

Infrastructure Automation: A Strong Starting Point for Agentic AI

With a strong foundation, enterprises can move agentic AI from pilot to production while maintaining control of risk, cost, and accountability.

Author

Arindam Dutta Roy
Dane Anderson





Table of Contents

Title	Page
1. Executive Summary	01
2. Why Infrastructure Is the Right Starting Point	02
3. Determining the Right Level of Autonomy	04
4. Governance Sets the Pace for Scaled Autonomy	08
4.1. Governed Autonomy in Practice	09
5. What Autonomous Operations Mean for Infrastructure Leaders	10
6. Conclusion	11
7. References	13
8. About the Author	14

Executive Summary

Autonomous operations have always served as an aspirational goal that enterprises strive towards. Over the past few years, AI technologies have made this objective more attainable. Today, as enterprises increasingly adopt agentic AI, autonomous operations are widely seen as a reality within immediate reach.

Agentic AI is a key driver behind the shift from traditional automation technologies that follow a fixed script. Its ability to reason, plan sequential actions, and adjust working processes as conditions change allow enterprises to bring intelligence into running operations and connected workflows.

Across industries, enterprises are actively piloting agentic AI programs with the intent to move them into production and scale across different functions. But moving from pilot to production continues to be a key point of contention among leaders striving to extract business value from their AI investments.



Why Infrastructure Is the Right Starting Point

A key reason AI pilots struggle to scale is that most are still seen as a one-time activity, rather than an ongoing capability.

To effectively move to production and achieve scale, AI programs have several requirements they need to satisfy. These include aligning to business outcomes that can be measured, identifying the right stakeholder who can release the necessary budget, owning the outcome beyond the pilot stage, and operating with data that the enterprise already trusts.


These become critical when we see the gap between AI adoption and scale. McKinsey notes that while 88% of organizations report regular AI use in at least one business function, only 23% are scaling an agentic AI system anywhere in the business.¹ The first key decision for enterprise leaders, therefore, is identifying the starting point where agentic AI can move beyond experimentation and demonstrate measurable outcomes.





Infrastructure operations fit these requirements particularly well. Routine activities such as incident management, service desk operations, capacity planning, and cloud cost optimization involve high volumes of operational activity. The outcomes associated with these activities are often tracked through measures such as incident volumes, resolution times, infrastructure utilization, service availability, and cost. The operational data available is also closely monitored in real time and validated for integrity.

Gartner expects agentic AI deployment in IT infrastructure operations to grow from under 5% of enterprises in 2025 to 70% by 2029, moving from a niche experiment to a mainstream operating pattern within a few years.² This positions infrastructure automation as a practical starting point for enterprises looking to build, pilot, and scale their agentic AI programs.



Determining the Right Level of Autonomy

Identifying the starting point is only one part of the decision that determines the success of any AI initiative. Enterprises also need to determine the right level of autonomy that can be granted to their agents, based on the complexity and the risks involved.

With infrastructure automation, the level of agentic autonomy increases progressively when agents prove their ability to scale. During initial stages of the program, agents can begin in an assisted role by identifying issues and recommending actions for human approval. As their performance becomes more reliable and consistent, agents can move into supervised execution, take on defined activities independently, and eventually operate with governed autonomy across wider workflows.

These progressive levels provide infrastructure teams with a practical way to determine how much responsibility can be assigned to an agent at each stage, while retaining the level of human involvement required for the activity.



Progressive levels of agentic autonomy in infrastructure operations




Figure 1: Progressive Levels of Agentic Autonomy in Infrastructure Operations



Gartner projects that human-in-the-loop involvement in IT operations workflows will fall to 40% by 2028, down from 95% in 2025.² This shift reflects the progressive nature of autonomy, where human involvement reduces as agents move from assisted and supervised roles towards delegated and governed autonomy.

As agents achieve higher autonomy levels, infrastructure teams will need to continuously review where agents produce reliable outcomes, where human intervention remains necessary, and whether existing controls can support a higher level of autonomy. Based on this analysis, teams can assess their readiness to move to production with the right level of autonomy.





Identifying Early Wins for Agentic AI

Assessing the effectiveness of an agentic AI program during its nascent stages is crucial for enterprises. One of the ways to achieve this is to identify areas where early wins for agentic AI emerge. For infrastructure automation, these areas involve functions that manage high volumes of repetitive activities.

Service desk management is one area where managing a significant volume of repetitive tickets, such as password resets and access requests, is considered routine activity. Since these activities require limited human intervention, agents can address these requests autonomously and free up support teams to focus more on high-priority requests and address complex incidents.

Cloud cost optimization is another area where early gains can be identified. Agents can continuously monitor infrastructure usage, identify idle resources, right-size instances, and adjust capacity based on changes in demand and utilization. Similarly, incident management and capacity planning also provide opportunities for delivering early wins.

Early wins from agentic AI become easier to measure when applied to infrastructure processes where outcomes are already closely monitored.

As infrastructure teams see consistent improvements across these areas, these early gains provide these teams with a clear indication of where agentic AI can be introduced next. This builds enterprise confidence in the performance of their agents first, before seeking expansion of their role into more complex areas of infrastructure operations.



Governance Sets the Pace for Scaled Autonomy

When agentic AI starts to scale across different infrastructure processes, the level of governance defined around these agents also requires maturation. Any action taken by the agent outside its defined parameters can cause operational and business risks, highlighting the necessity of established and continuous governance.

A robust governance framework clearly defines what an agent can do independently, the actions that require human approval, and the controls that need to remain in place as the level of autonomy increases.

A few mechanisms to consider when defining this framework include:



Decision boundaries define the actions an agent can perform independently and those that require human approval, depending on the impact and level of risk involved.



Audit trails provide operational teams with a record of the actions taken by an agent, including the data it used and its reasoning behind it.



Rollback mechanisms ensure that actions can be reversed quickly whenever the outcome is incorrect or creates an unexpected impact.



Access controls limit agents to the systems and environments required for their role, following the same least privilege principles already used for human credentials.

Ethics and risk management also need to be factored into this governance framework. These ensure agents do not access sensitive information such as customer data, employee details, and financial systems. Implementing automated controls and circuit breakers offers another layer of security and assurance for infrastructure teams, through the ability to pause an agent's action whenever a defined threshold is breached.

Deloitte's 2026 State of AI in the Enterprise survey found that only 21% of organizations currently have a mature governance model for agentic AI in place, even as the use of agents is expected to increase over the next two years.⁴ This highlights the need for enterprises to ensure that the maturity of their governance model continues to grow in tandem with the higher levels of autonomy given to agents.

Governed Autonomy in Practice

- **Shared ownership model:**

A leading Europe-based telecom operator partnered with Tech Mahindra to build and implement a shared ownership model that helped their operations move towards higher levels of agentic autonomy.⁵

- **Tech Mahindra security practice:**

TechM Orion and the VerifAI validation framework helped implement audit logging, access controls, and continuous monitoring into the architecture layer to support automated incident response while retaining oversight.⁶

What Autonomous Operations Mean for Infrastructure Leaders

As agentic AI takes on a larger role across infrastructure operations, the role and responsibility of infrastructure leaders will also begin to evolve. Responsibility for operational issues, ticket volumes, service performance, and infrastructure availability can gradually move to agents as the agents begin to demonstrate higher levels of reliability.

This frees up infrastructure leaders' time to focus more intently on strategic decisions, including defining the right level of autonomy, reviewing agent performance, and determining which decisions should continue to require human intervention.


Leaders will also need to work more closely with risk, compliance, security, and finance teams. This becomes increasingly important as agents operate across a wider range of systems and workflows, making accountability and oversight an integral aspect of infrastructure leadership.



Conclusion

Infrastructure automation offers enterprises an ideal starting point to determine what autonomous decision-making looks like in practice. As agents take on more operational responsibilities, the insights gained from the value they deliver help enterprises define how autonomy should be governed and assess their readiness to scale across different business functions.

For infrastructure leaders, this creates a larger responsibility around how autonomous operations are implemented and managed. The decisions leaders make regarding governance, ownership, and human intervention will help establish the appropriate operating discipline that broader adoption of agentic AI will eventually depend on.




Third Eye Analyst Perspective

The critical test of an AI pilot is whether operational improvement becomes a measurable business benefit. BCG's 2025 global study found that 60% of companies reported minimal revenue and cost gains from AI, while only 5% qualified as "future-built," consistently generating substantial value.⁷ Third Eye Advisory's view is that every scaling decision should include a finance-validated bridge from workflow metrics to business outcomes. Hours saved represent capacity released; their financial value depends on whether that capacity reduces expenditure, increases throughput, or supports additional revenue. Buyers should therefore require a baseline, a named benefit owner, and a net-value calculation that includes integration, AI consumption, human oversight, and ongoing change costs. This makes value realization a funding discipline throughout deployment.

Evidence snapshot — AI maturity and value remain uneven

Limited gains	Scaling and beginning to generate value	Substantial value
60% — Laggards	35% — Scalars	5% — Future-built





References

1. McKinsey & Company. (2025, November 5). The state of AI in 2025: Agents, innovation, and transformation. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
2. Gartner. (2025, December 4). Predicts 2026: AI agents will transform IT infrastructure and operations (C. Haight, A. Banerjee, J. Antelmi, J. Forest, & C. Mukhyala, Authors). Gartner, Inc. <https://www.gartner.com/en/documents/7242030>
3. Forrester Consulting. (2024). The business value of IBM AIOps and observability (commissioned study). In IBM, How can you make the case for observability to your entire organization? <https://www.ibm.com/resources/automate/observability-implementation>
4. Deloitte. (2026). Business and IT leaders report AI agents are scaling faster than their guardrails: The state of AI in the enterprise, 2026 edition. Deloitte Insights. <https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html>
5. Sullivan, P., & Madhok, R. (2026). Maximizing value in large-scale transformation: A co-creation approach [Whitepaper]. Tech Mahindra & Third Eye Advisory.
6. Mehrotra, S., & Kell, N. (2026). Security by design, secure in practice: An end-to-end framework for GenAI [Whitepaper]. Tech Mahindra & Third Eye Advisory.
7. Boston Consulting Group. (2025, September 30). AI leaders outpace laggards with double the revenue growth and 40% more cost savings. <https://www.bcg.com/press/30september2025-ai-leaders-outpace-laggards-revenue-growth-cost-savings>

About the Author

Arindam Dutta Roy is the Global Sales Head for Digital Core Services (DCS) at Tech Mahindra, helping enterprises modernize their core technology environments through cloud, infrastructure, network, hyperscaler, and cybersecurity solutions. Previously, as Chief Growth Officer for CIS, he played a pivotal role in driving growth, service excellence, and customer success across Tech Mahindra's \$1.3 billion Cloud and Infrastructure Services portfolio.

With over 30 years of industry experience, including more than two decades at Tech Mahindra, Arindam has led global growth and large-scale transformation initiatives across multiple industries. He is recognized for building high-performing teams, driving innovation, and leveraging partner ecosystems to create business value through technology modernization.

A member of the Forbes Business Council and a regular industry speaker on Cloud, Infrastructure, AI, and DigitalOps, Arindam also champions talent development and upskilling initiatives aimed at building future-ready technology leaders.



Arindam Dutta Roy
Global Sales Head for Digital
Core Services (DCS),
Tech Mahindra

About the Author

Dane Anderson leads Third Eye Advisory's research into infrastructure, cloud, AI infrastructure, and outsourced infrastructure operations. He spent 21 years at Gartner, most recently as Managing Vice President for Business and Technology Insights for Tech Services Leaders, and before that leading a 12-analyst team covering infrastructure and application outsourcing, cloud services, and product support for hardware and software. Dane works directly with senior executive teams at service provider companies of all sizes on the current and future state of their business, and has guided startup and emerging-provider CEOs to execute rapidly and drive sustainable growth.



Dane Anderson

Infrastructure, Cloud, AI
Infrastructure & Outsourced
Infra Operations Research,
Third Eye Advisory

About **Tech Mahindra**

Tech Mahindra (NSE: TECHM, BSE: 532755) offers AI-powered technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 146,000+ professionals across 90 countries, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies.

For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit <https://www.techmahindra.com/>



www.techmahindra.com

www.twitter.com/tech_mahindra

www.linkedin.com/company/tech-mahindra

Copyright © Tech Mahindra Ltd 2026. All Rights Reserved.

Disclaimer: Brand names, logos, taglines, service marks, tradenames and trademarks used herein remain the property of their respective owners. Any unauthorized use or distribution of this content is strictly prohibited. The information in this document is provided on "as is" basis and Tech Mahindra Ltd. makes no representations or warranties, express or implied, as to the accuracy, completeness or reliability of the information provided in this document. This document is for general informational purposes only and is not intended to be a substitute for detailed research or professional advice and does not constitute an offer, solicitation, or recommendation to buy or sell any product, service or solution. Tech Mahindra Ltd. shall not be responsible for any loss whatsoever sustained by any person or entity by reason of access to, use of or reliance on, this material. Information in this document is subject to change without notice.