The Right-Sized Frontier Roadmap

A Corporate IT Leader's Guide to Microsoft's Agentic AI Journey

Building an AI-fueled organization with the team you have.

Microsoft

softwareone

Contents

1. Introduction 3
2. Two Analyst Firms, One Signal 4
3. What the Research Means for a Smaller Team 5
4. Building Governance Into the Rollout 6
5. Planning for Usage-Based Cost 7
6. What “Right-Sized” Really Means 9
7. Where Does Your Team Stand? 10
8. Choosing Where to Begin 11

Introduction

1

Frontier is Microsoft's framework to help companies securely transition into the agentic AI era on a platform built to perform, with governance holding it all together. Corporate IT teams have weathered a lot of change over the last decade, but agentic AI might be the biggest shift yet.

Microsoft designed Frontier for enterprise, and its own playbook for becoming a Frontier Firm names five pillars for doing it well: governance, implementation, change management, measurement, and support. Each pillar assumes a level of staffing that a large organization can accommodate. Most of the current research from IDC, Gartner, and Forrester applies more organically to organizations with dedicated AI governance leads, application teams that build custom agents, and budget cycles built around a multi-year rollout.

A Corporate-segment IT leader encountering those industry reports faces that translation in whatever hours remain after the rest of the team's work is completed. The director might carry the accountability alone, while a handful of generalists split licensing, security, and support across whatever else comes with a deadline this quarter. (Nobody owns data science because the new hire has not yet been approved.)

In this space, SoftwareOne excels at taking what Microsoft built for the largest organizations and making it practical and workable for midsize businesses. If you're a director already covering licensing, security, and support with Frontier newly added to your list, this guide is for you.

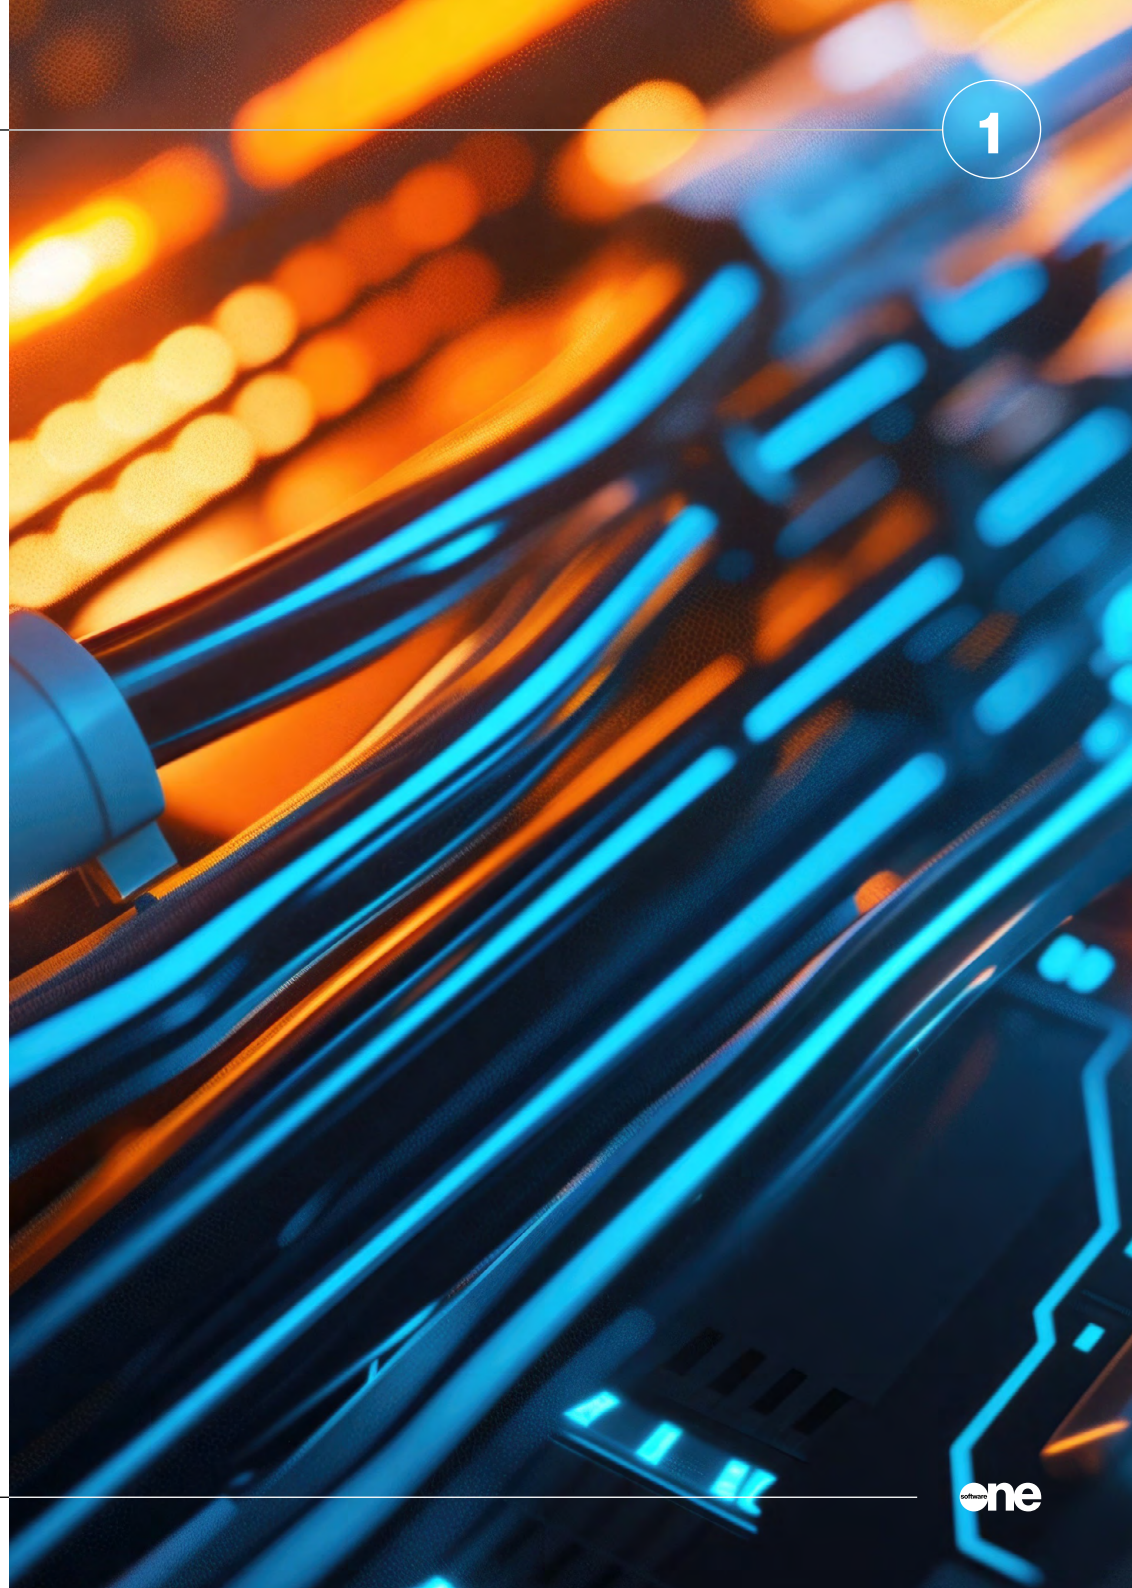

Two Analyst Firms, One Signal

IDC's FutureScape 2026 research projects that by the end of 2026, 40% of all Global 2000 job roles will involve working with AI agents.[1] IDC frames this as the start of a longer arc, moving from isolated pilots today to enterprise-wide orchestration by 2030.

Gartner's most recent forecast lands on a similar timeline from the application side, predicting that task-specific AI agents will be built into 40% of enterprise applications by the end of 2026, up from under 5% in 2025.[2]

When two of the largest independent technology research and advisory firms in the world measure the same shift, that shift is real and deserves attention. But a five-person team reading the same reports has one practical question: What does any of this mean at our size?

1. IDC. "IDC FutureScape 2026 Predictions Reveal the Rise of Agentic AI and a Turning Point in Enterprise Transformation." Business Wire, October 23, 2025.

2. Gartner. "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025." Press release. August 26, 2025.

What the Research Means for a Smaller Team

Large enterprises currently lead agentic AI deployment for several structural reasons. Specialist headcount and multi-year budget cycles fund custom agent infrastructure that a smaller team lacks the capacity to build, though such scale comes with its own drag: more approval layers, more legacy systems to integrate around, and a slower crawl from pilot to production.

Smaller organizations start from a different place. Frontier removes the need to build agent infrastructure from scratch, which changes what a lean IT department has to do to get a pilot running. A smaller organization also answers to fewer approval layers than an enterprise while keeping more budget room than a very small business, and that combination is the real advantage: it lets a lean team stand up a live pilot in weeks instead of quarters, as long as governance is built in before any agents go live.

Being differently positioned is not the same as being behind. In fact, it can carry a structural advantage if the rollout is approached deliberately. The “right-sized” framework ahead turns that positioning into a practical way to locate exactly where a specific organization sits before committing to a pilot.

Building Governance Into the Rollout

4

The research is also specific about why a Corporate-segment IT leader should treat governance as a first priority.

Forrester's 2026 technology and security predictions describe an AI market facing more financial and governance demands, as the gap between vendor promises and delivered value comes under closer scrutiny.

Weak governance of AI agents could expose organizations to real disruption by 2030, IDC's report suggests, including lawsuits, fines, and leadership turnover across a meaningful share of large enterprises.

Large organizations can respond to such exposure by creating dedicated AI governance roles built around this risk, but a team of five needs to squeeze governance work within the same hours as everything else on their calendars.

Planning for Usage-Based Cost

5

Budget and governance are connected more directly than most vendor conversations recognize.

Forrester's 2026 predictions expect enterprises facing greater financial scrutiny to defer one-fourth of their planned AI spend into 2027, and fewer than one-third of AI decision-makers can currently tie their AI investment to a measurable financial return.[3] That pattern describes the kind of correction most new technology categories eventually go through: a market moving away from hype and toward disciplined adoption.

Corporate-segment IT teams already operate with less budget flexibility than an enterprise and a less forgiving approval process. AI spending has to show results early, not later, because there's no slack in the process to absorb a slow-burning pilot.

3. Forrester. "Forrester's 2026 Technology & Security Predictions: As AI's Hype Fades, Enterprises Will Defer 25% of Planned AI Spend to 2027." Press release. October 28, 2025.

Planning for Usage-Based Cost

5

On top of that, Frontier adds another wrinkle. With usage-based billing replacing the fixed per-seat licensing most IT budgets are built around, usage-based cost forecasting needs to start on day one. Copilot Cowork, the agentic layer that runs multi-step tasks across Microsoft 365 rather than answering a single prompt, bills on usage in Copilot Credits, on top of the standard Copilot seat. What any one task costs depends on the model it uses, how much organizational context it pulls in, how many tools it calls, and how long it runs. Microsoft's own published estimates put a light task around one to three dollars, a medium task at four to seven, and a heavy one well past that. A single heavy task could cost more than a full day of the user's seat license.

Teams that do the math from day one aren't surprised when the first invoice arrives. The Microsoft 365 Admin Center provides a cost-management dashboard for tracking credit consumption by user and team, spend limits set at the tenant level, and a per-task cost readout for anyone who wants to see what a job cost.

> For a Corporate-segment team, SoftwareOne's Managed FinOps practice can help turn that forecasting habit into a predictable line item (instead of a surprise Teams ping from finance).

What “Right-Sized” Really Means

6

Most agentic AI maturity models describe four basic stages. These stages, while accurate, assume a larger team will staff each phase separately and have time for the analyst reports and business cases.

A five-person IT team can’t do any of that, but the four stages still apply if they’re sized accordingly.

Stage 1: Explore.

A realistic timeline for a lean team: 2-4 weeks of reading, internal alignment, and a hard look at how Copilot is already being used across the organization before committing to a pilot team. Most Corporate-segment shops have some Copilot footprint already, whether licensed deliberately or adopted ad hoc, and that existing usage is the most honest signal of where an agent pilot will land first. No new agent goes live in this stage.

Stage 2: Pilot.

A single team, one measurable process, and one agent or small group of agents, built on Copilot and, where the work calls for a custom agent, assembled in Foundry. The governance charter and access model both get written before day one: who approves the agent, what data it can reach, who owns the escalation path. 6-8 weeks from kickoff to a readable result.

Stage 3: Scale.

One additional team or function gets added at a time. No org-wide rollout happens in a single motion. Each addition reuses the governance charter and access model built in Stage 2, so the second pilot starts from a template instead of a blank page.

Stage 4: Operate.

Agent governance becomes a standing agenda item in an existing IT operations rhythm where Agent 365 does the heavy lifting. Its controls for managing, securing, and monitoring agents from one place allow a lean team to supervise a growing roster of agents. This is also where usage-based cost forecasting matters most, since operating costs become the recurring line item to manage.

Where Does Your Team Stand?

Five questions help sort out which stage accurately fits your situation:

01 Is there a named business sponsor outside IT for a candidate pilot team?

02 Do you know who approves a new agent before it goes live?

03 Is the current default agent access setting known, and is it curated or set to allow all?

04 Is there one process with a measurable, currently-tracked baseline (time, volume, or error rate)?

05 Has finance ever seen a usage-based cost model before, in any context?

Three or more “no” answers point to Stage 1. Mostly “yes” answers point to Stage 2 (ready to pilot). The most important aspect of this is taking an honest read of your agentic program’s current state.

Choosing Where to Begin

8

Governance and budget both point toward the same first decision: one team, one measurable process with real stakes, and a business sponsor who has something to gain from the result.

More than the specific agent capability deployed, that choice shapes almost everything about how the 90-day pilot moves forward. A well-chosen pilot team makes governance and budget forecasting easier to execute because the scope narrows to one team's access and one team's usage curve, instead of what amounts to org-wide guesswork.

Getting frameworks built for enterprises to fit a leaner team, with governance and cost control in place from the first pilot, is where an experienced partner demonstrates its value. Every day, SoftwareOne helps midsize businesses and Corporate-segment IT teams take Microsoft's enterprise-grade Frontier stack and make it workable at their size.

Explore SoftwareOne's Data and AI Services to see how that support takes shape, or reach out about a Frontier readiness assessment before you begin your next pilot.

Any questions?

Let's talk

Find out more:

www.softwareone.com

SoftwareOne, Inc. | US Headquarters

T. +1 800 444 9890

E: connect.us@softwareone.com

Copyright © 2026 by SoftwareOne AG, Riedenmatt 4, CH-6370 Stans. All rights reserved. SoftwareOne is a registered trademark of SoftwareOne AG. All other trademarks are the property of their respective owners. SoftwareOne shall not be liable for any error in this document. Liability for damages directly and indirectly associated with the supply or use of this document is excluded as far as legally permissible. © Imagery by: Adobe Stock and Getty Images.