


NEXT FRONTIER: AI AGENT-FIRST AUTONOMOUS WORKPLACE MANAGED SERVICES

THE JOURNEY TO AI AGENT-FIRST MANAGED SERVICES



Executive Summary

Digital Workplace investments over the last decade span across cloud productivity, mobility, collaboration, endpoint management, and self-service improved access and connectivity. The next hurdle is operational independence: services that detect friction early, understand user and device context, execute governed actions, and continuously improve employee experience.

For Digital Workplace and Managed Services leaders, the opportunity is clear: shift from reactive support to proactive

resolution, from fragmented automation to reusable automation fabric, and from AI chat experiences to closed-loop agentic operations. The winning model is governed autonomy: the right mix of automation, workflows, and human in loop for every action based on business impact, risk, confidence, and policy.

AI Agent-first Managed Services introduce an operating model where intelligent agents sense issues, reason over enterprise context, plan actions, invoke approved tools, validate outcomes, and escalate exceptions to humans for judgment and governance.

Gartner predicts that by 2028 at least 15% of day-to-day work decisions will be made autonomously through agentic AI, up from 0% in 2024, and 33% of enterprise software applications will include agentic AI, up from less than 1% in 2024.^[1]

Forrester similarly highlights that 40% of highly regulated enterprises will combine data and AI governance, and that three out of four firms building aspirational agentic architectures on their own will fail.^[2]

The Evolution of Managed Services: From Support to Autonomy

Workplace technology has evolved into the foundation of employee experience. Today's workplace is a highly interconnected ecosystem of user devices, virtual desktops, Microsoft 365, SaaS applications, collaboration tools, identity services, network access, ITSM platforms, and security controls,

all of which shape the employee experience and contribute to productivity friction. Traditional ticket-driven operations react after disruption occurs, whereas AI agent-first operations proactively prevent issues, personalize support, and continuously enhance the employee experience.

The Maturity Curve: Five Eras of Workplace Managed Services

Era					
	Operating Model				
	Limitation				
	Target Outcome				
	Manual	Automated	Intelligent	Agentic	Autonomous
					
	Human-led support and troubleshooting	Workflow automation and self-service	AI-assisted support, analytics, and recommendations	AI agents & AI Assistants orchestrate tasks across approved tools	Self-healing, self-optimizing, outcome-led operations
	Slow, inconsistent, hard to scale	Rule-based, limited exception handling	Humans still drive most decisions and actions	Requires governance, telemetry, and automation fabric	Requires mature identity and observability controls
	Basic service continuity	Faster fulfillment, lower effort	Better insight and user experience	Actionable AI with human oversight	Resilient, proactive, trusted managed services

Core Pillars of AI Agent-First Autonomous Workplace

The Autonomous Workplace is built on the collaboration of Human, Digital, and AI Workforces. Humans provide strategy, governance, and judgment; the Digital Workforce delivers execution and automation at scale; and the AI Workforce adds intelligence, reasoning, and orchestration. Together, they create an operating model that shifts workplace services from reactive support to proactive, autonomous, and experience-led operations.



The Autonomous Workplace is built on the collaboration of Human, Digital, and AI Workforces. Humans provide strategy, governance, and judgment; the Digital Workforce delivers execution and automation at scale; and the AI Workforce adds intelligence, reasoning, and orchestration. Together, they create an operating model that shifts workplace services from reactive support to proactive, autonomous, and experience-led operations.

Pillar	💡 What It Enables		🎯 Managed Services Outcome
 Agentic AI	 Agentic AI and Goal Planning	AI Assistants & AI Agents understand goals, decompose work into tasks, select next actions, and adapt based on outcomes.	 Moves support from static scripts to contextual resolution paths.
	 Enterprise Automation Fabric	Reusable, versioned, approved automations across ITSM, endpoint, identity, collaboration, virtual desktop, and cloud platforms.	 Prevents one-off automation sprawl and accelerates reuse across accounts
 Action Layer	 MCP-style Action and Tool Layer	A standardized way for agents to discover capabilities, invoke approved tools, exchange context, and return traceable outcomes.	 Converts AI assistants from response engines into action-oriented managed services interfaces.
	 Digital Experience Intelligence	Endpoint, app, network, sentiment, and service telemetry that detects friction and prioritizes impact.	 Enables proactive remediation and experience-led service management.
 Experience and Context Layer	 Context and Knowledge Foundation	RAG, knowledge graphs, service history, CMDB, SOPs, known errors, policy context, and customer-specific rules.	 Improves accuracy, consistency, and personalization of AI decisions.
	 Context-Aware Decisioning	Decision logic incorporates user persona, incident priority, device state, policy, risk level, business criticality, and confidence.	 Routes the right action to full automation, approval, recommendation, or human intervention.
 Governance Layer	 Agent Governance and Trust	Responsible AI Policies, permissions, risk tiers, approval flows, auditability, explainability, testing, and rollback mechanisms.	 Creates safe, compliant, and auditable autonomy.
	 Security, Identity, and Data Protection	Least privilege, scoped service principals, secrets management, data classification, privacy controls, and tenant isolation.	 Prevents agent overreach and protects customer environments.
 Optimization Layer	 Observability and FinOps for AI Ops	Tracks agent decisions, tool calls, success rates, cost, latency, token usage, failure modes, and business value.	 Connects AI operations to SLA, XLA, productivity, and cost outcomes.
	 Continuous Learning and Improvement	Feedback loops from users, engineers, telemetry, incidents, and post-action validation improve future performance.	 Builds compounding managed services value over time.

Making AI Actionable: The Role of MCP and Action Layers

The value of agentic AI appears when AI moves beyond answering questions and becomes a governed execution layer for managed services. MCP-style action layers give AI Assistants and AI agents a controlled way to discover tools, understand allowed actions, and invoke workflows.

This changes the service model from ask-and-wait to sense-decide-act-validate. A user can describe an issue once; the agent can enrich the context, check policies, select the right automation, seek approval where required, execute the action, validate the experience signal, and update the ticket or knowledge base.

 Current State	 Agent-First Shift	 Business Impact
AI answers questions	AI initiates approved actions	Resolution speed improves
Bots work in silos	Agents invoke shared action catalogs	Reuse increases across accounts
Manual follow-up after advice	Closed-loop execution and validation	Tickets become outcome records
Custom integrations per use case	MCP-style contracts and reusable tools	Scale improves with lower engineering effort



Agent Governance, Trust, and Control Model

Agent governance defines which agents can act, which tools they can use, which data they can access, what approvals are required, and how every decision is audited. It converts agentic AI from experimentation into an enterprise-scale managed services capability.

Responsible AI strengthens this model by ensuring agent behavior remains fair, reliable, secure, inclusive, transparent, and accountable.

Governance Principles



Risk-tiered actions:
classify actions as advisory, approval-based, supervised, or autonomous.



Least privilege: use scoped identities and approved wrapper APIs.



Human oversight: route sensitive, irreversible, or low-confidence actions to people.



Responsible AI by design: embed fairness, safety, privacy, inclusiveness, transparency, and accountability across the agent lifecycle.



Evidence by design: log rationale, tool calls, approvals, results, and validation.

Agent Governance Operating Model


Governance Area	Responsible AI Lens	Core Control	Managed Services Outcome
 Strategy and Ownership	Accountability	Define agent council, business owner, technical owner, risk appetite, and priority domains.	Clear ownership and decision rights.
 Agents and Agent Registry and Lifecycle	Transparency + Accountability	Track agent purpose, owner, risk tier, tools, data sources, deployment status, and retirement plan.	Prevents shadow agents and improves lifecycle control.
 Action Catalog and Tool Control	Reliability and Safety	Expose only approved APIs, scripts, workflows, and MCP-style tools with defined read/write/execute scope.	Governed execution across enterprise systems.
 Risk and Autonomy Matrix	Fairness + Safety	Map each action to advisory, approval-based, supervised, or autonomous mode based on risk and impact.	Balanced speed, safety, and consistency.
 Identity, Access, and Data Protection	Privacy and Security	Use least privilege, scoped identities, vaulting, data classification, masking, and secure logs.	Protected enterprise and employee data.
 Responsible AI Review	Fairness + Inclusiveness + Transparency	Review agent behavior for bias, accessibility, explainability, user impact, and policy alignment.	Trusted and inclusive employee experience.
 Audit, Validation, and Rollback	Transparency + Reliability	Record decisions, prompts, tool calls, approvals, outcomes, exceptions, and rollback evidence.	Auditable and resilient operations.
 Value and Experience Dashboard	Accountability	Measure auto-resolution, MTTR, XLA, cost-to-serve, adoption, RAI findings, and employee experience.	Business value with continuous improvement.

High-Value Autonomous Workplace Use Cases by Industry

Industry context matters. The same agentic platform should be positioned differently based on workforce type, regulation, business operations, and employee experience priorities.

Industry	High-Value Use Cases	Business Outcome
 Banking and Financial Services	Access triage; device compliance; collaboration support; audit-ready ticket updates.	 Controlled productivity with stronger evidence.
 Healthcare and Life Sciences	Endpoint remediation; app access recovery; VDI support; engineer assist.	 Continuity for critical users.
 Manufacturing	Shared workstation self-heal; app deployment; plant collaboration support.	 Less downtime across plants.
 Energy and Utilities	Field device support; incident summarization; endpoint and collaboration monitoring.	 Resilient support for distributed teams.
 Retail and CPG	Store device self-heal; device use cases; onboarding support.	 Reduced support peaks and store disruption.
 Telecom and Media	Collaboration monitoring; endpoint compliance; service desk triage.	 Higher productivity in service operations.
 Professional Services	Client workspace setup; Collaboration use cases; user device automation.	 More billable productivity.
 Public Sector	Policy-governed fulfillment; compliance; accessibility and virtual workplace support.	 Transparent, accountable service delivery.
 Education	Self-service knowledge; device checks; virtual lab/desktop support.	 Better learning continuity.





Conclusion

Digital Workplace connected employees to tools, devices, and platforms. Autonomous Workplace connects intelligence, context, action, governance, and outcomes. The future managed services model is AI agent-first, governed, and experience-led.

Start with the use cases that create visible employee experience impact, expose trusted automations through controlled action layers, measure outcomes, and scale toward supervised and then trusted autonomy.

Sources

1. <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
2. <https://www.forrester.com/blogs/predictions-2025-artificial-intelligence/>

About the Author



Raman Nambirajan

Principal Technology Architect, Infosys Limited

Raman Nambirajan is a Principal Technology Architect with over 20 years of experience leading Microsoft Workplace Solutions and AI transformation initiatives. He partners with global organizations to reimagine workplace operations through Agentic AI, AI Agent-First Managed Services, and intelligent automation, helping enterprises transition from traditional support models to autonomous, experience-driven operations that deliver measurable business value.

For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.