


Beyond the Pilot: Building Agentic AI You Can Govern, Trust, and Own

Infosys
topaz
fabric



Infosys
Navigate your next



How Infosys Topaz Fabric helps you turn agentic ambition into reality

Every enterprise has an AI mandate – but few have a governed path from pilot to production. This Infosys point of view lays out why winners in the Agentic AI race won't be the ones with the most agents, but the ones who can prove value, control cost, and govern at scale, using the technology they already own. Further, it introduces Infosys Topaz Fabric, an open platform for building, governing, and running AI agents on the technology enterprises already have, and explores the Topaz Fabric Marketplace, where teams can pick from ready-made agents and blueprints to launch their own governed industrialized pilot in as little as 12 weeks.

The enterprise AI reality

Every enterprise has an agent.
Few have an advantage.

Walk into almost any large enterprise today and you'll find an AI agent running somewhere – in a support queue, a code pipeline, a finance workflow. The demo stage is over. What's missing isn't ambition or even capability; it's a governed, repeatable way to take that first working agent and turn it into a fleet the business can actually trust, audit, and defend.

That distinction between an agent that works in a demo and one that survives a security review, a cost audit, and a model change six months later – is where most transformation programs quietly stall. [Gartner's 2026 Hype Cycle](#) captures this same tension: agentic AI sits at the peak of inflated expectations, a signal that ambition has outrun proof.

Key shifts

~17%

of organisations have actually deployed AI agents today.

60%+

intend to deploy agents within two years – the steepest curve Gartner has measured

~20%

annual growth in the AIOps market now heading toward ~US \$11.8B by 2034.

The gap between those two numbers - 17% deployed versus 60% planning to be - is the real story. It's not a talent gap or a tooling gap. It's a governance gap: most organizations can build an agent faster than they can prove it's safe to run.



What is truly blocking organizations?

Positioning paralysis: Assistive pilots and agentic claims are everywhere. Leaders struggle to frame a differentiated AI strategy that is more than a bundle of point tools and demos.

Fear in governance and control: Buyers now raise safety, guardrails, auditability, and regulatory alignment (EU AI Act, NIST, ISO) before scale, not after.

Runaway and opaque cost: Token spend, model sprawl, and vendor lock-in make the total cost of agentic operations hard to predict and even harder to defend.

Execution risk at scale: The gap between a working prototype and a governed, production-grade agent estate is where most transformation programs hit a stall.

The differentiator in 2026 is no longer whether an enterprise is "agentic". It's whether it can answer the three questions every buyer is quietly asking: can you prove this is working, can you prove it's safe, and can you prove it's worth the cost. The enterprises pulling ahead aren't the ones asking how many agents they can deploy. They're the ones asking how many they can defend.

Topaz Fabric: From Ambition to Accountability

If the problem is proving agents are working, safe, and worth the cost, the solution has to be a foundation that treats all three as first-class, not afterthoughts bolted on once something breaks.

That's the thinking behind [Infosys Topaz Fabric](#), an open, flexible platform, delivered as Services-as-Software, that brings together everything an enterprise needs to design, run, and govern AI agents in one place. It pairs a space to build and refine agents (covering agent creation, workflow design, context management, testing, and reporting) with a space to actually run them day to day. All of this is backed by Infosys' delivery expertise – engineers embedded alongside the client, with human oversight built into the process at every step. It works across IT operations, transformation, quality engineering, and cybersecurity, and is designed to run on the models, cloud, and tools a client already has, so there's no need to rip out and replace existing investments.

Key impact



25–45%

higher productivity through more efficient delivery across the engineering and operations lifecycle



20–40%

better customer experience, measured through Net Promoter Score (NPS) improvement, from agentic experiences tailored to each user



600+

ready-to-use IT and industry agents, plus 25+ pre-built blueprints, all built with Responsible AI (RAI) guardrails from the start

One destination, multiple opportunities


That combination of scale and governance doesn't stay confined to one function. Because Topaz Fabric is built as a shared foundation rather than a set of point solutions, the same platform can be pointed at different parts of the enterprise – from the engineering teams building and running agents, to the industry and business teams looking for outcomes specific to their function. In practice, this shows up as three connected layers:

Fabric Foundation:

The core layer that powers agentic experiences, covering multi-agent coordination, enterprise context, small language models, a mix of AI services, knowledge engineering, and responsible-AI guardrails.

Fabric Business Services:

Applies the same foundation to industry- and function-specific outcomes, spanning Financial Services, Manufacturing, Human Resources, Marketing, and more.



Fabric IT Services:

Applies that foundation across the IT value chain, from transformation and engineering through to operations and innovation.



Build, run, and govern

The three pillars we discussed above only hold together if the underlying architecture is built the same way from the ground up; not assembled from parts that were never designed to talk to each other. That's the thinking behind our layered services stack for enterprise agentic AI: governance, evaluation, and observability aren't features added at the end, they run through every layer by design.



Decoding the services stack

LAYER
07

Agentic AI Experiences

The front end: portals, co-pilots, and AI-powered workspaces that adapt to the person using them

LAYER
06

Agent Orchestration

The coordination layer: planning, routing, and getting multiple agents to work together on a task

LAYER
05

Marketplace, Repositories and Registries

The reuse layer: ready-made context packages, shared tool connections, skill libraries, and a registry of available agents

LAYER
04

Enterprise Agent Core, Context and Harness

The building layer: the underlying agent framework, memory, knowledge bases, and the tools developers use to build, test, and ship agents

LAYER
03

Enterprise AI Data Foundation

The knowledge layer: a digital representation of the business, connections to enterprise tools and APIs, and how that knowledge is organized and managed

LAYER
02

Model Gateway

The routing layer: choosing between models, managing cost, and handling fallback and rate limits automatically

LAYER
01

Poly-AI Cloud Platform and Runtime

The foundation layer: where everything actually runs – sandboxed, able to scale up or down, with storage, disaster recovery, and managed services built in.

What ties every layer of the stack together are three commitments that run underneath all of it. Governance, safety, and compliance are handled through guardrails, role-based and attribute-based access control (RBAC and ABAC), audit trails, PII handling, human-in-the-loop review, and readiness for regulatory compliance. Evaluation and observability keep the system accountable,

through logs, reasoning and tool traces, metrics, continuous evaluations, cost attribution, and drift detection. And deployment is a matter of choice, not constraint: enterprises can run on an open stack or a cloud-managed one, on Kubernetes or serverless (Knative), with first-class support for Azure.

Make scaling repeatable

A layered architecture only delivers on its promise if the people building on it can actually work quickly and consistently, which is where Topaz Fabric Studio comes in. It's an enterprise-grade workspace for building and orchestrating agents, designed to work whether someone prefers a no-code interface or wants to write code directly, through both a web interface and a VS Code extension. Built on multiple existing frameworks (Google ADK, LangGraph, AWS Strand), it's organized around five capability pillars that make reuse, controlled access, and day-to-day visibility the default rather than an afterthought:



Create :

Build using ready-made components and repositories for agents, tools, MCP servers, knowledge bases, and guardrails, with the ability to build, import, and export across frameworks - so teams can build something once and reuse it in different contexts.

Orchestrate:

Design workflows visually using the Flow Builder, test them in a simulator before going live, and add human approvals, scheduled triggers, and coordination between multiple agents working together.

Govern:

Manage access through role-based and identity access controls (RBAC/IAM, via Keycloak or Microsoft Entra ID), handle secrets securely, and apply responsible-AI guardrails consistently across models, agents, and tools.

Operate & Improve

Test and refine agents through multiple evaluation methods, including a model reviewing another model's work, along with explainability, diagnostics, and ongoing observability to catch issues before they affect reliability.

Integrate from Day 1

Connect out of the box to a broad tool ecosystem, including ServiceNow, BMC Helix, ClickHouse, Microsoft Teams, Datadog, Splunk, SAP, Cisco, Nokia, and more.

Capability powerhouse: Combining Core, Studio, and Fabric

Agent Builder

Build agents with zero-code or pro-code using Google ADK, LangGraph, or AWS Strand. Test configurations, orchestrate multi-agent workflows, add human approvals, and export code to ZIP or VS Code.

Flow Builder

Visually design, simulate, and orchestrate workflows using ARGO or LangGraph. Support long-running processes with Timer nodes and auditable HITL reviews.

Context Builder

Create and manage Prompts, Knowledge, and Skills with versioning, governance, built-in libraries, and import/export support.

Agent Evaluation

Test agents with Black-box, White-box, or Grey-box evaluations, configurable metrics, multi-model reviews, and detailed reports.

Shell, Client SDK, Apps

Access a unified experience with SSO and RBAC, extend with Python and TypeScript SDKs, and deploy through a built-in apps marketplace.

Enterprise Twin

Build a queryable digital representation of your business to give agents trusted business context and semantic understanding.

Insights (Observability and FinOps)

Monitor agent health, adoption, and infrastructure, while tracking usage, budgets, spend, rate limits, and cost forecasts.



From ambition to advantage

Where the market competes on features,
we compete on outcomes

Every layer of Fabric, from the architecture to the Studio workspace, exists because of the same five problems enterprises keep running into.

Your challenge



Positioning paralysis

Undifferentiated pilots and point tools, with no clear way to stand out from the crowd.



Governance & control fear

Concerns around safety, auditability, and regulatory alignment.



Runaway & opaque cost

Token sprawl and vendor lock-in driving up cost.



Execution risk at scale

A gap between a working prototype and a production-ready system.



Keeping investments

No appetite to rip out and replace what's already in place.

Our solution

Breadth across IT operations, transformation, quality engineering, and cybersecurity, all on one composable platform.

Built-in agent evaluation (black-box, white-box, and grey-box, with model-led review), responsible-AI guardrails aligned to the EU AI Act, NIST, and ISO, role-based access control, human-in-the-loop review, and full audit trails.

Built-in cost management across budgets, rate limits, and forecasting by team or model, plus the freedom to choose models, frameworks, and clouds - including smaller or open-source models to keep running costs down.

Infosys delivery backed by engineers embedded with the client, 600+ existing agents and 25+ blueprints, a Flow Simulator for testing before go-live, and a structured process for maturing pilots into production.

A platform that works above existing models, clouds, and tools - where connections to enterprise systems go beyond simply rebuilding what's already there.

Taken together, this is a reframe: from deploying more agents, to owning, extending, and governing the agents an enterprise already has, on the stack it already owns, with Infosys engineers accountable for the outcome.

The value lifecycle

Innovate and operate for what's next

Topaz Fabric treats value as a continuous loop, not a one-off project: it helps turn an idea into a shipped, verified outcome, then keeps that outcome running securely, moving through discovery, delivery, and realization, powered by the Vibe Agent and grounded in enterprise context that compounds over time.

Discovery starts with building the right thing - through deep research, product discovery, and PRD creation, ideated quickly in [Living Labs](#) and incubated through vibethons and hackathons.


Delivery is about building it fast and building it right, with agent-enabled requirements, grounded code generation, and agent-driven testing and deployment.

Realization means running it securely, efficiently, and reliably, through proactive monitoring, predictive triage, automated remediation, and continuous learning - reliable through full-stack observability and anomaly detection, secure through agentic

vulnerability scans and rapid remediation, efficient through zero-touch automation with human review where it matters, and business-aligned through a Business Command Center that ties IT signals directly to business outcomes. At the center of it all is Vibe Studio, anchored by the Vibe Agent, which turns natural-language instructions into working, enterprise-ready software - fully version-controlled in Git, fully portable, with no lock-ins.

From Idea to Impact in 12 Weeks, with the Infosys I3 Framework

The pilot project stage has a shelf life - and most enterprises are letting it expire. With the Infosys I3 Framework and Topaz Fabric, the path from a promising idea to a governed, production-ready agentic AI pilot doesn't have to take a year of committees and re-pilots. It takes 12 weeks, with measurable business outcomes to show for it.



12-week pilot: Put it into action

Start with proven building blocks, then create, test, and govern your own pilot in a live environment.

Ready to Start →

About Infosys

Infosys is a global leader in next-generation digital services and consulting. Over 320,000 of our people work to amplify human potential and create the next opportunity for people, businesses and communities. We enable clients in 59 countries to navigate their digital transformation, powered by cloud and AI. We enable them with an AI-first core, empower the business with agile digital at scale, and drive continuous improvement through always-on learning across our innovation ecosystem.

Visit www.infosys.com to see how Infosys (NSE, BSE, NYSE: INFY) can help your enterprise navigate your next.

Experience Fabric

Explore the Fabric Marketplace and Vibe Studio to build, evaluate and govern your first agents.

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.