks clear ownership, causing it to degrade quickly across systems.
- Organizations carry fatigue from failed CMDB and SDM efforts, creating skepticism toward similar initiatives.
- Reconciling data across tools requires sustained effort, slowing early value realization.
- ITSM, DevOps and platform teams operate with different priorities, making alignment difficult.
- Vendors describe capabilities inconsistently, making it hard for buyers to define and evaluate solutions.
- TDH challenges existing notions of “source of truth,” creating resistance as teams are forced to align on data precedence.
- If treated as another tool or integration layer, TDH adds complexity instead of resolving it.
- Poor source data is exposed, not corrected, requiring governance that many organizations are not prepared to enforce.

User Recommendations

- Position TDH as a decision alignment layer, not a replacement for CMDB, observability or discovery.
- Define ownership at the decision level and establish clear rules for which data sources take precedence.
- Start with a single high-impact use case, such as change impact analysis, and validate outputs in live scenarios.
- Integrate TDH outputs into existing workflows rather than creating parallel processes.
- Ensure the model supports DevOps and platform workflows, not just ITSM processes.
- Avoid expanding scope before trust is established in the output.
- Do not treat TDH as a data aggregation or integration project; its value comes from reconciliation and governance.

Sample Vendors

BMC Software; Dynatrace; Faddom; Freshworks (Device42); ServiceNow

Gartner Recommended Reading

[Drive CMDB Transformation Through Continuous Discovery and Service Mapping](#)

[Solution Criteria for AIOps Platforms](#)

[Innovation Insight: Data Observability Enables Proactive Data Quality](#)

[Move Beyond the AIOps Hype: The State of AI in IT Operations](#)

Agent-Centric I&O

Analysis By: Roger Williams

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Agent-centric infrastructure and operations (agent-centric I&O, previously referred to as agent-native I&O) is a pattern for organizing infrastructure and IT operations designed to exploit agentic AI in making autonomous and semiautonomous decisions to act toward organizational goals. Humans and machines work together to understand the current and future needs of key stakeholders and adapt systems and practices to meet those needs.

Why This Is Important

Agentic AI is poised to be the biggest disruptor of I&O since the introduction of the cloud. Agent-centric I&O provides the technical and organizational environments required for AI agents and multiagent systems to fulfill their potential. This includes seamless delivery of infrastructure platforms and services that minimize waste (such as overprovisioning) when responding to demand spikes, and observability of current versus expected states of I&O offerings to secure and improve system performance.

Business Impact

Agent-centric I&O enables AI agents to create the following competitive advantages:

- AI agents can self-serve infrastructure for experimental offerings within human-defined limits.
- People and AI agents can use simulated environments to explore disruptive possibilities.
- AI agents can administer systems and practices to create the resilience and adaptiveness required to meet business needs now and in the future.

Drivers

- AI agent capabilities have already started to drive significant differences in work adjacent to I&O such as software engineering, making them a topic of great interest.
- Heads of I&O are interested in AI agents to address chronic challenges with complexity, talent shortages and expensive cost structures, yet are unclear on what they need to do to best obtain those results.
- Agent-centric I&O provides a compelling vision for how I&O can play a critical role in business success, which addresses perennial I&O concerns about being viewed as a strategic value contributor.
- Agentic AI poses a dilemma for I&O leaders regarding how much of their effort should be devoted to preparing for agentic AI versus addressing current needs. Agent-centric I&O goes beyond just the introduction of AI agents to address this dilemma by providing a comprehensive vision that builds future readiness for agentic AI adoption while addressing current challenges.
- Agentic AI needs trustworthy data about the entire infrastructure estate and the IT operating model to make accurate, timely decisions with minimal costs and risks. The intelligent simulation element of agent-centric I&O addresses this by giving AI agents needed feedback on actions in complex environments to support them at the speed required for effective experimentation.
- I&O staff are at risk of not developing vital infrastructure expertise as AI agents take on easier tasks, which agent-centric I&O accommodates through its use of intelligent simulation.

Obstacles

- **Inadequate governance of I&O:** AI agents need explicit guidance about goals, boundaries and rules that are often not documented, and organizational staff need ways to measure and assess the performance of AI agents to improve results.
- **Staff readiness for AI agents:** Agent-centric I&O requires people to have the motivation, expertise and time to understand how best to work with AI agents.
- **Lack of agentic AI standardization:** Differences in vendor offerings will require trade-offs in how AI agents are matched to work that needs to be done.
- **Lack of I&O effectiveness in managing infrastructure platforms:** The ability for I&O to effectively manage infrastructure platforms is mandatory for obtaining full value from the use of AI agents. Gaps frequently exist in the infrastructure estate, operating model and documented knowledge about the organization that limits I&O effectiveness in managing infrastructure platforms and, in turn, their ability to adopt agent-centric I&O.

User Recommendations

- Create a vision for how agentic AI would change I&O by identifying where I&O is constrained due to limited resources for decision making and taking action, and imagining what would be different if these constraints were broken by agentic AI.
- Provide clarity about the mission and the constraints that AI agents must adhere to in order to fulfill it.
- Use Gartner's [IT Score for Infrastructure and Operations](#) to understand and close the most significant gaps.
- Require new infrastructure investments to be based on intelligent infrastructure and platforms that AI agents can then use to perform tasks and augment their capabilities.
- Address quality gaps in technology asset data required for effective agentic AI use by strengthening ITAM, service management, observability and digital employee experience practices.
- Sponsor staff development vital skills for working with AI agents and instill confidence that they will still be treated as valuable members of the organization as AI agent use increases.

Gartner Recommended Reading

[Blueprint for Platform-Centric I&O](#)

[The Future of I&O 2030: The Impact of AI](#)

[The Future of I&O Automation Is Agentic So Begin Piloting Now](#)

Policy as Code

Analysis By: Paul Delory

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

Policy as code (PaC) languages express governance and compliance rules as code, so they can be enforced programmatically by automation tools. PaC languages are often domain-specific and declarative. With PaC, policies are treated as software, making them subject to version control, code review and functional testing. The most mature PaC tools can render any business logic in code. You can use PaC today to enforce infrastructure compliance, authorization, Kubernetes admission control and more.

Why This Is Important

Platform engineers use PaC to build optimization, governance and compliance controls into automation pipelines. Infrastructure and security teams have used it for years to build guardrails around infrastructure and data while preserving a separation of duties that mirrors a typical IT organization chart. With the rise of generative AI (GenAI), PaC is poised to become a way to control AI agent behavior and enforce standards programmatically, which current GenAI tools often struggle to do.

Business Impact

- **Security, compliance and automation:** PaC, combined with automation, enforces policies with implicit compliance guarantees.
- **Alignment of security and operations teams:** PaC allows security and operations teams to interface directly with automation pipelines.
- **Visibility and auditability:** PaC documents policies. PaC tool logs can be audited to prove policies are being enforced.
- **Time and effort spent:** PaC means less toil for operators because it forestalls configuration drift and out-of-spec elements.

Drivers

- **Agentic AI:** AI agents can make writing and enforcing policy as code easier. PaC tools often require the user to write policies in a dedicated logic language. But an AI agent trained on existing natural-language policy documentation could make effective enforcement decisions with much less effort. The relationship between AI agents and PaC is symbiotic. PaC can provide both effective control and meaningful testing and auditing of agents' outputs. It can help manage AI agents even as AI agents write policy.
- **Government focus:** New government initiatives, most prominently the EU's Interoperable Europe Rules-as-Code, require machine-readable policy definitions. Worldwide, new regulations continue to increase both the difficulty of compliance and the pressure on compliance teams. PaC allows compliance teams and auditors to document their policies in detail and verify that they are being enforced.
- **Continued growth of DevOps and DevSecOps:** As more companies are embracing DevOps and DevSecOps, they are also encountering the hard governance problems of automation. Many teams that implement infrastructure as code quickly find that they need better policy enforcement, and PaC can help.
- **Cloud optimization and cost control:** Besides their benefits for security and compliance, PaC tools can also be used to enforce the build standards for infrastructure, including budgets. In the public cloud, where oversized or unnecessary infrastructure incurs direct out-of-pocket costs, programmatically enforced policies can help to control spending.

Obstacles

- **Scarcity of community-generated content:** PaC tools will not gain traction until they have extensive libraries of community-generated content from which users can download the policies they need rather than having to write their own. Over time, as the user base expands, PaC tools will reach a critical mass of downloadable content that supports real-world uses.
- **Skill set:** Many technical professionals lack the skills to operate PaC tools effectively. As the learning curve might be steep, you may need to accept some flubbed policy enforcement due to lack of experience.
- **Integration challenges:** Integrating with existing tools is complex and often requires additional configuration.
- **Organizational inertia:** In some organizations, collaboration between infrastructure and operations teams and security or compliance teams is actually unwanted. This dynamic may slow the rate, scope and scale of PaC initiatives.
- **Costs:** Even if PaC tools themselves are free, you need support, training and/or consulting.

User Recommendations

- **Start small:** Choose a pilot use case where PaC will likely provide real business benefits, then expand to others once PaC has proven its value.
- **Upskill staff:** PaC languages are not always intuitive. Technical staff will need practice and/or training. Adopt the four-eyes principle to prevent flawed policies from impacting operations.
- **Promote reusability:** Focus your PaC efforts on use cases that have ready-made implementation templates — ideally, downloadable content. For example, almost every PaC tool on the market has a canned implementation of the customer information systems benchmarks.
- **Integrate PaC into automation pipelines:** Use PaC to build guardrails for automation tools, so that they cannot take actions that are out of compliance.
- **Measure before and after:** Use observability tools and value stream mapping to define your starting state, then compare it to the end state. Collect real data to quantify the value of PaC.

Sample Vendors

HashiCorp; Palo Alto Networks; Progress Software; Pulumi; Styra

Gartner Recommended Reading

[Use 'Policy as Code' to Secure Application Deployments and Enforce Compliance](#)

[Magic Quadrant for DevSecOps Platforms](#)

[How to Protect Your Clouds With CSPM, CWPP, CNAPP and CASB](#)

[Innovation Insight for Continuous Compliance Automation](#)

[Innovation Insight for Cloud-Native Application Protection Platforms](#)

Operations Assistant

Analysis By: Chris Matchett, Matt Crossley

Benefit Rating: Low

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

Operations assistants are conversational interfaces that provide human IT agents with data-driven insights and also execute ITSM actions to help them carry out their role. They provide access to AI and GenAI capabilities, such as case summarization and automatic communications.

Why This Is Important

While traditional virtual support agents have focused primarily on helping business consumers, there is a growing demand for solutions tailored specifically to the needs of human IT agents. This has driven interest in operations assistant solutions, which offer user experiences similar to other AI assistants but are tailored toward helping I&O workers increase their productivity.

Business Impact

- Optimized and improved IT support processes through insight and automation
- Faster resolutions and improved accuracy in triage, categorization and expert identification
- Easy access to deeper insight by analyzing processes, tickets and workflows to identify opportunities for improvement
- Improved experience for human IT agents, leading to fewer frustrations and improved morale

Drivers

- While business-consumer-facing virtual support agents have previously been the main focus, heads of I&O are now prioritizing AI assistant solutions that deliver similar user experiences for human IT agents, aiming to enhance their productivity.
- Some organizations have purchased or are piloting solutions such as Microsoft Copilot or Anthropic's Claude Cowork and hope to use similar technology within I&O.
- Increasing numbers of vendors in IT operations management (ITOM) and AI applications in ITSM markets have branded part of their offering as AI assistants.
- Some heads of I&O believe that conversational interfaces will become a primary method of agent interaction with IT operations and service management software.
- Heads of I&O believe that AI-enabled assistant technology can help them meet the challenges of macroeconomic conditions to find cost savings through increases in productivity.

Obstacles

- General-purpose conversational AI assistants primarily designed for meeting transcription and summarization use cases are not suited to be an operations assistant to facilitate I&O workflows.
- Operations assistants are too new to be considered as a must-have. Consequently, the high cost of current commercial solutions is hard to justify for experimental use.
- Human IT agents may not want a simplified conversational interface to carry out their tasks, instead preferring fully featured forms or command line user interface (UI).
- IT agent advisory capabilities do not require an operations assistant; instead, they more commonly show up as other UIs, such as the classic web form UI of ITSM platforms.
- The term “copilot” has seen reduced use by vendors from 2025 to 2026, as it has become more associated with Microsoft’s offering than a technology category.
- Operations assistants rely on other underpinning AI capabilities, such as agent advisory.

User Recommendations

- Seek operational assistants that are designed for ITOM and ITSM use cases because general-purpose AI assistants would need extensive customization.
- Confirm that existing IT agents can benefit from conversational interfaces rather than preferring traditional user and command line interfaces to interact with tools.
- Do not rely on operations assistants to replace domain specialists. Ensure that novice or new target user groups have sufficient domain knowledge to identify erroneous generative output.
- Since stand-alone operations assistant products have not emerged into the mainstream and do not contribute to a drop in overhead, heads of I&O should wait before committing to emergent products.

Sample Vendors

Automation Anywhere; BigPanda; Freshworks; OpenText; PagerDuty, ServiceNow

Gartner Recommended Reading

[Magic Quadrant for Artificial Intelligence Applications in IT Service Management](#)

[Critical Capabilities for Artificial Intelligence Applications in IT Service Management](#)

[AI Use-Case Assessment for IT Service Desk](#)

AI Applications in IT Service Management

Analysis By: Chris Matchett

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Adolescent

Definition:

AI applications in IT service management are tools that augment and enhance IT service management (ITSM) workflows using AI. These analyze ITSM data and metadata to provide intelligent advice and execute actions on ITSM practices and workflows, such as IT service desk and support activities. The software can be delivered as an incumbent ITSM platform with integrated AI features, a native AI add-on, or a third-party product that integrates with an ITSM platform.

Why This Is Important

Heads of I&O seek the benefits of AI, but they are not replacing ITSM platforms to obtain AI capabilities. Instead, they choose between incumbent ITSM platform capabilities and third-party, specialist AI solutions. The market is continually evolving, driven by advances in AI and natural language technologies. AI applications in ITSM are distinguished by their reliance on ITSM data and metadata as the primary input, ensuring alignment with ITSM practices.

Business Impact

- Faster resolution and improved accuracy in triage, categorization and expert identification.
- Tangible cost reductions and labor savings through automation of support issues and requests.
- Improved employee-facing user experience and stronger business consumer relationships.
- Deeper insight into service, practice and staff performance.
- AI advisory, automation and agentic features to help avoid disruptions, reduce downtime and improve IT service reliability.

Drivers

- The popularity of GenAI is prompting heads of I&O to have higher expectations of conversational AI platforms and look beyond previous-generation chatbot features in ITSM platforms.
- Buyer desire for agentic automation and “zero-touch” IT support has accelerated, and vendors are promoting their AI applications in ITSM products that offer agentic ITSM features.
- Buyer requirements have expanded beyond virtual support agents (VSAs) for self-service to include IT agent advisory, with a dual focus on improving both employee experience and ITSM team efficiency.
- ITSM platforms provide a variety of AI features, but heads of I&O often purchase add-ons or integrate stand-alone products from third-party vendors when they require more extensive capabilities.
- GenAI capabilities are increasingly sought by heads of I&O to automate content generation and improve communications. Examples include summarizing information, such as knowledge base articles or case work-log updates, and generating major incident notifications.
- This market saw several significant acquisitions in late 2025, which reflect a trend of vendors seeking to acquire missing capabilities and strengthen their ITSM automation and workflow orchestration solutions.

Obstacles

- AI applications in ITSM products have room for significant improvement, notably in machine learning analytics and agentic ITSM.
- I&O groups lack skills in areas such as prompt engineering to accelerate deployments using LLMs.
- I&O organizations lack sufficient foundational data to train their AI capabilities — whether consistently categorized incidents or good-quality knowledge for their virtual support agent.
- As heads of I&O must choose between incumbent ITSM platforms and add-on products specialized in AI, several smaller vendors have been acquired and merged into larger offerings in this market, reducing choice for buyers.
- Most AI applications in ITSM have yet to deliver on agentic ITSM capabilities beyond marketing hype.
- Costs associated with acquiring and implementing AI applications in ITSM are generally considered high and unpredictable.

User Recommendations

- Use the [AI Use-Case Assessment for IT Service Desk](#) to determine which AI use cases will provide the most transformational value, while being sufficiently feasible to implement quickly.
- Invest in AI applications for ITSM to achieve significant enhancements across a broad range of ITSM practices.
- Consider general-purpose conversational AI solutions for commodity features, such as using a virtual support agent to deflect calls from the service desk.
- Verify whether the incumbent ITSM platform can already meet these needs. If an add-on or upgrade is required, then assess the licensing and implementation costs before comparing with third-party solutions.

Sample Vendors

Atlassian; Automation Anywhere; BMC Helix; ServiceNow

Gartner Recommended Reading

[Market Overview for AI Applications in IT Service Management](#)

[Magic Quadrant for AI Applications in IT Service Management](#)

[Critical Capabilities for AI Applications in IT Service Management](#)

[AI Use-Case Assessment for IT Service Desk](#)

[Infographic: Choose an AI Solution for ITSM](#)

At the Peak

AI-Powered IT Agent Advisory

Analysis By: Chris Matchett, Siddharth Shetty

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

AI-powered IT agent advisory is the use of AI to analyze ITSM data and metadata and generate recommendations that accelerate and enhance IT agent activities. Example features include intelligent categorization, triage, routing, escalation, swarming, risk advisory and sentiment analysis.

Why This Is Important

IT agent advisory features are a core capability of AI applications for ITSM that differentiate them from business consumer-facing conversational assistants such as virtual support agents. This expands capabilities to meet the needs of IT agents through advice and actions embedded in the UI of an ITSM platform or via an operations assistant conversational interface.

Business Impact

- Enhances decision making, helping with IT support tasks
- Reduces costs through labor savings and faster resolutions
- Improves accuracy in triage, classification, swarming and routing
- Provides change-risk advisory to avoid disruptions and ensure reliability
- Sentiment analysis improves service experiences
- Frees up human IT agents' time from routine tasks
- Optimizes ITSM practices through actionable insights

Drivers

- Heads of I&O believe that AI can help them find cost savings through automation, leading to interest in AI for IT agent use cases and capabilities such as agent advisory.
- AI can be used to improve the operational efficiency of I&O teams using agent advisory features from ITOM vendors.
- AI technologies, such as GenAI and LLMs that enable AI-powered IT agent advisory features, are continually evolving.

Obstacles

- Products on the market have significant room for improvement, particularly in areas such as intelligent escalation, which is part of agent advisory (see [Critical Capabilities for AI Applications in IT Service Management](#)).
- I&O organizations often lack sufficient foundational data (e.g., consistently categorized incidents or knowledge) to train their AI capabilities.
- AI skills such as prompt engineering are needed to maximize the benefits of agent advisory, but are not yet sufficiently developed within I&O.
- AI-powered IT agent advisory often requires an upgrade or additional software purchase that heads of I&O can struggle to justify without evidence of proven ROI.
- Many AI applications in ITSM rely on LLMs for AI but lack machine learning techniques like cluster analysis and pattern recognition, which are needed for the full range of IT agent advisory features.

User Recommendations

- Use [AI Use-Case Assessment for IT Service Desk](#) to determine which AI use cases will provide the most transformational value and are feasible for quick implementation.
- Invest in AI applications for ITSM to gain AI-powered IT agent advisory features. Pair this capability with an operations assistant to use a conversational interface to leverage the advisory.
- Verify whether your incumbent ITSM platform already meets these needs. If an add-on or upgrade is required, assess the licensing and implementation costs before comparing them with third-party solutions.

- Optimize existing processes and remove bottlenecks before using AI-powered IT agent advisory to optimize ITSM practices, as AI will not fix a fundamentally broken process.
- Treat IT agent advisory as suggestions and maintain a human-in-the-loop verification of the results to ensure quality and reliability.

Sample Vendors

Automation Anywhere; BMC Helix; ServiceNow

Gartner Recommended Reading

[AI Use-Case Assessment for IT Service Desk](#)

[Magic Quadrant for AI Applications in IT Service Management](#)

[Critical Capabilities for AI Applications in IT Service Management](#)

[How to Achieve Success With AI at the Service Desk](#)

IT Knowledge Generation

Analysis By: Rich Doheny

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

IT knowledge generation creates knowledge articles with GenAI by summarizing attempted and successful solutions from sources, formatted in accordance with enterprise style guidelines and templates. These sources can include organic conversations, documents, and historical cases. These articles can then be leveraged in a management system (e.g., ITSM platform or conversational AI tool) for both IT and end-user consumption.

Why This Is Important

An effective knowledge base enhances IT self-service adoption and IT support staff efficiency, saving I&O time and money. Most development and ongoing maintenance of the knowledge assets is manual, relying on time-intensive human effort after an issue is resolved. Therefore, many organizations struggle to deliver an effective and sustainable knowledge management (KM) practice. By leveraging GenAI techniques, heads of I&O can automate the creation and maintenance of knowledge assets.

Business Impact

An effective KM system improves IT support quality and self-service. Heads of I&O can adopt GenAI for knowledge creation to:

- Accelerate the authoring and publication process
- Identify IT knowledge gaps
- Increase the number of useful knowledge assets
- Deliver knowledge-article content and format standards
- Reduce IT support overhead

Drivers

- Heads of I&O have been looking to automate the knowledge creation process for years, and the proliferation of GenAI now makes that possible.
- Given the manual nature of building and maintaining knowledge articles, I&O leaders find it challenging to keep IT staff engaged..
- Heads of I&O want more responsive KM practices to effectively support the proliferation of technology caused by digital business initiatives.
- Product teams have led to more distributed and federated support models. Automated knowledge generation can reduce the risk of knowledge silos by simplifying the process of knowledge authoring.
- AI applications in ITSM increasingly provide support for knowledge generation.
- Deep research techniques offer the promise of drafting articles sourced from multiple sources with full attribution, without relying on extensive case documentation.
- ITSM platform vendors, AI apps in ITSM, and conversational AI providers are increasingly embedding knowledge-generation features in their products.

Obstacles

- IT support teams often lack ticket documentation standards, which will reduce the value of knowledge generation.
- Automating the production of knowledge can increase the risk of knowledge bloat through low-value or redundant knowledge due to poor management controls on newly generated knowledge.
- Few knowledge-generation solutions automate the publishing pipeline and life cycle management for knowledge, as most solutions only focus on creating the knowledge asset text.
- LLMs risk misinterpreting the data they are fed, causing hallucinations (i.e., making up facts).
- There are data privacy, data ownership and security concerns about LLM data usage.
- Solutions offering this capability are often bundled with other GenAI features that are priced at a premium, requiring a broader strategy to drive the ROI.
- Many organizations use third-party support services that may not be incentivized for ticket documentation and/or KB creation.

User Recommendations

- Select an AI-enabled knowledge-generation solution that is integrated into the publishing life cycle and knowledge base and provides domain-trained LLMs.
- Have a human review and validate any AI-generated content, even previously published information that is only being updated.
- Investigate how GenAI can flag related knowledge and automate the maintenance of existing knowledge to keep knowledge assets current.
- Identify areas to automate the capture of information for ticket documentation, such as call transcripts or log data, which can be fed into LLMs to build knowledge drafts.
- Avoid dissolving the existing KM practice; instead, evolve it to account for system-generated knowledge.
- Encourage knowledge managers to work closely with the AI platform teams.
- Ensure that contracts with external support partners include terms to support knowledge capture and knowledge portability at the termination of the agreement.

Gartner Recommended Reading

[Revitalize IT Support With GenAI-Powered Knowledge Management](#)

Service Operations

Analysis By: Chris Laske

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Service operations capability focuses on the provision of highly available and performant services by detecting and responding to potential or actual interruptions to a service or a reduction in service quality. This is achieved through converging observability practices with incident management practices to create a dynamic and responsive end-to-end discipline that more effectively diagnoses and resolves incidents.

Why This Is Important

Rapid incident diagnosis and resolution are vital for effective management of digital services. By leveraging technologies like AI and automation, traditional silos between monitoring and the service desk can be dismantled. This integration creates unified service operations, closely linking event detection through observability with IT service management (ITSM) platforms. The result is improved incident identification, faster resolution, and higher service availability and resilience.

Business Impact

The convergence of ITSM and observability results in a more productive and seamless incident management practice. The contextual analysis of alerts resulting in anomaly detection and escalation can result in fewer incidents and improved mean time to resolve (MTTR), especially when integrated with the ITSM platform. Products and services have higher availability and are less likely to suffer from poor performance or outages and disruption, costs are reduced, and overall productivity is improved.

Drivers

- Business pressure for reliable and stable services drives the demand for incidents to be either eliminated before they can cause disruption or diagnosed and resolved as quickly as possible.
- The closer integration of observability with the service desk creates a far more effective and seamless environment for managing outages and disruption.
- Many ITSM platforms now feature observability and correlation functions to identify incidents as part of their newly available monitoring capabilities.
- Event intelligence platforms are leveraging AI/ML to perform predictive analysis, proactively providing guidance regarding the potential for an event to become a business-impacting incident.
- Staff can operate more efficiently as one team of specialists collaborating to manage the environment.
- ITSM platforms enrich the events by applying the context of recent incidents and changes to the impacted configuration items (CIs) to radically improve the identification of the probable root cause.
- The use of AI allows anomalies and events from different monitoring systems to be swiftly correlated, analyzed and presented to technicians.

Obstacles

- AI-based monitoring systems can take longer than expected to understand the context of the services, the traffic and the environment, creating delays in improving availability.
- Achieving improved diagnostics and associated resolution times requires time and effort because of the need to better manage dependencies, including effective service mapping, the degree of automation and the monitoring quality.
- Aligning data feeds and telemetry from different sources may result in data disparity, leading to delays in correlating incidents.
- Accurate service mapping and comprehensive, up-to-date configuration management databases are necessary to support end-to-end incident processes needed to achieve enhanced availability and resilience.
- Automated and AI-guided responses generally require human oversight to ensure that their decisions match the use case.
- Many monitoring and service desk teams operate in silos, making it difficult to create a collaborative and seamless experience.

User Recommendations

- Develop a service operations strategy founded on business priorities that focuses on effective incident management. Consider the implications on the monitoring and support staff (including the service desk), their integration strategy, and the need to adapt service practices.
- Identify the appropriate tool strategy and define the end-to-end process, including dependencies, and use this to build a roadmap to deploy and implement the solution.
- Develop runbooks and deploy AI agents to provide scripted solutions for common incident scenarios.
- Explore vendor options for solutions offering AI, automated responses and self-healing to ensure alignment and integration with your environment.
- Adopt a pragmatic approach and pilot the initiative to understand the capabilities and drawbacks.
- Establish aligned goals and shared KPIs to bring the teams together and challenge them to improve service levels and technology availability.

Sample Vendors

BigPanda; BMC Helix; Datadog; Dynatrace; Grokstream; LogicMonitor, PagerDuty; ScienceLogic; ServiceNow

Gartner Recommended Reading

[Magic Quadrant for Observability Platforms](#)

Multiagent Systems

Analysis By: Leinar Ramos, Anthony Mullen, Pieter den Hamer

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Multiagent systems (MAS) are collections of AI agents that interact to achieve individual or shared goals. These agents can operate within a single environment or be independently developed and deployed across distributed environments.

Why This Is Important

Multiagent systems are a powerful design architecture for managing complex workflows. By breaking workflows into modular components, these systems enable individual AI agents to specialize in specific decisions and tasks. Applying multiple AI agents together allows them to tackle complex tasks that individual agents cannot while creating more adaptable, scalable and robust solutions.

Business Impact

MAS can be used in:

- **Software development** for automating complex tasks across the software delivery life cycle.
- **Complex business workflows** such as customer service, marketing and sales.
- **Robotics** for warehouse optimization, search and rescue, and environmental monitoring.
- **Supply chain** operations for scheduling, planning, routing and supply chain optimization.
- **Transportation** for traffic flow optimization and autonomous vehicle coordination.
- **Telecom** for network optimization and fault detection.

Drivers

- **Evolution of MAS:** MAS can be built on a single platform, across platforms or across the internet, forming networks of AI agents. These emerging design patterns provide flexibility and access to diverse skills and tools as agents form dynamic collaborations to address complex challenges.
- **Multiagent frameworks:** The rise of multiagent frameworks is increasing the feasibility of experimenting with and deploying these systems, particularly those based on LLM-powered agents. These frameworks simplify the creation, orchestration and management of multiple AI agents.
- **Limitations of single AI agents:** Current AI agents are not reliable enough to perform well across a broad set of tasks. As a result, it is often more effective to break a process into smaller tasks and assign each task to a narrow, specialized agent coordinated through MAS.
- **Increased decision-making complexity:** AI is increasingly used in real-world engineering problems that involve complex systems, where large networks of interacting components exhibit emergent behavior that is difficult to predict. The decentralized nature of MAS makes them more resilient and adaptable to complex decision making.
- **Agent communication protocols:** The emergence of agent-to-agent communication standards and protocols is increasing potential interoperability among agents built on different platforms.
- **Critic agents:** Agents that apply standards and guardrails can evaluate workflow quality and determine whether to proceed to the next step or execute final actions, enhancing overall decision-making quality and task execution.

Obstacles

- **Training complexity:** MAS are harder to train and build than individual AI agents. These systems can exhibit emergent behavior that is difficult to predict in advance, increasing the need for robust training and testing.
- **Monitoring and governing multiple agents:** Coordinating and collaborating across agents is challenging. Effective oversight requires careful monitoring, governance and a shared grounding to ensure the system behaves as intended.
- **Reliability:** Multiagent approaches without some form of centralized planning are often unreliable. Successful implementations typically enforce tighter workflow control across agents, which improves reliability but reduces system flexibility.

User Recommendations

- Use MAS for complex problems that single AI agents cannot solve, including tasks that require multiple perception steps, decisions and actions to achieve higher accuracy. Break each step of the workflow into modular tasks to produce accurate results across complex processes.
- Shift to a multiagent approach gradually since this is an emerging research area and its risks and benefits are not yet fully understood.
- Invest in technologies that support collaboration among AI agents to harness the full potential of MAS as cross-platform agent capabilities evolve.
- Establish clear guardrails when implementing MAS, including legal and ethical guidelines on autonomy and liability, as well as robust security and data privacy measures.
- Educate AI teams on MAS: how they differ from single-agent designs and the frameworks available to build and manage these systems.

Sample Vendors

Amazon Web Services; CrewAI; Google; LangChain; Maisa AI; Microsoft; OneReach; Openstream.ai; Salesforce; Thunk.AI

Gartner Recommended Reading

[Case Study: Multiple AI Agents' Recommendations Improve Efficiency](#)

[Innovation Insight: AI Agents](#)

Innovation Insight: AI Simulation

XLA

Analysis By: Karl Rosander

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Experience-level agreements (XLAs) are designed to measure and improve the quality, performance, and experience of technology and related services. By leveraging and analyzing a combination of technology monitoring, sentiment analysis, and traditional service-level metrics to monitor the timeliness and effectiveness of supporting processes, XLAs support continuous improvement and the achievement of business outcomes.

Why This Is Important

Modern technologies and their associated services are often delivered by multiple parties, both internal and external to an organization, which requires collaboration for a shared vision of success. XLAs allow organizations to measure and analyze performance of technology and services, providing actionable insights focused on improved digital experiences.

Business Impact

The goal of XLAs is to achieve tangible results in achieving business outcomes. Traditional SLA measures, which are often transactional in nature, center around operational results and have limited success in measuring the entire experience. Actual business outcomes, measuring technology, services, and sentiment beyond transactions, can include lower attrition, increased productivity, cost optimization, digital adoption, and much more.

Drivers

- An XLA creates a composite of metrics for various KPIs that are analyzed to provide recommended actions. When acted upon, XLAs will result in improved experience. The underlying metrics should measure individual factors that make up an experience, such as network performance, application workflows, and the sentiment of provided technology and tools.
- Traditional IT-centric performance measures fail to capture the total experience (TX) properly, as they rely on specific data points (generally when an incident occurs, or a service request is made) that are a fractionable component of the TX. Outcome-centric experience measures are required to fix this challenge.
- Digital experience tools can measure each touchpoint in the overall user journey, which enables the XLA implementation and drives the motivation for capturing the touchpoints. These tools should also incorporate some measure of user and stakeholder feedback.
- The need for business-outcome-based performance measures drives XLA adoption. XLAs enable the impact of IT systems and services to be mapped directly to business KPIs, which allows the business to create outsourced service deals that contract for defined business outcomes. XLAs deliver more executive-level metrics, like productivity improvement, talent retention, cost optimization.
- Outsourcing initiatives must demonstrate business value. Decisions about specific metrics to use should be made through co-creation with providers and internal stakeholders. The aim is to bridge the gap between the intended business performance improvements and the technologies and services that will be delivered and measured through the XLA metrics.
- The impact of technology, such as DEX and AI tools, capture the experience of various signals and support the analysis of experience-related challenges, making the adoption of XLAs easier.

Obstacles

- **Lack of awareness:** Many enterprises are unaware of how XLAs are contracted. This approach often misses the added value of prioritizing employee experience over operational efficiency.
- **Resistance to change:** Some enterprises are unwilling to embrace the workplace transformation required to improve the user experience. Champions of XLAs often encounter resistance where buy-in is critical.
- **Tooling costs and complexity:** High initial implementation costs, due to tools deployment, are limiting factors for clients that focus more on cost than value. Implementing XLAs often requires extra preparation and due diligence, along with change management, which can discourage clients from taking the next steps.
- **Variability in implementation:** When it comes to contracting for XLAs, no two vendors use the same framework for how to measure or how to be held accountable for outcomes. This context makes it difficult for end clients to harmonize the different approaches.

User Recommendations

- Identify and configure XLAs that measure end-to-end experiences instead of individual metrics (e.g., “speed of service,” which is a composite metric). The key to success is to implement monitoring and to identify bottlenecks that focus on the employee experience, instead of just IT deliverables.
- Select partners that have demonstrated referenceable business improvement with a strong track record in process mapping, analytics, and digital experience monitoring
- Improve internal operations by collaboratively identifying and addressing the leading causes of dissatisfaction.
- Consider implementing an experience management office (XMO) or an assigned team for managing, monitoring and evolving the XLAs. Work closely with providers and/or internal organizations to identify XLAs, develop improvement plans, and provide stakeholder involvement and governance during implementation.
- Ensure that contracts are structured properly to accommodate the changes required to govern and implement XLAs, as well as to reward and penalize regimes.
- Identify and select appropriate toolsets to trace the XLA data point.
- Get started using XLAs more quickly by implementing the recommendations found in Gartner research: [5 Steps for a Simple Start to Implement XLAs](#).

Gartner Recommended Reading

[Drive Valuable IT Services Outcomes With a Standardized XLA Framework](#)

[5 Steps for a Simple Start to Implement XLAs](#)

[Establish a DEX Management Office to Overcome Digital Friction](#)

[Build a Case for a DEX Management Office to Overcome Digital Friction](#)

Autonomous Endpoint Management

Analysis By: Tom Cipolla

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Autonomous endpoint management (AEM) is a powerful approach enabled by functionality within advanced endpoint management tools. AEM leverages configuration, compliance, risk, performance, and experience data to intelligently perform common endpoint management tasks. The most common use case for AEM is autonomous patching that accelerates patch deployment and compliance, which reduces IT overhead and degradation of digital employee experience (DEX).

Why This Is Important

Digital workplace leaders are simultaneously pressured to increase the velocity of endpoint management (patching, configuration management, and life cycle management) and maintain DEX. AEM is accelerating endpoint management and is set to streamline and automate operations while protecting DEX and reducing IT overhead.

Business Impact

Digital workplace leaders can further automate endpoint and DEX management tasks and reallocate efforts toward business-value-added work. Specific impacts include:

- Reduced IT overhead through automatic resolution of operations issues that impede employee productivity.
- Reduced cyber risk by automating patch and configuration management.
- Automated software and configuration deployment based on policy and persona.

Drivers

- IT staff is overwhelmed with the growing number of endpoint devices, operating systems, and applications.
- Largely driven by increased AI capabilities, application vendors have accelerated development and release cadence, and IT cannot keep pace.
- Increased vulnerabilities and greater risk exposure demand faster patch deployment, better device configuration compliance, and closer alignment with vendor life cycles to reduce vulnerabilities.
- Adoption of DEX practices and tools continues to grow rapidly as many accelerate their focus from a technology-centric to an employee-centric experience.
- Cloud-based endpoint management and DEX tools are demonstrating how AI- or ML-powered intelligence can quickly process a significant amount of data, provide actionable insights and recommendations, and execute automations.
- Developing and maintaining automation for common administrative tasks and applying standard policies and configurations are time-consuming and require integration across multiple tools.
- Digital workplace leaders struggle with utilizing and consolidating data from other endpoint management tools and agents.
- AEM directly supports the end-user services leader's goal of speed and agility.
- AEM use cases are promising in addressing the management of applications and replacing human execution of routine IT processes.

Obstacles

- Overly complex environments with too many disparate tools lack integration.
- Highly customized environments require extensive testing of every update prior to deployment.
- Fragile environments have a significant amount of technical debt — including legacy operating systems or applications that depend on unsupported browsers, runtime environments, or plug-ins.
- Low-to-mid maturity organizations lack the foundational processes and concepts to adopt AEM.
- Device operating system limitations or controls may prohibit experience and automation capabilities.
- AEM is not possible on-premises, so cloud-averse organizations will not be supported.
- Organizations that lack experience with agile methodologies and automation skills operate with a legacy mindset prioritizing control and customization.

User Recommendations

- Avoid vendor lock-in by ensuring strategic endpoint and DEX management vendors have a roadmap that directly provides or includes necessary partnerships to enable AEM.
- Reduce location dependence by migrating experience and endpoint management, and security and identity solutions to the cloud.
- Prepare your organization by assessing current and future skills requirements, updating existing roles and defining new ones, and implementing strategies for upskilling and professional development.
- Eliminate inertia by promoting a human-centric and enablement mindset, while adopting modern management principles and agile methods.

Gartner Recommended Reading

[Magic Quadrant for Endpoint Management Tools](#)

[Critical Capabilities for Endpoint Management Tools](#)

Innovation Insight: Autonomous Endpoint Management

AI Agents

Analysis By: Tom Coshow, Haritha Khandabattu

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

AI agents are autonomous or semiautonomous software entities that use AI techniques to perceive, make decisions, take actions and achieve goals in their digital or physical environments.

Why This Is Important

AI agents have the ability to make decisions and take action in their target environment to achieve organizational goals. By using AI practices and techniques such as LLMs, organizations are creating and deploying AI agents to achieve complex tasks.

Business Impact

AI agents have the potential to:

- Revolutionize a broad range of industries and environments with their ability to automate tasks from consumer, industrial, data analytics, content creation and logistics.
- Make informed decisions and interact intelligently with their surroundings.

Drivers

- **Generative AI breakthroughs:** Reasoning models and LAMs advance the ability to plan a complex series of actions.
- **Multimodal understanding:** The ability to use diverse modalities like vision, audio and language enables more general and flexible AI agents. This allows automatic adaptation to changes in the workflow, user interface or API. With this, one can create advanced workflows without explicit programming, significantly reducing the development time and effort for automation.
- **Increased decision-making complexity:** AI is increasingly used in real-world engineering problems containing complex systems, where large networks of interacting parts exhibit emergent behavior that cannot be easily predicted. AI agents can learn, plan and execute in complex environments.
- **Composite AI, including neurosymbolic models:** Advances in models that improve planning and problem solving are enabling more complex AI agents. AI agents can utilize a wide variety of AI practices to forecast, make decisions and plan.

Obstacles

- **Vulnerabilities:** Due to the complexity of the AI agent system, all components face various potential vulnerabilities such as access security, data security and governance.
- **Lack of trust:** Users are unsure whether they can trust the technology to accurately predict and execute tasks independently. Without a human in the loop, agents may take multiple consequential actions in rapid succession and bring about significant impacts before a human notices.
- **Interpretability and oversight:** Action policies may be opaque and have poor explainability, requiring mechanisms for human interpretability, oversight and control.
- **Pace of change:** The technologies deployed, from models to tooling, and the framework options available are changing rapidly, making it difficult for organizations to define their roadmap.

User Recommendations

- Incorporate AI agents into strategic planning by investing in understanding their capabilities and potential applications in various environments, considering their increasing autonomy and wide-ranging usability.
- Investigate the possibilities of utilizing multiagent systems, collectives of AI agents, that can operate both collaboratively and independently, enhancing adaptability and flexibility in response to different tasks and scenarios.
- Promote the development and integration of the use of a variety of AI practices, enabling learning, negotiation and decision-making capabilities.

Sample Vendors

Amazon; Anthropic; CrewAI; Google; LangChain; Maisa; Microsoft; OneReach.ai; OpenAI; Salesforce

Gartner Recommended Reading

[Innovation Insight: AI Agents](#)

[Build AI Agent Services to Revolutionize Client Operations](#)

[Innovation Insight for the AI Agent Platform Landscape](#)

[AI Agents: Are You Ready to Set Your AI Free?](#)

[Emerging Patterns for Building LLM-Based AI Agents](#)

Enterprise Service Management

Analysis By: Siddharth Shetty

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Enterprise service management (ESM) applies service-management principles, tools and best practices to enhance service delivery and support for non-IT functions, such as HR, facilities and legal.

Why This Is Important

ESM helps organizations harmonize the enterprise support experience for business users. By reducing application sprawl, ESM streamlines operations, driving stronger cross-functional collaboration and the sharing of best practices between IT and non-IT functions. Clients evaluating ITSM vendors increasingly assess the vendor's ability to support ESM across their organization.

Business Impact

Traditional enterprise service delivery models for IT, HR, facilities, legal and other functions rely on separate tools and portals, inconsistent user experience (UX), and varied automation levels. ESM reduces application sprawl and harmonizes the UX when consuming enterprise services by using a shared platform with similar support channels. This harmonized approach to service delivery drives stronger collaboration across support units and a more consistent user experience for business users.

Drivers

- Most prominent ITSM vendors actively position ESM as part of their marketing strategies to extend workflow and self-service capabilities to other support functions.
- Heads of IT and lines of business (LOBs) increasingly call on IT leaders to lead ESM initiatives that consolidate best practices in end-user support and service orchestration to deliver enterprise services beyond IT.
- Reducing application sprawl helps address inconsistent user experiences, user fatigue and the talent requirements needed to support enterprise service tools.
- Opportunities to automate manual workloads across enterprise support groups shorten service delivery times and free up capacity for support teams.
- Traditionally, each support function used a separate channel to support end users. ESM provides multichannel support through a harmonized interface, giving business users choices that align with demographic diversity and hybrid workplace needs.
- Employees often experience fragmented service delivery when services rely on multiple disparate groups, such as onboarding and offboarding, due to limited coordination among support teams.
- ESM fosters collaboration across enterprise functions. Processes such as employee onboarding, service reporting and service-level management benefit from improved coordination and shared accountability across functions.
- The growing interest in AI drives leaders in non-IT departments to leverage IT's AI expertise to accelerate AI adoption through ESM.
- CRM vendors are increasingly expanding into the ITSM market with emerging ESM capabilities.

Obstacles

- ESM programs are usually pitched as a “cost-savings” exercise. The primary value of an ESM program lies in its ability to drive better experiences for business users and foster stronger collaboration across support units.
- ESM programs often focus primarily on tool or platform selection, which delivers limited value unless organizations concurrently define business outcomes and develop a program roadmap.
- Vendor definitions of ESM closely align with their respective market offerings, resulting in a general lack of standards and a clearly defined scope for the ESM market and its capabilities.
- Clients often underestimate the organizational change management and collaboration efforts required to implement an ESM program, which can undermine program success and erode stakeholder confidence.
- ESM solutions must support data segmentation for sensitive and confidential information, particularly within HR, legal and finance functions.

User Recommendations

- Start by defining the ESM strategy and business objectives that incorporate both short-term and long-term demand for enterprise services. This should include the current business challenges that ESM must address.
- Involve stakeholders early from other enterprise functions to address organizational change management challenges. Secure buy-in for a cross-functional collaboration team that develops the ESM strategy, program plans and business communications.
- Identify ESM practice capabilities that support the ESM strategy, including how the organization will address process, workflow and collaboration challenges.
- Use ESM practice capabilities to inform platform selection and reduce the risk of overbuying functionality.
- Identify business value in terms of digital UX and customer satisfaction metrics to complement ROI-based justifications, rather than relying on ROI alone.

Sample Vendors

Atlassian; BMC Helix; EasyVista; Freshworks; Ivanti; ManageEngine; ServiceNow

Gartner Recommended Reading

[How to Build a Successful Enterprise Service Management Program](#)

[4 Key Criteria for Selecting Your Enterprise Service Management Solution](#)

[Determine If Enterprise Service Management Is Right for Your Organization](#)

[Market Guide for IT Service Management Platforms](#)

Threat Exposure Management

Analysis By: Pete Shoard, Jeremy D'Hoinne, Mitchell Schneider

Benefit Rating: Transformational

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Threat exposure management encompasses processes and technologies that allow enterprises to continually and consistently assess the visibility and validate the accessibility and exploitability of their digital assets. It must be governed by an effective continuous threat exposure management (CTEM) program.

Why This Is Important

The diversity of modern infrastructure affects organizations' ability to accurately assess cyber risks. Security teams often struggle to identify modern exposure types, leaving security gaps due to accelerated AI usage and lack of Identity oversight, SaaS and cyber-physical systems security. Threat exposure management addresses these challenges by enabling identification, prioritization and validation of issues across diverse attack surfaces and ensuring comprehensive visibility and mitigation.

Business Impact

Threat exposure management governs and prioritizes risk reduction for the modern enterprise. It requires assessment of all business-related systems, applications and subscriptions, broadening risk understanding for today's digital landscape. CTEM programs factor in business importance, likelihood of attack, visibility of vulnerability and validation of the existence of an attack path, enabling businesses to mobilize responses to genuine, impactful risks.

Drivers

- Organizations seek to upgrade vulnerability management programs to demonstrate to the business that cyber risk is more pervasive than simply traditional IT assets that require patching. For instance, there are new threats and increased exposure from rapid technology development in areas such as AI. Threat exposure management provides business alignment through the scoping process and helps reduce volumes of irrelevant or nonbusiness-critical issues.
- Exposure is a factor of not only where vulnerabilities exist, but also where impacts to normal business processes can be most greatly felt. Threat exposure management aims to allow reprioritization of treatments as environments shift in a rapidly changing and expanding IT landscape.
- Organizations commonly silo exposure activities, such as penetration testing, threat intelligence management and vulnerability scanning. Siloed views provide little or no awareness of the complete picture of cyber risk.
- Modern technologies, application development and innovation efforts with AI, as well as the increasingly pervasive use of third-party SaaS applications, is making organizations susceptible to invisible and high-impact exposures. Vulnerability management approaches have focused too much on systems and software-based exposures, where a much wider set of visibility and new remedial responses are required to reduce the exposure to new threats.
- Vendor offerings to identify threat exposures have consolidated into exposure assessment platforms (EAP), offering broader visibility into exposures as standard. End users now frequently consider nonpatchable technology issues as potential threats and are planning response actions for new scenarios.

Obstacles

- The increased scope of CTEM programs over traditional VM introduces many new complexities often not previously considered or budgeted for.
- While evaluating new exposures is necessary, building effective response processes and the ability to mobilize a gradient of countermeasures, such as threat monitoring and control configuration, is lacking. Many organizations are not prepared to respond to modern threat scenarios.
- Processes to manage end-to-end awareness (from visibility of attack vectors to response to breaches) are virtually nonexistent in most organizations, which often simply scan their networks for compliance reasons. Regulations rarely factor in the exploitability of exposures.
- Assessing the complexity of attacks requires new skill sets. Market areas, such as adversarial exposure validation (AEV), make it simple to test the out-of-the-box scenarios using simulation tools. But users need new skills to be effective at using these capabilities and customizing scenarios.

User Recommendations

- Evolve vulnerability management programs with CTEM to manage a wider set of exposures. Focus on business importance when building target scopes. Don't scan for or try to find every possible exposure, but drive toward reducing risk in areas of significant business importance.
- Preagree routes to mobilize responses to exposures with a variety of stakeholders, as success depends on it. Automated remediation of gaps in security posture remains unlikely to have a significant impact due to end-user trust in solutions and the complexity of resolution paths.
- Communicate exposure risk to the board. Executives must be made aware of changes to exposure and be encouraged to allocate adequate resources to proactive security.
- Validate that exposures genuinely exist and that security controls are functioning as expected with automated testing tools such as AEV.
- Include assets that your organization doesn't directly own, such as social media accounts, SaaS and data held by supply chain partners.

Gartner Recommended Reading

[Reduce Threat Exposure With Security Controls Optimization](#)

[Market Guide for Adversarial Exposure Validation](#)

[Key Differences Between CTEM and Vulnerability Management](#)

[Strategic Roadmap for Continuous Threat Exposure Management](#)

[Reference Architecture Brief: Exposure Management](#)

Sliding into the Trough

Infrastructure Platform Engineering

Analysis By: Hassan Ennaciri, Paul Delory

Benefit Rating: Transformational

Market Penetration: More than 50% of target audience

Maturity: Early mainstream

Definition:

Infrastructure platform engineering (IPE) is the discipline of building infrastructure platforms as products that abstract complex IT infrastructure, making it easily consumable for users and other systems. Infrastructure platforms provide self-service tools to deploy and manage infrastructure, while I&O retains governance, security and compliance. These platforms form the foundation for higher-order self-service layers such as internal developer platforms.

Why This Is Important

Organizations are pressured to innovate and deliver products faster while mitigating costs to meet customer needs. This requires adopting new operating models and modern practices to deliver scalable, reliable platforms that enable faster product delivery. IPE provides automated delivery of curated, secure, reliable and scalable infrastructure services. These are made available via self-service or APIs, and reduce the effort and cycle time for users to request and access the products.

Business Impact

IPE simplifies complex infrastructure, delivering evolving platforms as products to meet customer needs. This agile approach empowers developers with self-service tools for on-demand access to environments, services and tools, boosting productivity and improving customer experience. It balances agility with control, enabling efficient, governed infrastructure management while accelerating digital transformation.

Drivers

- **Boosting business agility and innovation:** IPE transforms infrastructure from a bottleneck into a high-velocity delivery engine. By providing “Golden Paths,” it enables teams to pivot and deploy features in minutes, drastically shortening the feedback loop between ideation and customer value.

- **Enabling AI & ML Workloads:** Deploying AI requires more than just a model; it requires a governed environment. IPE delivers the platforms that act as the control plane for AI, providing the specialized storage (vector DBs) and compute (GPU clusters) needed for modern inference and AI workloads.
- **Modernizing infrastructure platform delivery:** I&O must shift from reactive ticket-takers to strategic “Product Managers.” By treating the platform as a product, teams provide self-service primitives that empower developers while maintaining central oversight and long-term architectural integrity.
- **Improving developer experience and productivity:** Abstracting infrastructure complexity eliminates cognitive load on developers. Providing seamless, self-service access to environments and tools maximizes productivity and allows developers to focus on feature development.
- **Strengthening compliance and security:** Integrating automated security and compliance controls directly into infrastructure delivery pipelines, in collaboration with security and risk teams, strengthens the organization’s security posture.
- **Optimizing costs and efficiency:** IPE drives scalability, reliability and security, improving operational efficiency and reducing resource costs associated with manual work and change-related downtime. Platform standardization further optimizes resource utilization and eliminates redundant tooling expenses.
- **Scaling operations effectively:** As organizations grow, managing infrastructure becomes increasingly complex. Infrastructure platforms provide a scalable solution, enabling I&O teams to support growth without a proportional increase in headcount or manual effort.

Obstacles

- **Specialized skills gaps:** Beyond basic automation, platforms require a product mindset and deep expertise in cloud native architecture. Finding talent that balances software engineering with operations and security remains a significant barrier.
- **Agentic AI adoption:** As platforms adopt agentic AI for automated and autonomous provisioning, mature IPE is required to enforce the necessary guardrails.
- **Domain I&O silos:** Traditional I&O structures are divided by networking, storage and compute that hinder the delivery of unified, end-to-end services. IPE requires a shift to cross-functional “stream-aligned” teams that many legacy orgs resist.

- **Demonstrating value:** Demonstrating the return on investment from IPE can be challenging in the early stages.
- **Tooling complexity:** The landscape of infrastructure automation tools is vast and complex. Choosing the right tools, integrating them effectively and avoiding tool sprawl can be difficult.

User Recommendations

- **Start small and evolve:** Define clear goals and objectives of the platform by understanding common user demand and delivering viable products that continuously evolve to meet those needs.
- **Assemble a product-centric team:** IPE requires a fusion of software engineering and infrastructure expertise. Appoint a product owner to align technical roadmaps with developer demand. Prioritize hiring talent with strong interpersonal skills and an automation-first mindset, while fostering continuous learning to upskill existing team members for AI-native operations..
- **Establish agentic guardrails & governance:** As you integrate AI agents, implement Policy-as-Code to act as a real-time safety net, ensuring AI-generated changes are automatically validated for cost, security, and compliance before execution.
- **Measure and demonstrate value:** Track key platform effectiveness metrics (deployment frequency, cost, adoption and user satisfaction) and regularly report performance/value to stakeholders.

Gartner Recommended Reading

[Adopt Platform Engineering to Improve the Developer Experience](#)

[How to Implement Cloud Platform Operations](#)

[Video: Demystifying Infrastructure Platform Engineering](#)

[How to Develop Infrastructure Platform Engineering Teams](#)

Site Reliability Engineering

Analysis By: George Spafford, Daniel Betts

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Site reliability engineering (SRE) is a collection of systems and software engineering principles used to design and operate scalable, resilient systems. Site reliability engineers collaborate with customers or product owners, using operational data and feedback, to define service-level indicators (SLIs) and service-level objectives (SLOs). Site reliability engineers work with product or platform teams to design, operate and continuously optimize systems that meet defined SLOs.

Why This Is Important

SRE is vital because it shifts reliability from an “IT problem” to a business strategy. By using reliability-focused metrics, it provides a data-driven framework to balance fast feature delivery with system stability. SRE drives efficiency by aggressively automating manual “toil,” allowing teams to scale without linear headcount growth. Ultimately, SRE ensures that reliability is a shared responsibility, protecting the customer experience while enabling sustainable innovation.

Business Impact

SRE transforms reliability into a growth driver. By focusing on critical customer journeys, it ensures that IT performance directly protects revenue and brand loyalty. The real impact is the ability to scale without proportional cost increases while accelerating time to market through safer, automated deployments. This shifts I&O from a “cost center” to a competitive advantage that manages risk at the speed of the business.

Drivers

- Organizations are under pressure to meet customer requirements for reliability while scaling their digital services.
- SRE has evolved since its initial implementations. Reliability remains the focus and integrates a rich body of knowledge that complements agile, DevOps and platform engineering approaches.
- Organizations that have adopted highly skilled automation practices (usually DevOps) and usage of infrastructure-as-code capabilities to deliver digital business products expand to incorporate reliability into these products.
- The most common use case, judging from inquiry calls with Gartner clients, is to leverage SRE concepts to improve the reliability of existing systems that are not meeting customer requirements for availability or performance, or proving difficult to scale.
- As more autonomous agents are deployed, SRE provides the guardrails (SLOs) for nondeterministic AI behavior.
- Moving from automation to autonomous operations uses agents to detect, diagnose and self-heal outages.
- SRE is required to manage unpredictable dependencies created by multiagent architectures.

Obstacles

- Belief that reliability is an I&O responsibility only, rather than being a shared organizational goal.
- Organizations have difficulty shifting mindsets from service-level agreements with blanket availability statements to defining and measuring reliability in terms of customer-focused SLOs.
- Defining appropriate SLOs that reflect both business needs and user expectations can be complex.
- Finding SRE role candidates with the right mix of development, operations and people skills is a challenge for clients. This impacts initial adoption and scaling efforts.
- Clients have voiced problems with product owners who overly focus on functional requirements, thus hampering improvements.
- SRE requires close collaboration between development and operations teams. Organizational silos can make it difficult to implement SRE effectively.
- SRE initiatives require executive sponsorship and investment to effectively scale.

User Recommendations

- Define your goals for reliability that can guide implementation of SRE.
- Begin your SRE journey with a small, focused initiative and then iteratively grow adoption.
- Select opportunities that are politically friendly, will demonstrate sufficient value and have an acceptable risk profile.
- Automate for SRE efficiency.
- Work with the product owner to identify the SLIs and SLOs necessary to make customers and users happy.
- Implement and improve observability to objectively report on performance relative to error budgets and SLOs.
- Ensure product owners treat SLOs as a feature and are accountable for them.
- Instill collaboration among site reliability engineers, developers and other stakeholders.
- Create a community and leverage organizational learning practices while evolving SRE practices.
- Cultivate a blameless culture where mistakes are seen as opportunities for learning and improvement.
- Evaluate and invest in AI SRE tooling to lower cost of SRE adoption, meet operational demands and deliver effective reliability.
- Establish “agentic guardrails” by defining SLOs specifically for AI agents to ensure autonomous actions don’t breach safety or cost boundaries.

Gartner Recommended Reading

[7 Steps to Start and Evolve Your SRE Practice](#)

[Improve Customer Experience and Reliability With SLIs and SLOs](#)

[Organize for Resilience With an Optimal SRE Team Topology](#)

[Market Guide for AI Site Reliability Engineering Tooling](#)

AI, Autonomy, and Architects: The Future of Site Reliability Engineering

Swarming Support

Analysis By: Siddharth Shetty, Chris Matchett

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

Swarming support is a support model that bypasses a traditional hierarchical escalation structure by using collaboration capabilities to engage directly with individuals from different support teams to diagnose and resolve an incident. A swarm is usually an ad hoc, agile arrangement formed by those with specific skills.

Why This Is Important

IT operations teams use swarming support techniques to enable direct engagement with support technicians via collaboration channels to address higher-priority incidents in real time. Major incident managers and problem managers use similar outreach techniques to form temporary support teams for certain high-profile issues. Swarming is a key aspect of collaborative support hubs.

Business Impact

Swarming support:

- Enables more rapid resolution by directly accessing skilled, available support resource
- Allows for increased IT productivity and reduced negative business impact, as higher-priority incidents are resolved more quickly
- Helps facilitate collaboration between the IT service desk and expert groups
- Helps create knowledge articles more easily, as the agent is part of the dialogue
- Improves efficiency, as fewer tickets are sequentially passed between teams

Drivers

- Swarming support is used to respond to high-priority incidents when major incident managers form agile response teams.
- As agentic AI matures, service desk agents will need to focus on more complex and higher-value interactions. Swarming support provides a mechanism to coordinate these tasks across different teams.
- Swarming occurs organically in peer IT support.
- The emergence of AI capabilities, including intelligent swarming, and the integration with event intelligence may make swarming automation more accessible.

Obstacles

- Swarming support is perceived as only suitable for higher-priority incidents and unable to scale to support all types of incidents.
- Scaling responses for swarms of more than a few participants requires investment in tools or platforms that include collaboration capabilities.
- Ad hoc swarming doesn't guarantee that the right expertise is brought to bear. If experts are already busy, swarms may go unanswered, resulting in inconsistent support experiences.
- An open and collaborative approach across the support environment is required to make swarming work.
- Swarming may be counterculture in some organizations where technology-specific support is leveraged.
- Inappropriate swarming requests can be intrusive and time-intensive, and can distract experts from other important work.

User Recommendations

- Catalog the skills and capabilities of participating IT staff, including knowledge and abilities that are not related to their formal job role.
- For 24/7 operations, document IT service desk staff shifts and on-call schedule of experts.
- Provide clear prioritization guidance to swarming participants and allocate time in their routines to contribute without disrupting other work and initiatives.
- Pilot swarming with a select set of applications and associated tickets to establish their viability and surface improvement opportunities.
- Use the pilot to agree with constraints, as well as the type of incidents in scope.
- Measure the value of swarming in the improved mean time to resolve.
- Reward participation in swarming outreaches, as these may fall outside the scope of standard reporting.
- Establish predefined swarm teams for highly critical systems and frequently recurring, business-impacting incidents.

Gartner Recommended Reading

[How to Integrate ITSM Teams With Product Teams](#)

[Quick Answer: How Can I Reduce the Impact of Major Incidents?](#)

[Innovation Insight: Collaborative Support Hubs](#)

[Essential Roles and Competencies for a Collaborative Support Hub](#)

DEX Tools

Analysis By: Dan Wilson, Stuart Downes, Lina Al Dana, Robin Milton-Schonemann

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Digital employee experience (DEX) management tools measure and help continuously improve the performance and employee sentiment related to company-provided technology. Near-real-time processing of data from endpoints, applications, employee sentiment and organizational context surfaces actionable insights and drives self-healing automation, optimized support and employee engagement. Insights and self-healing also enhance interactions with self-service portals and chatbots.

Why This Is Important

DEX is critical because, if poorly managed, it directly undermines the ROI of AI and digital workplace initiatives, creating a hidden tax on employee productivity and threatening talent retention. Client interest in DEX persists, but primary use cases remain IT operations focused. More mature organizations strategically invest in experience-focused use cases and a cross-functional DEX management office (DEXMO) to scale and help the company retain and attract top talent.

Business Impact

DEX tools shift IT's focus from technology management to experience enablement and more business value-added work. Specific impacts include:

- Reduced employee digital friction
- Improved IT visibility and efficiency
- Quicker IT support resolutions
- Increased reliability through self-healing
- Improved endpoint compliance
- Better balance of objective and subjective success measures
- Optimized device costs by embracing performance-based refresh cycles
- More proactive, predictive and human-centric IT

Drivers

- A DEX strategy and DEX tools play a significant role in shaping the overall employee experience as the workforce relies heavily on technology to carry out job responsibilities.
- Poor DEX is limiting the ability to achieve returns from AI and other technology investments as unstable or frustrating DEX slows adoption, productivity and trust in new tools.
- When employees silently endure or find workarounds for issues themselves, IT is unaware and unable to resolve them, leading to negative sentiment for IT.
- IT leaders demand broader measurement to supplement internally focused activity KPIs that have proven incomplete.
- IT administrators are looking for better visibility into how devices and applications are performing.
- Employee sentiment toward technology cannot be measured effectively with periodic or transactional surveys alone. Feedback must also include how employees feel about and engage with specific devices or apps, and how technology changes impact their work.
- The need for faster support interactions requires service desk and other IT support analysts to have faster access to device configuration and performance data.
- Increasing threat of cyberattacks demands faster identification and remediation of configuration issues and missing patches.
- Organizations are looking to reduce costs and e-waste by replacing scheduled device life cycle strategies with performance-based strategies that increase lifespans and maximize ROI.
- Increased importance of IT sustainability goals is driving demand for better visibility into endpoint power consumption and greenhouse gas emissions, automation to adjust power settings, and ways to communicate directly with employees to encourage sustainable behaviors.
- AI has significantly increased the perceived value and capability of SaaS-based DEX tools.
- Organizations seeking to improve employee retention require continuous, experience-level monitoring to identify digital friction before it escalates.

Obstacles

- IT administrators distrust AI-powered insights and recommendations and are sometimes threatened by automation.
- An “ignorance is bliss” mindset, fearing that a sudden unveiling of chronic issues will make IT look bad.
- Like AI, DEX offers a strong return on employee (ROE) but ROI remains difficult to objectively measure.
- Cost to acquire, implement and integrate new tools is often unbudgeted or perceived as high.
- Insufficient staffing levels and training, no clear owner for or skills to operate a DEX tool.
- Limited organizational change management hinders process and IT staff’s behavioral changes.
- Many organizations overpay for market-leading tools when a less expensive (and mature) solution may meet near-term needs.
- Some struggle to integrate with a fragmented landscape of IT service management, monitoring and asset management systems, which can reduce the impact of insights and successful self-healing automation.

User Recommendations

- Develop EX-focused requirements, use cases and success metrics by collaborating with business and IT peers.
- Ensure the business case focuses on objective and measurable impacts by minimizing reliance on vendor-provided ROI templates.
- Choose a DEX tool that best fits needs and budget by reviewing the [Magic Quadrant for Digital Employee Experience Management Tools](#) and [Critical Capabilities for Digital Employee Experience Management Tools](#).
- Assign dedicated leadership and resources to focus on DEX by reallocating staff as proactive automation reduces support volumes.
- Transform the IT operations culture to incorporate proactive support and incentivize new behaviors by adapting IT performance measures to focus more on outcomes than activities.
- Avoid diminishing returns by adding features and use cases as the team and DEX tool mature.
- Build a business case for and implement a DEX management office to scale DEX efforts.

Sample Vendors

ControlUp; HP; Ivanti; Lakeside Software; Nexthink; Omnissa; Riverbed; ServiceNow; Tanium; TeamViewer

Gartner Recommended Reading

[How to Successfully Deploy a DEX Tool](#)

[Magic Quadrant for Digital Employee Experience Management Tools](#)

[Critical Capabilities for Digital Employee Experience Management Tools](#)

Digital Employee Experience

Analysis By: Tori Paulman, Dan Wilson

Benefit Rating: Transformational

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Digital employee experience (DEX) refers to the employees' perceptions of, and interactions with, the technology they use for work, including applications, devices, services and, increasingly, AI. A DEX strategy enables IT to empower employees to adopt new ways of working by minimizing technology friction, supporting meaningful adoption, applying usage and performance metrics and driving continuous improvement.

Why This Is Important

Most workers are spending more time using technology to get their work done. Additionally, organizations continue to deploy AI capabilities, making the digital employee experience a baseline component to employee willingness to adopt technology to boost productivity and satisfaction. A strategy for continuously understanding, measuring and improving DEX is essential to attracting and retaining talent, improving employee engagement, reducing IT overhead and maximizing workforce productivity.

Business Impact

DEX strategies leverage best practices like personas, journey mapping, and observability to minimize digital friction. A coordinated, holistic DEX approach across IT and non-IT partners reduces silos and maximizes benefits, including improved workforce digital dexterity, well-being and talent attraction/retention. Proactive, DEX-leading IT teams demonstrate higher employee satisfaction and worker productivity than reactive, supporting IT teams.

Drivers

- IT leaders must understand how workers and leaders perceive and value their relationship with IT and improve partnerships.
- The proliferation of applications, devices and services has resulted in a workforce frequently toggling between applications. DEX strategies must reduce friction and improve the productivity of workers by helping them spend more of their time doing their core work functions.
- There is a wealth of data about employees' workstyles, productivity and friction. A DEX listening strategy turns the noise of telemetry and sentiment into insights, or call workstyle analytics, while protecting employee privacy.
- Personas, journey mapping, experience design and effective listening can ensure technology investments deliver expected benefits.
- Business leaders are increasingly looking for guidance on how technology can help boost employee productivity, engagement, well-being and digital dexterity.
- IT leaders are realizing that poor DEX also inhibits the organization's workforce AI ambitions.
- IT leaders are increasingly investing in DEX tools that collect, combine and derive actionable insights from employee feedback with measurements of technology performance, stability and use. These tools generate actionable insights and drive self-healing automations and optimized support.
- A focus on DEX often leads to reduced operational costs as large systemic issues are removed from the environment, enabling IT to shift focus from resolving technical issues to higher-value employee enablement initiatives.

Obstacles

- DEX strategies intentionally shift IT's focus from technology to humans, which is difficult for tech-centric IT staff and leaders.
- Focusing too much on DEX tools as an all-in-one fix often leads to suboptimal outcomes where stakeholders cannot see the value of your DEX investments.
- Proving a return on investment for investing in DEX tools requires rethinking traditional IT metrics.
- Trying to improve DEX is nearly impossible without HR partner buy-in to connect overall employee experience to DEX strategy.
- Software, services and staffing costs to acquire, implement and integrate DEX tools often require collaborative funding from both IT, HR and line of business.
- Many workers still see IT as a transactional break/fix service provider, so DEX leaders must more effectively communicate the value they are providing.
- Low-maturity IT organizations often discover their DEX ambitions are impossible if they are still working toward consistent backlog management and reliable service delivery.

User Recommendations

- Appoint a DEX leader to establish a charter and vision that connects to mission-critical priorities.
- Strengthen partnerships with IT, HR, communications, facilities and business peers by connecting digital and employee value proposition and developing broader success metrics.
- Formalize and structure crossfunctional efforts with a DEX management office.
- Prioritize the highest-value activities by developing a DEX listening strategy that blends quantitative and qualitative data collected from various sources.
- Prioritize and build high-impact digital personas and journey maps that identify key moments in the employee journey.
- Relentlessly pursue digital friction by looking for activities that require excessive effort, result in pain points or wait times and lead to worker questions or errors.
- Mine insights from interactions with self-healing, self-service portals and chatbots for opportunities to improve DEX at scale.

Gartner Recommended Reading

[Build a Case for a DEX Management Office to Overcome Digital Friction](#)

[Establish a DEX Management Office to Overcome Digital Friction](#)

[Nationwide Thrives With a Data-Driven Level 4 Digital Workplace](#)

[Tool: Digital Employee Experience Journey Map](#)

Generative AI

Analysis By: Svetlana Sicular

Benefit Rating: Transformational

Market Penetration: More than 50% of target audience

Maturity: Early mainstream

Definition:

Generative AI (GenAI) technologies can generate new derived versions of content, strategies, designs and methods by learning from large repositories of original source content. Generative AI has profound business impacts, including on content discovery, creation, authenticity and regulations; automation of human work; and customer and employee experiences.

Why This Is Important

GenAI adoption has leapt in the enterprise, even where IT budgets are shrinking. Low-maturity firms focus on reducing costs and improving productivity and high-maturity firms pursue enterprise evolution and transformation. Most application vendors have already incorporated GenAI features. LLM providers keep shipping large-step upgrades rather than small tweaks for reasoning, inference, context, coding assistance and agentic capabilities.

Business Impact

Executive leadership is the top advocate for GenAI adoption. \$10 billion-plus revenue firms, automotive, life sciences, oil and gas, consumer goods and transportation are leading in GenAI adoption and receiving significant value. Organizations deploy frontier models for broad capabilities and cost-effective, smaller models for specialized business. This year, the enterprise focus is on LLM-enabled automation of complex tasks and software engineering productivity. Governance is the top challenge.

Drivers

- Agentic AI is a top driver of a GenAI value proposition this year due to automation benefits and combining GenAI with other techniques.
- Copilot and ChatGPT Enterprise are the top GenAI productivity assistants of choice for the enterprise.
- GenAI drives the wave of AI-enabled applications that support business functions from customer service and regulatory compliance to content generation and decision support. This helps enterprises achieve greater responsiveness in language-based tasks.
- As per the 2025 Gartner Generative and Agentic AI in Enterprise Applications Survey, 54% of firms require demonstrable ROI to justify broad GenAI investment. Improvement in business outcomes is the most common measure of GenAI success (64%). Just over one in two enterprises get significant value from their internally customized LLM.
- Organizations that apply a formal, centralized strategy consistently to all enterprise applications get the most value from their LLM implementations.
- IT departments have become a top use case for GenAI. They leverage GenAI for code generation, cybersecurity testing, help desk support, and to accelerate legacy modernization.
- Governments worldwide, spurred by the GenAI promise, compete for computational dominance and military applications. Divergence between U.S., EU, and Chinese regulatory approaches is increasing.
- Intense GenAI vendor competition continues to push boundaries, fueled by the shift toward autonomous agents that can execute complex, multihour, multistep workflows. Frontier models are evolving from simple chatbots into sophisticated thinking systems that prioritize deep reasoning and integrated audit trails. To manage extreme compute costs, architectural optimizations like mixture of experts and post-training paradigms such as reinforcement learning with verifiable rewards have accelerated progress in reasoning. Inference scaling techniques now improve accuracy through parallel sampling and self-refinement without altering underlying model weights.

Obstacles

GenAI is in the trough due to the mounting obstacles for the enterprises:

- Top challenges in GenAI adoption are security and governance, lack of AI literacy and difficulty integrating into workflows.
- Only 22% of organizations report significant value from GenAI tools. They often see productivity leakage.
- GenAI licensing and pricing remain confusing and constantly evolving, often catching customers by surprise. Escalating compute and memory costs driven by increases in context length and inference scaling threaten the financial viability of model deployments at scale.
- Enterprise adoption lags frontier technical capability, constrained by security, legal compliance, change management and the need to rewrite legacy systems.
- Governance challenges include trustworthiness, accountability and establishing a governance operating model. Uneven regulations across various jurisdictions may impede GenAI initiatives.

User Recommendations

- **Align IT, business and leadership for GenAI adoption:** Jointly reach an understanding about what problems GenAI can solve for tangible benefits in innovation and operational excellence. Prioritize innovation (trial and error).
- **Adopt an agile strategy:** Regularly update your approach to integrate relevant evolving technological, operational and governance perspectives for long-term growth.
- **Design for model flexibility:** Build loosely coupled solutions to enable easy model switching.
- **Execute an AI-ready data strategy:** Take advantage of your unique proprietary data. Ensure data is well-understood, relevant and accessible for model consumption.
- **Invest in talent:** Prioritize AI literacy and upskilling to ensure the workforce can effectively utilize GenAI technologies.
- **Emphasize governance:** Establish an operating model with clear decision rights, policies and technical oversight.
- **Budget for total cost:** Plan for infrastructure, compute and ongoing maintenance throughout the GenAI life cycle.

Sample Vendors

Alibaba Cloud; Anthropic; Cohere; DeepSeek; Google; Meta; Mistral AI; OpenAI

Gartner Recommended Reading

[Forecast: Generative AI Models, Worldwide, 2024-2030, 1Q26](#)

[Emerging Tech: AI Vendor Race: Adoption Trends for Generative AI](#)

[Emerging Tech: AI Vendor Race: Where Generative AI Works, Will Soon Work, and May Never Work](#)

[Emerging Tech: Open Generative AI Models Threaten and Reshape Vendor Product Strategies](#)

Event Intelligence Solutions

Analysis By: Martin Caren

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Event intelligence solutions (EISs) apply AI and data analytics to system and service events to augment, accelerate, and automate manual efforts in the event management process. EISs are defined by the key characteristics of cross-domain event ingestion, topology assembly, event correlation and reduction, pattern recognition, and remediation augmentation. They help reduce event noise and improve productivity by processing event streams into actionable insights and enabling proactive responses.

Why This Is Important

Rising application complexity and event volumes from multiple tools and platforms overwhelms manual processes. EISs use AI to classify and cluster cross-domain events in near real time, delivering insights that speed up incident response, reduce risk, and help ensure business continuity. EISs help keep IT operations running more efficiently in large, heterogeneous IT environments.

Business Impact

EISs deliver value through:

- **Improved agility and productivity:** Operators focus on critical issues, resolve incidents faster, and boost efficiency.
- **Enhanced service availability and reduced triage costs:** Faster root cause identification and automated remediation minimize downtime and lower operational expenses.
- **Maximized return on monitoring investments:** Unified events across tools enable informed decisions, rapid resolution, and better team collaboration.

Drivers

Demand for EIS capabilities is accelerating and fueled by:

- **Increasing complexity:** Organizations rely on a growing portfolio of monitoring and observability solutions to ensure the reliability and resilience of increasingly complex and distributed, hybrid and multicloud workloads.
- **Increasing monitoring expectations:** Investments and improvements in monitoring, and the pursuit of observability, generate more data from more sources. This increase in data presents operators with extremely detailed views into their applications, business services, and the end-user experience. Effective use of this additional data requires near-real-time analysis and correlation of events from related assets and services.
- **Demand for reliability:** Shifts in roles and responsibilities are driven by modern operating models, like DevOps and site reliability engineering (SRE), in the pursuit of greater availability and faster incident resolution. EISs enable agility by offloading some mechanical tasks of event triage, root cause analysis, and solution identification, accelerating response to common issues and freeing up human creative capacity for novel events and business priorities.

Obstacles

- **Unrealistic expectations and ROI:** Hype makes it hard for clients to separate AI claims from reality, making ROI difficult to demonstrate. EISs are often seen as operator tools, not as strategic assets, so they are vulnerable during budget cuts.
- **Maturity and complexity:** EIS success typically requires high maturity in dependencies like change management and automation. Without this, it adds a layer of complexity that can outweigh the benefits. Furthermore, fragmented data access remains a persistent bottleneck.
- **Market overlap and AI SRE:** EISs overlap with observability and IT service management (ITSM), creating a crowded and sometimes confusing market. Newer, AI-driven SRE models may also leapfrog traditional EISs, which must evolve beyond correlation to avoid becoming a legacy utility.

User Recommendations

- Layer the cross-domain analysis of an EIS with maturing your monitoring and observability strategy. This approach creates a solid foundation of valuable data for ingestion and analysis, and the surfacing of insights across domains.
- Do not focus solely on automated remediation; this is rarely achieved at scale. Accelerating response and augmenting human decision making have tremendous value. These approaches often avoid the challenge of the probabilistic uncertainty, combined with automated change in production environments.
- Explore EISs that make use of agentic AI for enhanced diagnostic capabilities and the possibility of accelerated remediation and root cause analysis (RCA).

Sample Vendors

BigPanda; BMC Software; Cisco; GreySkies; IBM; Interlink Software; PagerDuty; ServiceNow; Virtana; Vitria

Gartner Recommended Reading

[Market Guide for Event Intelligence Solutions](#)

[Observability Tool Consolidation Must Begin With Event Intelligence](#)

[Choosing the Right Observability Architecture: Blueprints and Decision Criteria](#)

Climbing the Slope

Technology Change Automation

Analysis By: Chris Laske

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Technology change automation streamlines change management by automating activities, such as change ticket creation, risk assessment, conflict detection and approvals. Integrations between IT service management (ITSM) platforms, continuous integration/continuous delivery (CI/CD) tools and event intelligence platforms enable rapid delivery while ensuring change visibility and minimizing business disruption. Organizations also use artificial intelligence to further mitigate change risk.

Why This Is Important

Growing demand for rapid deployments and low tolerance for change-related outages make manual change approval, execution and follow-up process insufficient. In response, infrastructure and operations (I&O) teams use technology change automation to autogenerate change records, assess change risk, identify change conflicts, streamline approvals and improve change visibility, resulting in higher change velocity and quality while minimizing business disruption.

Business Impact

Technology change automation delivers value to I&O and agile organizations by:

- Integrating ITSM platforms with CI/CD pipeline tools to streamline change record creation.
- Using AI-driven risk analysis to improve change classification and decision accuracy.
- Automating change approval workflows to reduce delays and accelerate execution.
- Running automated conflict detection to prevent change collisions.

- Feeding change record data into event intelligence platforms to quickly correlate incidents with production changes.

Drivers

- Rising digital business demands and widespread adoption of agile development require faster and more frequent deployments.
- Manual change practices are labor-intensive, placing significant strain on internal resources.
- Higher change volumes increase the risk of conflicts and collisions, undermining stability.
- Inaccurate change risk assessment leads to business disruption and unplanned outages.
- Limited visibility into production changes delays service restoration, especially during major incidents.
- Audit and compliance requirements continue to mandate change records, even as deployment speeds increase.
- ITSM vendors increasingly extend AI applications beyond IT support, using them to automate change risk analysis and evaluate historical change data to identify standardization opportunities.

Obstacles

- Independent change management practices weaken outcomes when development teams underestimate the need for an integrated approach. Without unified change management, organizations increase the risk of conflicting changes, service disruptions and reduced operational resilience.
- The absence of foundational change management practices, such as clearly defined standards for routine changes, makes it difficult to implement automated approval workflows and effectively leverage AI.
- Integrating multiple release and deployment tools with the ITSM platform often remains challenging, limiting end-to-end visibility and coordination across change activities.

- Organizations that fail to track business value and change velocity metrics, including value enablement and the cost of disruptions caused by failed changes, struggle to justify the investment required to enable change automation.

User Recommendations

- Track and report business value and change velocity metrics to quantify outcomes such as faster value enablement and delivery, as well as the cost of business disruption from failed changes.
- Define clear request-for-change review criteria to convert policies into workflow rules by identifying recurring decision patterns and embedding them directly into change evaluation processes.
- Implement automated approval workflows or use AI agents for standard changes, incorporating dynamic routing to increase autonomy and reduce bottlenecks. Leverage ITSM capabilities such as conflict detection, preapproved change window validation and CI/CD tool integration.
- Integrate ITSM and event intelligence platforms to improve correlation between incidents and production changes using shared change management data.
- Adopt intelligent risk advisory capabilities that apply predictive analytics from prior releases to assess change risk and business impact.

Infrastructure Automation

Analysis By: Chris Saunderson

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Definition:

Infrastructure automation enables DevOps, platform, and infrastructure and operations (I&O) teams to deliver automated infrastructure services across on-premises and cloud environments. This includes the definition of environments, life cycle of services through creation, configuration, operation and retirement. These services are then made available through automation workflows, specialized platforms, self-service catalogs, direct invocation, agentic AI and API integrations.

Why This Is Important

Infrastructure automation delivers velocity, quality, and efficiency through scalable, declarative approaches for deploying and managing infrastructure. These tools integrate into delivery pipelines and platforms targeting deployments that range from on-premises, hybrid and to the cloud, enabling infrastructure consumers to build what is needed when needed. Once deployed, IA provides day-two-and-beyond operational automation, and extends to policy compliance and policy enforcement capabilities.

Business Impact

IA services enable:

- **Agility, scalability and collaboration:** Fostering team autonomy, delivering what product teams need with security, scalability, cost and compliance requirements
- **Productivity:** Version-controlled, declarative, repeatable, and efficient deployments
- **Cost improvement:** Reductions in manual efforts through increased automation
- **Risk mitigation:** Compliance driven by standardized configurations and governance
- **Quality assurance:** Removal of human error with automated repeatable deployments

Drivers

IA tools support maturation beyond simple deployments through:

- Hybrid infrastructure delivery that is vendor- and platform-neutral
- Support for immutable and programmable infrastructures
- Predictable delivery, enabling automated operations
- Self-service and on-demand environment creation
- Integration into DevOps, platform and infrastructure platform initiatives
- Resource provisioning, including scale-up and scale-down cost optimization capabilities
- Operational configuration management efficiencies

- Policy-based delivery and assessment or enforcement of deployments against internal and external policy requirements
- Enterprise-level framework to enable maturing of automation strategies

Obstacles

- Combining the tools needed to deliver IA capability can increase tool count, cost and complexity.
- Software engineering skills and practices are required to get maximum value from tool investments.
- IA vendor capability expansion overlaps and confuses the tool landscape, resulting in overinvestment.
- Steep learning curves can cause developers and administrators to revert from familiar imperative methods using scripts for infrastructure-as-code-based declarative approaches.
- Demonstrating the return on investment (ROI) of automation projects can be challenging, especially in the short term.

User Recommendations

- Identify existing IA tools used to catalog capabilities. Recognize use cases and document overlaps to aid decision making.
- Assess existing internal IT skills to incorporate training needs that more fully enable IA, especially for an automation architect role to coordinate standards development and implementation. Evaluate the use of GenAI as a means to upskill teams.
- Baseline how managed systems and tooling will be consumed (e.g., by engineers, self-service catalog, API, agentic AI or on-demand).
- Integrate security and compliance requirements into the scope for automation and delivery activities.
- Develop an IA tooling strategy that incorporates current needs and evolution of the near-term roadmap.

Sample Vendors

Amazon Web Services; Broadcom (VMware); env zero; Gruntwork; IBM (HashiCorp); Microsoft; Perforce Software; Progress Software; Pulumi; RackN; Scalr; Spacelift; Upbound

Gartner Recommended Reading

[Market Guide for Infrastructure Automation and Orchestration Tools](#)

[How to Start and Scale Your Platform Engineering Team](#)

Automated Incident Response

Analysis By: Tanmay Bisht

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Automated incident response centralizes the routing of events or incidents through a policy or rule-based engine, on-call scheduler and streamlined collaboration. AIR improves operational efficiencies with action-oriented insights, shorter incident durations and automated workflows for event routing, seamless collaboration, remediation and escalations.

Why This Is Important

Manual processes for incident resolution are a challenge, especially when multiple experts need to be involved, time is of the essence and the organization wants to improve efficiency. For DevOps teams, juggling contact lists and lack of seamless collaboration inhibit speedy delivery and stability of application features. Automated incident response (AIR) solutions solve this by proactively automating tasks handled manually, streamlining handoffs between support personnel, and reducing cognitive load during high-severity incidents.

Business Impact

- Minimize customer impact and maintain service continuity through quick incident resolution.
- Improve process and operational efficiency insights into incidents and their responses.
- Driving a collaborative approach for I&O and DevOps by integrating with multiple tooling categories.
- Removes fatigue-driven mistakes and compresses resolution time through automated workflows.
- Efficient dissemination of near-real-time information across the organization during an incident.

Drivers

- **Streamlined and automated notification:** AIR automates communication across distributed teams, ensuring quick and accurate responder engagement. It enables organizations to address challenges of geographically distributed teams, remote workforce, complex on-call schedules and notification channel preferences, especially when dealing with incident triage when time is of essence.
- **Efficient data contextualization:** AIR tools enrich events, support remediation through run book automation and collect background context thereby saving time and enriching data that responders require for troubleshooting and incident resolution.
- **DevOps and site reliability engineer (SRE) requirements:** Traditional incident management models cannot meet the needs of agile cultures because of manual tasks in the incident response workflow. AIR caters to the need for seamless collaboration across various groups, enabling DevOps to underpin its offerings with an effective, consistent IT service management (ITSM) practice.
- **Transparent review and analysis:** AIR tools capture an incident's progress from identification through resolution, including the handoffs needed between various teams, and the time and action taken at each step of the incident. This provides vital information for postincident review and process review, thus enabling further enhancements.
- **Workflow automation:** AIR tools can automate workflows that are part of processes like converting actionable alerts into incidents, opening a communications channel in instant messengers for collaboration, updating the real-time status on a web-portal and one-click remediation for existing run books.
- **Wider reach:** Use of AI, including generative AI and agentic AI, drives adoption for a wider audience by translating technical jargon into easily understandable language, making it accessible for relevant stakeholders.

Obstacles

- **Overlapping capabilities:** Features of AIR solutions often overlap with observability, ITSM and AI SRE solutions, making it difficult to articulate the value of investing in a stand-alone solution.
- **Data dependency:** Since AIR solutions rely on topological data and service dependencies, ensuring reliable data feeds and discovery processes are in place is critical.
- **Market volatility:** The vendor landscape has been continuously in flux. With incumbents adding AIR features, startups entering and frequent M&A, vendor selection becomes complicated as well as raises replacement risks.
- **Maturity in I&O:** Large enterprises recognize the need for AIR tools, but many SMBs lack the I&O maturity needed to quantify the cost of time lost during manual incident resolution and justify the investment.

User Recommendations

- Evaluate the necessity of discrete AIR capabilities in a stand-alone solution by reviewing similar features in your current observability, EIS and ITSM solutions.
- Integrate AIR solutions with observability and ITSM platforms through bidirectional interfaces, and extend remediation capabilities by leveraging automation across DevOps toolchains.
- Integrate incident workflow processes with ChatOps tools, such as Slack or Microsoft Teams, to improve incident communication and collaboration. Leverage AIR solutions' GenAI-based capabilities for data collection and summarization for postincident review and analysis.
- Pilot AIR tool's agentic AI capabilities, evaluating its ability to make and execute decisions autonomously or semiautonomously, with humans in the loop.

Sample Vendors

All Quiet; AlertOps; Datadog; Thomas Bravo (Everbridge); Grafana Labs; incident.io; OnPage; PagerDuty; Rootly; Xurrent

Gartner Recommended Reading

[Innovation Insight: Increase Availability With Automated Incident Response](#)

Entering the Plateau

Software Asset Management Tools

Analysis By: Yolanda Harris, Jaswant Kalay

Benefit Rating: Moderate

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Software asset management (SAM) tools help maintain compliance with software, SaaS and cloud licensing agreements. They support efforts to optimize the assets, identifying opportunities to reuse, rationalize and downgrade software, monitor its consumption, and provide key data to support software and cloud negotiations. SAM tools aggregate enterprise entitlement and consumption data, reconciling it to establish an effective license position (ELP) and reporting to properly govern software use.

Why This Is Important

SAM tools discover, inventory, normalize and reconcile on-premises, cloud and SaaS software consumption with entitlement data, allowing organizations across any industry and geography to streamline software licensing management. Recognized in Gartner's ITSM Hype Cycle as mature, mainstream solutions, SAM tools provide near-real-time usage data, cost optimization and noncompliance risk reduction. These capabilities often extend to enable security, risk and cloud consumption.

Business Impact

- SAM tools help organizations address the universal challenge of managing software and SaaS licenses, a task made harder by large estates, decentralized purchasing, complex licensing and hybrid environments.
- SAM tools give IT and procurement leaders a unified view of the software estate, supporting decision making, risk mitigation, spend optimization and improved life cycle management.

Drivers

- Growth of cloud computing, SaaS and AI adoption has increased the complexity of SAM, with expanding requirements on cloud cost management, complex licensing rules, new and evolving consumption metrics and business-led technology buying.
- IT sourcing, procurement and vendor management (SPVM) leaders who invest in SAM tools gain visibility into software inventory and usage and improve the enterprise's ability to negotiate with software vendors and proactively manage risks.
- SAM tools help IT finance and budget owners gain insights into software expenditures for improved demand and forecasting, cost allocation and IT financial management.
- SPVM leaders are under pressure to minimize software licensing costs and identify unbudgeted compliance fees associated with software vendor audit risks.
- Enterprise architects and IT security teams look to SAM tools to provide visibility into the entire IT estate. Enriched information, such as known vulnerabilities and exposures, end of life and end of support, becomes available, which enables enterprises to mitigate risks and plan for upgrades.
- SAM tools, within infrastructure and operations, have become valuable because they help consolidate and normalize data from various discovery and inventory sources. This cleansed data is essential for a robust configuration management database and accelerates adoption and value for IT service delivery.
- Management of the SaaS life cycle to reduce sprawl and security risks drives the need for tools that expand traditional SAM capabilities or replace them with SaaS management platforms and SaaS security tools. These platforms are designed to address SaaS-specific challenges, such as visibility, usage optimization, entitlement management and enhanced security, across the SaaS portfolio.

Obstacles

- Most sourced software still does not comply with ISO/IEC 19770-3:2016 for software entitlement tags. Nonstandardized tags hinder integration with SAM tools, require manual license management, and limit adoption of AI-driven SAM tool automation and data sharing with adjacent platforms, such as IT service management and configuration management databases. Standardized entitlement data remains critical for next-generation SAM capabilities.

- Cloud services, hybrid environments and evolving compliance standards change the criteria for SAM tool features and leave gaps in their tool coverage. This often requires manual intervention and SAM managed service providers (MSPs) or additional data collection tools to produce an effective license position.
- Organizations often struggle to identify, differentiate and prove the value of SAM tools because of specific use cases and preexisting software tools that have adjacencies or significant overlapping capabilities with SAM tools.

User Recommendations

- Start with three to five priority vendors for initial SAM management, based on risk, spend and importance. Create key vendor management use cases for the tool, then expand to additional publishers over time.
- Clarify specifics about your environments and use cases — for example, SaaS, remote users, etc. This will help you select appropriate SAM tool(s) for your enterprise.
- Ensure out-of-the-box connectors are compatible, user-friendly and maintainable with existing data sources . Regularly validate data accuracy and suitability for SAM with data owners.
- Evaluate SAM MSPs that can support tool implementation and address its limitations. Augment resources for tool administration and operations.
- For most tools, importing entitlement and contract data is often manual and time-consuming. While some vendors are adding emerging AI features to streamline this, evaluate their effectiveness carefully.
- Remediations of nonoptimized states happen outside the tool. A SAM tool supports this but does not drive cost or risk management projects.

Sample Vendors

Eracent; Flexera; Matrix42; ServiceNow; USU; Xensam

Gartner Recommended Reading

[Advance ROI By Applying 3 Core Software Asset Management Lessons](#)

[Toolkit: How to Start a Software Asset Management Discipline](#)

[Toolkit: Build the Business Case for Software Asset Management to Prove ROI](#)

[Market Guide for Software Asset Management Tools](#)

[Magic Quadrant for Software Asset Management Managed Services](#)

Appendixes

See the previous Hype Cycle: [Hype Cycle for ITSM, 2025](#).

Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 2: Hype Cycle Phases

(Enlarged table in Appendix)

Phase	Definition
<i>Innovation Trigger</i>	A breakthrough, public demonstration, product launch, or other event generates significant media and industry interest.
<i>Peak of Inflated Expectations</i>	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the innovation is pushed to its limits. The only enterprises making money are conference organizers and content publishers.
<i>Trough of Disillusionment</i>	Because the innovation does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
<i>Slope of Enlightenment</i>	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the innovation's applicability, risks, and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
<i>Plateau of Productivity</i>	The real-world benefits of the innovation are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
<i>Years to Mainstream Adoption</i>	The time required for the innovation to reach the Plateau of Productivity.

Source: Gartner

Table 3: Benefit Ratings

Benefit Rating	Definition
<i>Transformational</i>	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
<i>High</i>	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
<i>Moderate</i>	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
<i>Low</i>	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner

Table 4: Maturity Levels
(Enlarged table in Appendix)

Maturity Levels	Status	Products/Vendors
Embryonic	In labs	None
Emerging	Commercialization by vendors Pilots and deployments by industry leaders	First generation High price Much customization
Adolescent	Maturing technology capabilities and process understanding Uptake beyond early adopters	Second generation Less customization
Early mainstream	Proven technology Vendors, technology and adoption rapidly evolving	Third generation More out-of-box methodologies
Mature mainstream	Robust technology Not much evolution in vendors or technology	Several dominant vendors
Legacy	Not appropriate for new developments Cost of migration constrains replacement	Maintenance revenue focus
Obsolete	Rarely used	Used/resale market only

Source: Gartner

Document Revision History

- [Hype Cycle for ITSM, 2025 - 17 June 2025](#)
- [Hype Cycle for ITSM, 2024 - 17 June 2024](#)
- [Hype Cycle for ITSM, 2023 - 14 July 2023](#)
- [Hype Cycle for ITSM, 2022 - 13 July 2022](#)
- [Hype Cycle for ITSM, 2021 - 21 July 2021](#)
- [Hype Cycle for ITSM, 2020 - 27 July 2020](#)
- [Hype Cycle for ITSM, 2019 - 11 July 2019](#)

[Hype Cycle for ITSM, 2018 - 18 July 2018](#)

[Hype Cycle for ITSM, 2017 - 19 July 2017](#)

[Hype Cycle for ITSM 2.0, 2016 - 4 August 2016](#)

[Hype Cycle for ITSM 2.0, 2015 - 17 July 2015](#)

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[Understanding Gartner's Hype Cycles](#)

[Tool: Create Your Own Hype Cycle With Gartner's 2024 Hype Cycle Builder](#)

[Market Guide for IT Service Management Platforms](#)

[AI Use-Case Assessment for IT Service Desk](#)

[Market Guide for Artificial Intelligence Applications in IT Service Management](#)

[A Blueprint for Modernizing Your ITSM Practice to Drive Digital Transformation](#)

[The IT Service Management Leader's First 100 Days](#)

[3 Keys to Enterprise Service Management Success](#)

[Key Functional Considerations to Identify Your Next ITSM Platform](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Priority Matrix for ITSM, 2026

Benefit ↓	Years to Mainstream Adoption			
	Less Than 2 Years ↓	2 to 5 Years ↓	5 to 10 Years ↓	More Than 10 Years ↓
Transformational	Digital Employee Experience Generative AI	Infrastructure Platform Engineering Multiagent Systems Service Operations Site Reliability Engineering	Agent-Centric I&O Threat Exposure Management	
High	Automated Incident Response Autonomous Endpoint Management Technology Change Automation XLA	AI Agents AI-Powered IT Agent Advisory DEX Tools Infrastructure Automation Policy as Code	AI Applications in IT Service Management Enterprise Service Management Technology Dependency Hub	Natural Language Case Extraction
Moderate	Event Intelligence Solutions Software Asset Management Tools	IT Knowledge Generation	Swarming Support	
Low		Operations Assistant		

Source: Gartner (June 2025)

Table 2: Hype Cycle Phases

Phase	Definition
<i>Innovation Trigger</i>	A breakthrough, public demonstration, product launch, or other event generates significant media and industry interest.
<i>Peak of Inflated Expectations</i>	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the innovation is pushed to its limits. The only enterprises making money are conference organizers and content publishers.
<i>Trough of Disillusionment</i>	Because the innovation does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
<i>Slope of Enlightenment</i>	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the innovation's applicability, risks, and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
<i>Plateau of Productivity</i>	The real-world benefits of the innovation are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
<i>Years to Mainstream Adoption</i>	The time required for the innovation to reach the Plateau of Productivity.

Source: Gartner

Table 3: Benefit Ratings

Benefit Rating	Definition
<i>Transformational</i>	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
<i>High</i>	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
<i>Moderate</i>	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
<i>Low</i>	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner

Table 4: Maturity Levels

Maturity Levels	Status	Products/Vendors
<i>Embryonic</i>	In labs	None
<i>Emerging</i>	Commercialization by vendors Pilots and deployments by industry leaders	First generation High price Much customization
<i>Adolescent</i>	Maturing technology capabilities and process understanding Uptake beyond early adopters	Second generation Less customization
<i>Early mainstream</i>	Proven technology Vendors, technology and adoption rapidly evolving	Third generation More out-of-box methodologies
<i>Mature mainstream</i>	Robust technology Not much evolution in vendors or technology	Several dominant vendors
<i>Legacy</i>	Not appropriate for new developments Cost of migration constrains replacement	Maintenance revenue focus
<i>Obsolete</i>	Rarely used	Used/resale market only

Source: Gartner