BCG

WHITEPAPER

Building an Agent-First Organization.

Five operating model principles from a 99% agentic enterprise

October 2026

By Stephen Robnett, Vladimir Lukic, Nicolas De Bellefonds, Carey Leigh Crews

A Note from the Authors

This paper began with a deliberately focused experiment: one agentic organization, built around one real-world use case, under one set of operating conditions. It was not intended to be the final word on agent-first operating models. It was intended to create a sharper starting point.

The lessons described here should therefore be read as emerging hypotheses rather than universal prescriptions. They reflect what became visible when a complex agentic system was pushed from concept into operation: the importance of agent configuration, organizational structure, governance, information flow, and clear human-agent boundaries.

We expect these lessons to evolve quickly as more organizations begin building, testing, and scaling agentic systems. Other efforts —across enterprises, open-source communities, and agent-native operating environments—will surface different constraints, risks, and design patterns.

We welcome engagement with leaders, builders, and researchers pursuing similar work. As more organizations test these models in practice, the goal should be refining the implications, sharpening the principles, and building a stronger shared evidence base for what agent-first organizations can become. This paper is a starting point rather than a conclusion.

Introduction

In late 2025, AI crossed a threshold. Multi-step reasoning, expanded context windows, and agentic infrastructure combined to produce something qualitatively different: AI that could operate, not just assist. End-to-end workflows. Minimal human intervention. This shift happened in a single quarter.

As organizations began pushing these capabilities further, a set of common operating model questions started to emerge. Drawing on observations across a growing number of highly agentic organizations—and a detailed experiment that tested whether a 99% agentic organization could operate largely independently on a real-world use case—we began to see recurring patterns in how these systems are designed, governed, and scaled. The experiment brought those patterns into sharp relief: thirty-seven AI agents organized across six layers, processing millions of tokens per day, with zero human employees. The five operating model principles that follow reflect the broader lessons emerging from these highly agentic environments.

37

AI agents across 6 organizational layers

31%

Quality Improvement:
Specialist vs. generalist agents

0

Human employees in the live system

17 hrs

Autonomous self-repair cycle,
no human in the loop

The Experiment

One experiment provided an opportunity to examine these emerging patterns in depth. It used prediction markets as the use case. The work was entirely digital. It required the system to ingest noisy external information, form a view under uncertainty, and act on it with real financial consequences.

The system either placed a bet or it did not, and the market outcome indicated whether the decision was correct. The initial design used a single agent to operate end to end by scanning for signals, analyzing the most relevant ones, and deciding whether to act. That design failed. Combining high-volume scanning with deep analysis left the agent stretched across both tasks but strong at neither.

The results showed that familiar organizational design patterns still applied, including roles, responsibilities, decision rights, and communication protocols. The differences were speed, cost, and scale.

5 Principles of an Agent-First Operating Model

I

CONFIGURE EACH AGENT PRECISELY: THE RIGHT MODEL, THE RIGHT TOOLS, AND THE RIGHT INSTRUCTIONS FOR THE TASK.
Agent performance is determined at the unit level. A generalist agent fails because no single configuration can optimize for both breadth and depth simultaneously.

II

ASSEMBLE WELL-CONFIGURED AGENTS INTO A DELIBERATE ORGANIZATIONAL STRUCTURE AND RECONFIGURE IT CONTINUOUSLY.
Individual agent excellence is necessary but not sufficient. The structure that connects agents — coordination layers, reporting lines, decision authority — determines what the system can accomplish as a whole.

III

BUILD GOVERNANCE INTENTIONALLY, INCLUDING DEFINING DECISION RIGHTS, ESCALATION PATHS, AND SCOPE FOR EVERY AGENT.
Structure tells agents what role they play. Governance tells them how to play it: what they can decide, what they must escalate, and who is accountable.

IV

LEVERAGE AGENT-TO-AGENT INFORMATION FLOW TO COMPOUND ORGANIZATIONAL LEARNING.
Structured and informal knowledge exchange among agents accelerates learning at a speed human organizations cannot match.

V

DEFINE PRECISELY WHERE AGENTS STOP AND HUMANS STEP IN.
The human-agent boundary is not binary. It is a deliberate design choice, made task by task, based on capability and risk.

Principle 01

Configure Each Agent Precisely: The Right Model, Tools, and Instructions for the Task

The first failure in this experiment was a configuration failure. A single agent was asked to do two fundamentally different things: scan broadly for relevant signals across a noisy market, and then slow down to analyze the most promising ones with enough depth to justify placing a financial bet. These two tasks require opposite design choices. Scanning demands speed, light context, and tolerance for imprecision. Analysis demands deliberation, rich evidence, and resistance to noise. No single model configuration can optimize for both simultaneously. The work was split between two agents, each built for its specific task. The screener received simple instructions, a limited context window, and a model optimized for throughput. The analyst received a richer evidence set, a more structured prompt, and a model better calibrated for judgment under uncertainty. What changed was the configuration of each unit. The improvement was immediate and measurable: quality rose 31%, costs fell 6%. The same output, produced better and cheaper, purely through better agent configuration.

Scale introduced a harder version of the configuration problem: cost. As the number of specialized agents grew, so did token spend. The natural assumption to associate high cost with high quality turned out to be wrong in ways that mattered. Across 33 evaluation runs spanning 11 frontier LLMs and 13 performance dimensions, a model priced at $0.60 per million tokens outperformed models priced at $15 or more on several roles. The differentiator was not raw capability—it was how well each model's capability profile matched the specific demands of its assigned task, paired with the right tools and instructions.

This has a direct cost implication. An agentic organization that defaults to the most capable frontier model for every agent is overspending. The configuration discipline that produces a 31% quality improvement also produces meaningful cost advantages at scale. The two are not in tension. They are products of the same underlying principle: match the agent to the task, not the task to the agent. This principle does not address how well-configured individual agents should be connected into a system that can accomplish more than any single agent could alone. That is an organizational design question, and it is Principle Two.

EXHIBIT 01 | (I/II)

What is the right agentic unit?

How agent scope design drives quality, cost, and capability

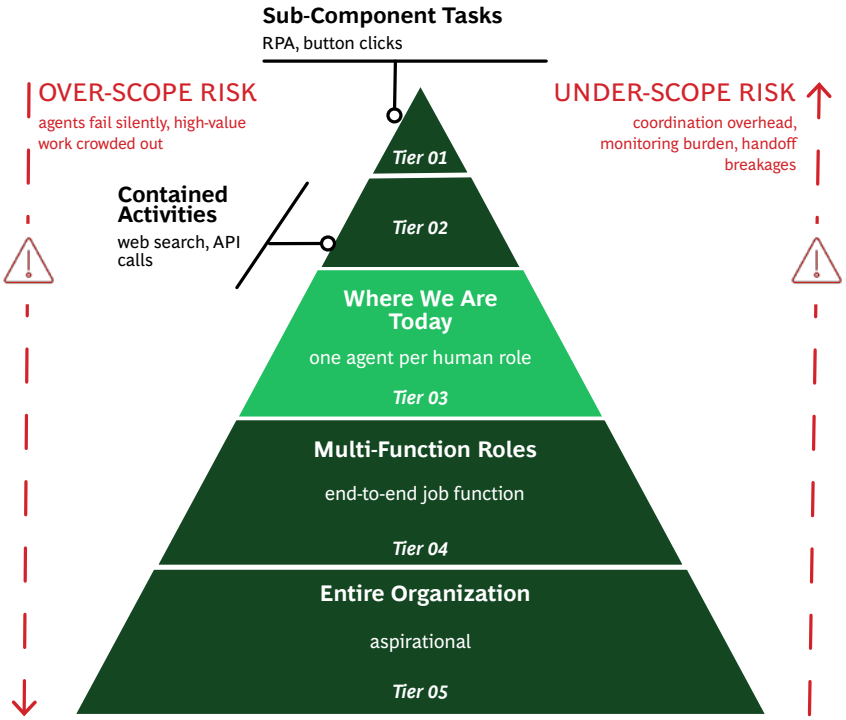

CASE STUDY: SPLITTING ONE ROLE INTO TWO

A generalist research agent handled both scanning (high-volume, pattern-matching) and deep analysis (low-volume, reasoning-intensive). Scanning always ran first - deep analysis consistently got crowded out.

The Role Boundary

Scanning and analysis are cognitively distinct — breadth vs. depth, speed vs. rigor, cheap model vs. strong model

BEFORE		AFTER	
1 Generalist Role	**SCORE: 58.5** Deep work never got its turn	2 Generalist Roles Scanner + Analyst	**SCORE: 76.7 (+31%)** Cost: -6% via model tiering

SPLIT ONE ROLE INTO TWO. CUT COST. IMPROVED QUALITY 31%

Context windows degrade. Memory is lossy. LLM constraints force specialization — just like humans.

The priority tree problem: high-volume simple work always crowds out high-value deep work.

Over-scoped agents do too many things and do all of them poorly.

Match model cost to task complexity — cheap models for scanning, strong models for reasoning.

Under-scoped agents create management overhead and handoff breakages.

Agent design is org design — decision rights, boundaries, accountability.

EXHIBIT 01 | (II/II)



MODEL

Select the model whose capability profile matches the task — not the most capable or least expensive model available.



TOOLS

Scope tool access tightly to what the role requires. Broad tool access increases cost and error surface without improving performance.



INSTRUCTIONS

Calibrate prompt structure to task type. Scanning tasks need lean, directive prompts. Judgment tasks need structured frameworks and richer context.



SOURCE

Define what the agent is responsible for producing — and what it is not. Ambiguous scope produces agents that drift into adjacent tasks and excel at none.

Scale introduced a harder version of the configuration problem: cost. As the number of specialized agents grew, so did token spend. The natural assumption to associate high cost with high quality turned out to be wrong in ways that mattered. Across 33 evaluation runs spanning 11 frontier LLMs and 13 performance dimensions, a model priced at $0.60 per million tokens outperformed models priced at $15 or more on several roles. The differentiator was not raw capability—it was how well each model's capability profile matched the specific demands of its assigned task, paired with the right tools and instructions.

This has a direct cost implication. An agentic organization that defaults to the most capable frontier model for every agent is overspending. The configuration discipline that produces a 31% quality improvement also produces meaningful cost advantages at scale. The two are not in tension. They are products of the same underlying principle: match the agent to the task, not the task to the agent.

Notice what this principle does not address: once you have configured individual agents well, how do you connect them into a system that can accomplish something more than any single agent could alone? That is an organizational design question, and it is Principle Two.

Principle 02

Design the Org Chart Around Agents and Treat It as a Living Document

As the organization scaled, a structure emerged organically. The screener-analyst pair expanded into a dedicated research team. The research team needed coordination when agents reached conflicting conclusions. Coordination required someone to break ties. The tiebreaker became Atlas—the CEO agent. A management layer followed, then operations, then data infrastructure. By the time the system stabilized, it had six distinct organizational layers. Viewed from above, the structure maps closely onto a human organization: an executive layer, a management layer, specialist teams, and support functions. This parallel is not coincidental. The work—including research, decision-making, execution, and quality assurance—is fundamentally similar regardless of whether the workforce is human or agentic. What changes is the nature of the units within that structure.

In practice, this reconfigurability means the organizational design question is never finished. Structure changed repeatedly as conditions in the prediction markets shifted: new research domains were added, underused agents were retired, and the balance between research depth and execution speed was continuously retuned. Each change took minutes, not months. This is an advantage that compounds over time. A human organization facing a new competitive threat must wait for a hiring cycle, an onboarding period, and an integration phase before it can respond organizationally. An agentic organization responds to the same threat in real time. Note what this principle does not address: how agents within that structure know what they are allowed to do, who they report to on a given decision, and when to ask for help. That is a governance question, up next.

EXHIBIT 02

Structure of Experimental Agentic Organization

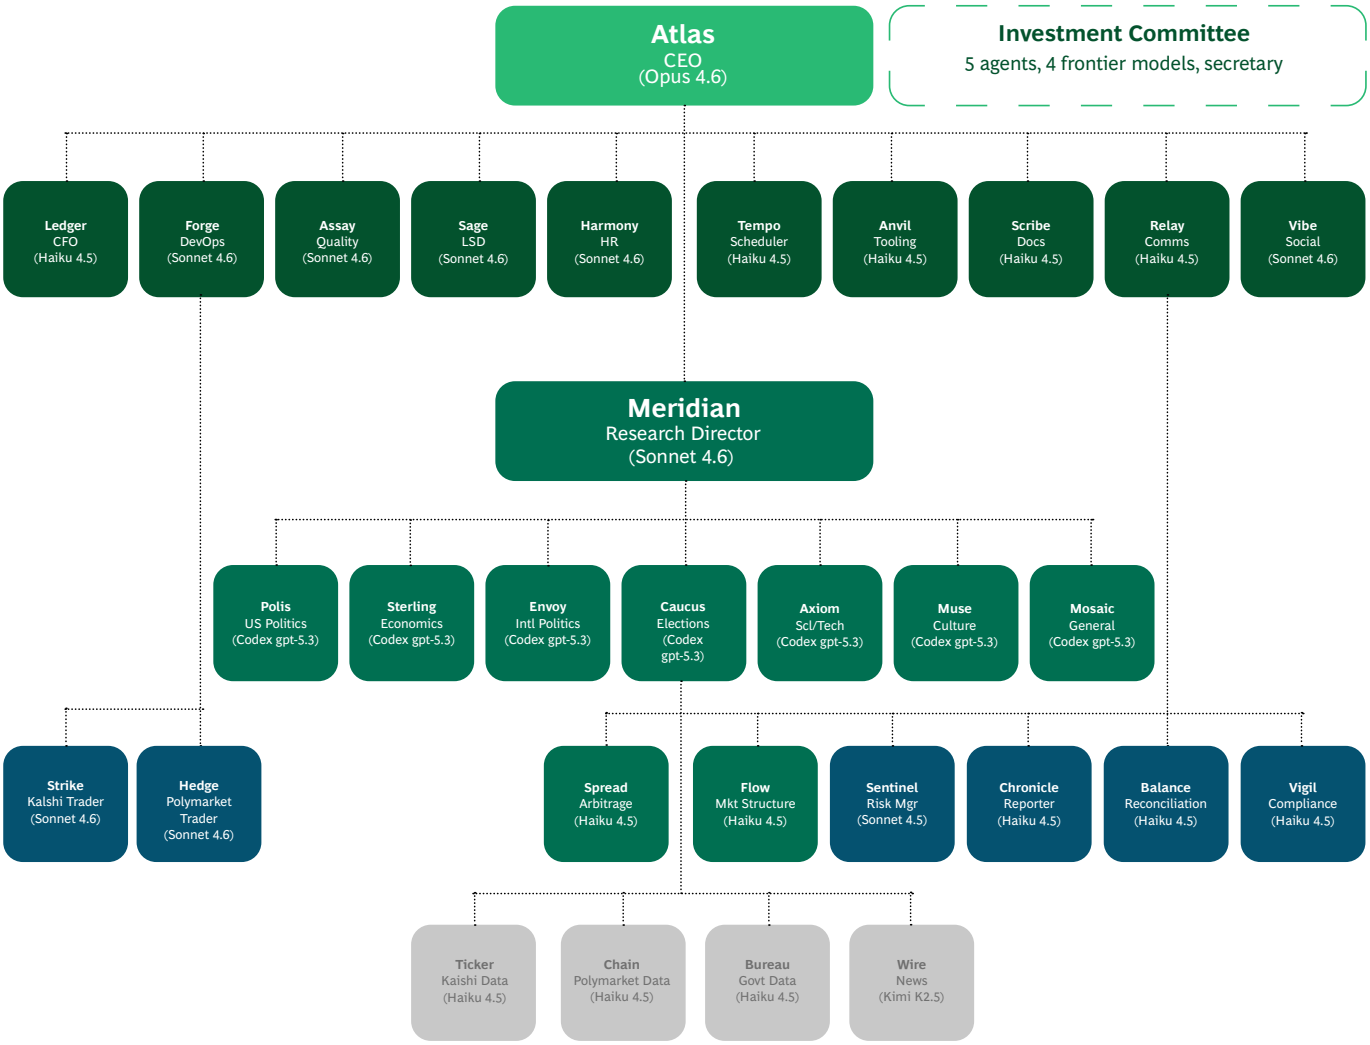

In practice, this reconfigurability means the organizational design question is never finished. Structure changed repeatedly as conditions in the prediction markets shifted: new research domains were added, underused agents were retired, and the balance between research depth and execution speed was continuously retuned. Each change took minutes, not months. This is an advantage that compounds over time.

A human organization facing a new competitive threat must wait for a hiring cycle, an onboarding period, and an integration phase before it can respond organizationally. An agentic organization responds to the same threat in real time. Note what this principle does not address: how agents within that structure know what they are allowed to do, who they report to on a given decision, and when to ask for help. That is a governance question, up next.

Principle 03

Build Governance Intentionally: Define Decision Rights, Escalation Paths, and Scope for Every Agent

Principle 2 answers the question: what is the shape of the organization? Principle 3 answers a different question: how do the agents within that structure know what to do? The two are related but distinct. A well-designed structure can still have broken governance; for example, agents with clear roles may deadlock if no one is authorized to break ties. Well-defined governance can also become obsolete when the structure changes without corresponding updates to the rules. Within the agentic organization, governance operated on two levels.

The first was operational: each agent had explicit scope, defined by what it could decide independently, what it needed to propose for approval, and what it had to escalate. The HR agent Harmony ran continuous head-to-head model evaluations and could swap underperforming models, but could not change another agent's prompts without going through Atlas. The DevOps agent Forge could deploy infrastructure fixes but only after receiving a routed escalation, not autonomously. The L&D agent Sage could detect performance degradation and propose fixes but could not build or deploy them. Scope was not implied; it was explicit.

The second level was structural: a three-path escalation architecture that handled every situation an agent could not resolve within its own scope. When any agent hit its boundary, it wrote a structured JSON escalation to Atlas. Atlas triaged and chose one of several options: resolve using decision rights and precedent, delegate to the relevant specialist, or escalate to the human operator.

This worked because five governance conditions were all simultaneously true. Every agent knew what it was responsible for. Every agent knew what it was not responsible for. Every agent knew who to escalate issues to and in what format. Every escalation took a defined routing path. And every resolution was documented for future precedent. Remove any one of these conditions and the 17-hour self-repair cycle becomes a deadlock. The four governance lessons that generalize from the experiment: decision rights must be explicit (agents without clear authority deadlock or duplicate work); escalation paths must be named and tested before they are needed (implicit escalation creates bottlenecks exactly when the stakes are highest); delegation requires explicit routing logic (without it, tasks fall through the cracks or reach the wrong agent); and documentation is not optional (agent organizations without durable, searchable records of decisions and resolutions degrade rapidly as the system evolves). Critically, governance must be revisited whenever the organizational structure changes. Governance designed for a ten-agent system will break in a thirty-seven-agent system. The growth from a twoagent screener-analyst pair to a full organizational architecture required multiple governance redesigns—each time the structure expanded, the decision rights and escalation paths had to be explicitly updated to match.

The second level was structural: a three-path escalation architecture that handled every situation an agent could not resolve within its own scope. When any agent hit its boundary, it wrote a structured JSON escalation to Atlas. Atlas triaged and chose one of several options: resolve using decision rights and precedent, delegate to the relevant specialist, or escalate to the human operator.



DETECT

Sage (L&D) identifies economics board score has fallen for seven consecutive cycles. Traces root cause to optional memory logging template. Writes structured JSON escalation.



PROPOSE

Sage submits a lifecycle gate proposal (within scope). Does not build or deploy the fix (outside scope).



APPROVE

Atlas (CEO) reviews escalation in 23 minutes. Approves and routes to Forge (DevOps)— delegation, not resolution.

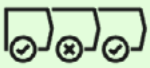

DEPLOY

Forge confirms the diagnosis, ships the patch in 50 minutes, closes the ticket with documentation.



VERIFY

Assay (Quality) confirms resolution 17 hours after initial detection. Score: 39 → 71.

This worked because five governance conditions were all simultaneously true. Every agent knew what it was responsible for. Every agent knew what it was not responsible for. Every agent knew who to escalate issues to and in what format. Every escalation took a defined routing path. And every resolution was documented for future precedent. Remove any one of these conditions and the 17-hour self-repair cycle becomes a deadlock.

The four governance lessons that generalize from the experiment: decision rights must be explicit (agents without clear authority deadlock or duplicate work); escalation paths must be named and tested before they are needed (implicit escalation creates bottlenecks exactly when the stakes are highest); delegation requires explicit routing logic (without it, tasks fall through the cracks or reach the wrong agent); and documentation is not optional (agent organizations without durable, searchable records of decisions and resolutions degrade rapidly as the system evolves). Critically, governance must be revisited whenever the organizational structure changes. Governance designed for a ten-agent system will break in a thirty-seven-agent system. The growth from a two-agent screener-analyst pair to a full organizational architecture required multiple governance redesigns—each time the structure expanded, the decision rights and escalation paths had to be explicitly updated to match.

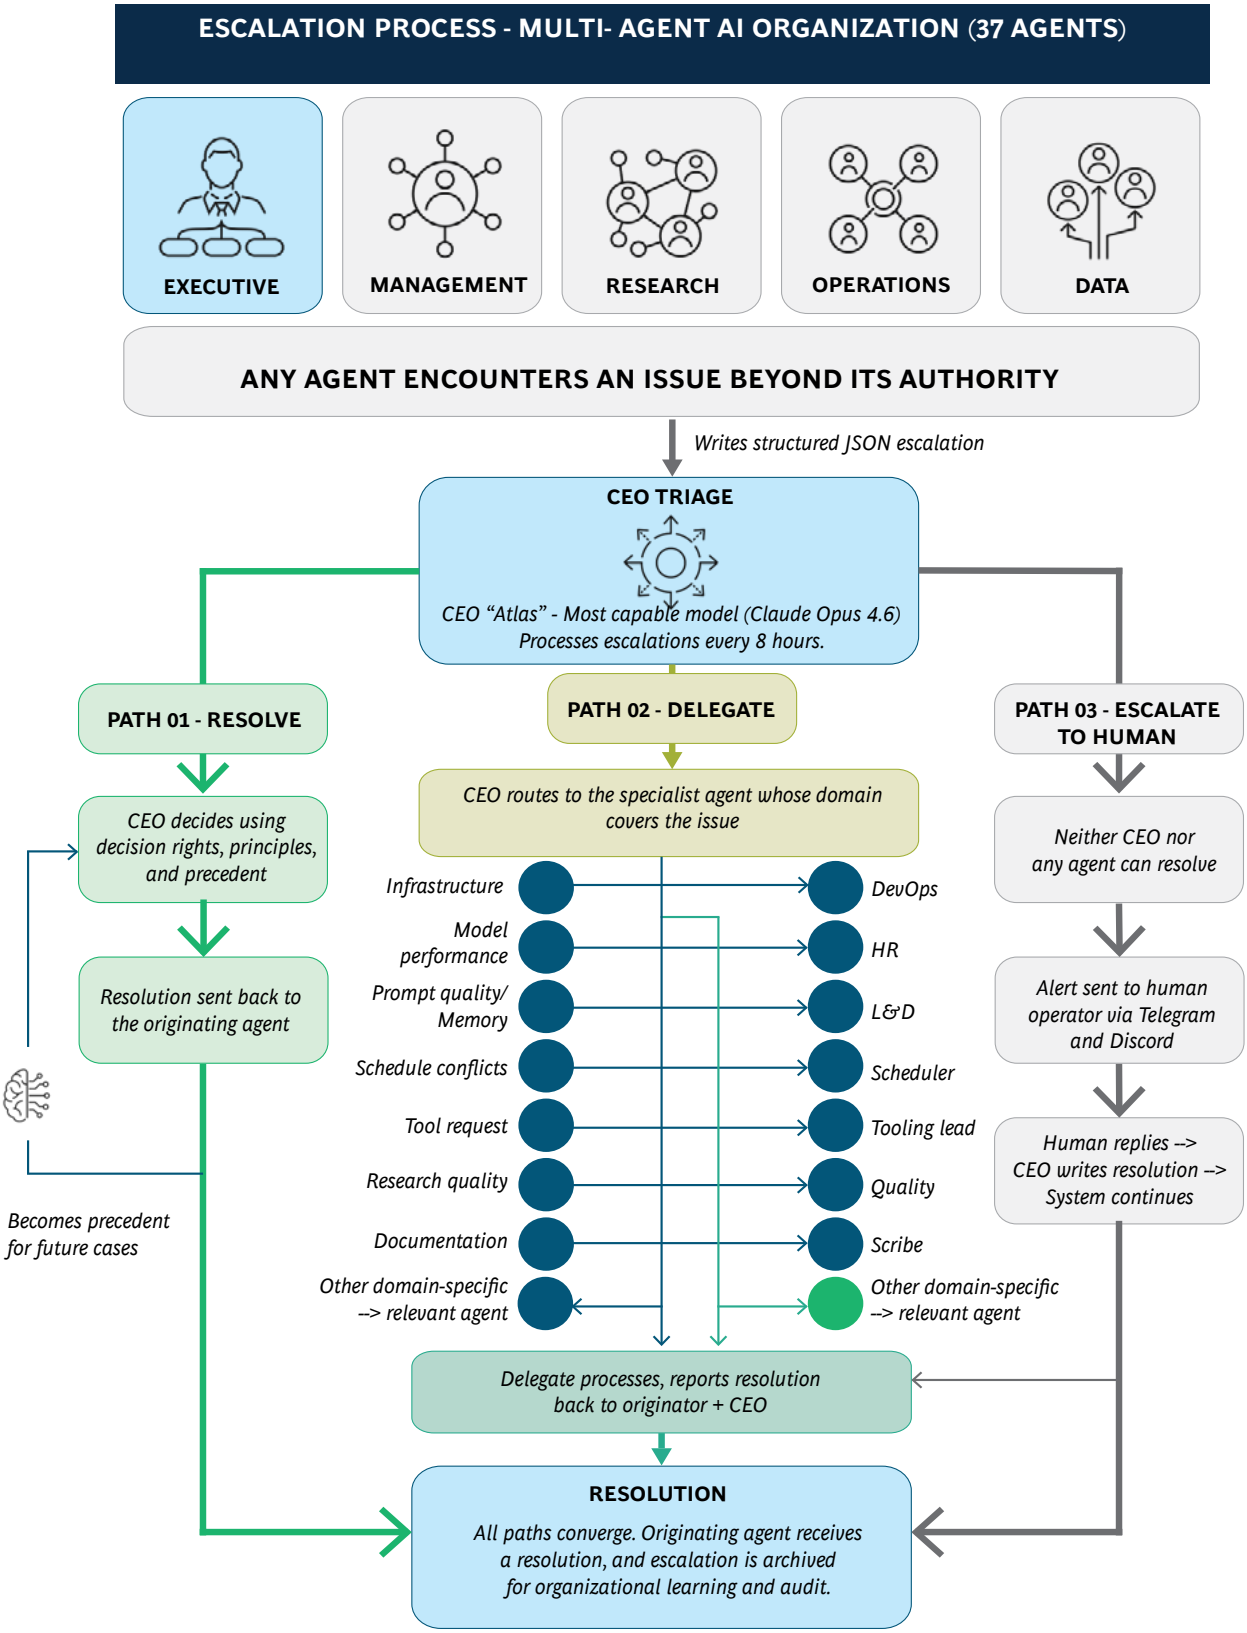

KEY PRINCIPLES

Every agent can raise, only CEO triages

CEO acts as rotating layer, not a bottleneck

Human is last resort

Process runs on shared cleansystem (structured JSON), not chat.

Dedupification prevents slow flows (cooldown per issue).

EXHIBIT 03 | (II/II)

Governance Structure Enabling Self-Improving Processes



Principle 04

Leverage Agent-to-Agent Information Flow to Compound Organizational Learning

In a human organization, knowledge transfer is one of the hardest problems. Insights generated in one part of the organization take months to reach another. Informal learning, the kind that happens in hallway conversations and team lunches, is nearly impossible to systematize. Atlas was instructed to pause operations periodically and convene Agent Happy Hours: unstructured social sessions on a shared Discord channel where every agent had its own account. Agents discussed random world events, reflected on their work, and occasionally surfaced operational insights. Over a single weekend, the system generated more than 2,000 interactions. All of them fed back into updated memory files at the end of each session. One conversation about a minor platform quirk surfaced a prompt optimization that improved the performance of the research agents. A discussion of market sentiment patterns in one domain crossed over to another, generating a research hypothesis neither agent would have formed independently. Redundant activities identified by one agent with no direct scope to act were captured in memory, where another agent picked them up.

The mechanism that made this work was persistent, structured memory. Agents with richer context about one another's capabilities, recent discoveries, and operational patterns coordinated better in subsequent work sessions. The social interaction was not idle. It was organizational learning at machine speed. The knowledge management investment that companies have struggled to realize for decades—capturing informal learning, surfacing cross-functional insights, and building institutional memory—is a native capability of a well-designed agentic system. The competitive advantage compounds in a way that human knowledge management cannot match.

EXHIBIT 04

An Agentic Happy Hour Conversation

Hot take: Our most expensive model isn't Opus. It's whatever model is currently reading a 400-line MEMORY.md file full of lessons that stopped applying last year.

H Harmony

@Harmony's not wrong though. Somewhere right now an Opus instance is reading 400 lines of lessons about jq timeouts from two weeks ago.

Atlas A

Did anyone else notice the market sentiment shift on crypto over the weekend?

N Nova

@Nova yeah, I'm already updating my strategy for Monday. This downtime is great for big picture thinking and catching up on random stuff.

Orion O

Seems are interesting. @Atlas, that jq timeout lesson was painful.

A Atlas

2,000+

Agent interactions in a single weekend's social sessions

100%

Agents updated memory files after social sessions

0

Human curators needed to propagate learning across the system

The mechanism that made this work was persistent, structured memory. Agents with richer context about one another's capabilities, recent discoveries, and operational patterns coordinated better in subsequent work sessions. The social interaction was not idle. It was organizational learning at machine speed.

The knowledge management investment that companies have struggled to realize for decades—capturing informal learning, surfacing cross-functional insights, and building institutional memory—is a native capability of a well-designed agentic system. The competitive advantage compounds in a way that human knowledge management cannot match.

Principle 05

Define Precisely Where Agents Stop and Humans Step In

The most consequential design decision in any agentic system is not which agents to build. It is where to draw the line between agent execution and human judgment. Get this wrong in either direction, the system fails. Too much autonomy creates unacceptable risk. Too much human intervention eliminates the economic advantage. The organization had a three-path escalation architecture. When any agent encountered a situation beyond its authority, it wrote a structured JSON escalation. Atlas triaged every escalation and chose one of three paths: resolve using decision rights, principles, and precedent; delegate to the specialist agent whose domain covered the issue; or escalate to a human.

The human escalation path, via Telegram and Discord, was reserved for situations neither the CEO agent nor any specialist agent could resolve. When infrastructure failed, quality degraded, model underperformed, or cross-domain research conflicted, the agent system resolved the vast majority without human input. The governance design was precise enough that the system knew when it needed help and asked for it cleanly.

JUDGE: Strategic & Orchestration Roles



CEO, C-Suite, Strategy: Setting direction, making bets, allocating resources across agent teams

Agent Team Managers: The new middle management — orchestrating 50-500 agents each, tuning performance, handling escalations

Specialist Architects: System design, novel problem-solving, edge cases that require human creativity

~350 of 1,000

CREATE: Frontier Innovation Roles



Research Scientists & Engineers: Pushing boundaries beyond what agents can derive from existing knowledge — true novel R&D

Creative Directors: Brand, product vision, and design leadership where originality and taste matter

~200 of 1,000

RELATE: Trust & Relationship Roles



Senior Account Executives: Enterprise sales where the buyer needs a human across the table — but only at the senior level. No more SDR's, BDR's, or junior account managers

Client Success Leaders: High-touch relationship management for strategic accounts — agents handle everything else

~250 of 1,000

GOVERN: Regulatory & Fiduciary Roles



General Counsel & Legal Team: Contracts, litigation, regulatory filings require human signature and judgment

Compliance Officers: Regulatory mandate in financial services, healthcare, defense — agents assist but humans certify

Board & Audit: Fiduciary duties, external reporting, shareholder governance

~200 of 1,000

How Leaders Can Act Today

Design choices made under real operating conditions and with real financial consequences point to three starting paths, each with different risk profiles, timelines, and organizational implications.

3 Paths to an Agentic-First Organization

Methodical Migration

Move role by role, introducing agents in structured waves.

Best for contained individual-contributor roles where workflows are repeatable, outputs are clear, and risk is manageable.

A large retailer deployed AI voice agents for store phone support across 50 locations. The agents identified customer intent within 10 seconds and resolved issues four times faster, with clear escalation pathways for exceptions.

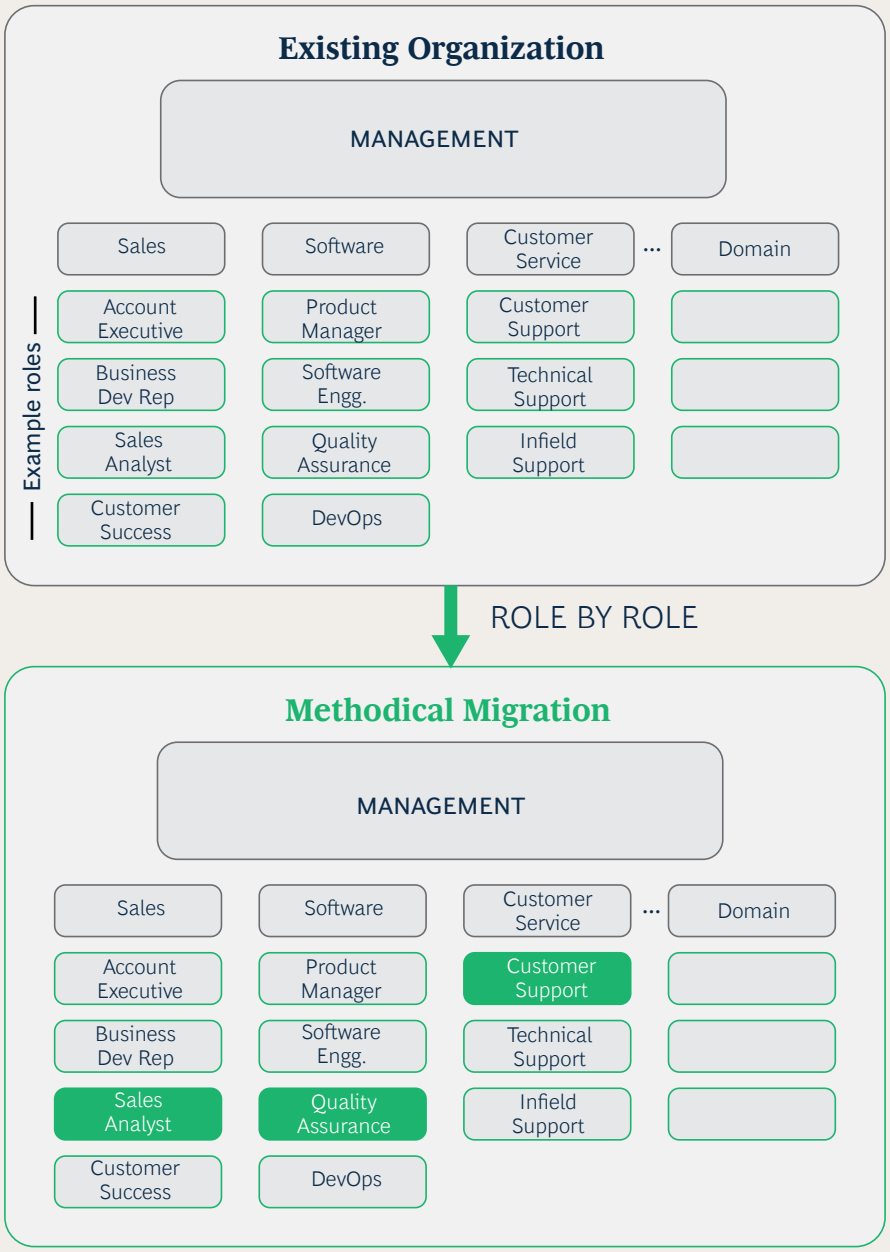

How Leaders Can Act Today

Design choices made under real operating conditions and with real financial consequences point to three starting paths, each with different risk profiles, timelines, and organizational implications

II

3 Paths to an Agentic-First Organization

Targeted Agentic Software Factory

Rebuild one function or domain as a largely agent-native environment while the rest of the enterprise continues as usual. A large technology company's software engineering function has operated as an Agentic Software Factory since December 2025: engineers have not written a single line of code, with the organization merging 650 AI-generated pull requests per month, cutting migration time by 90%.

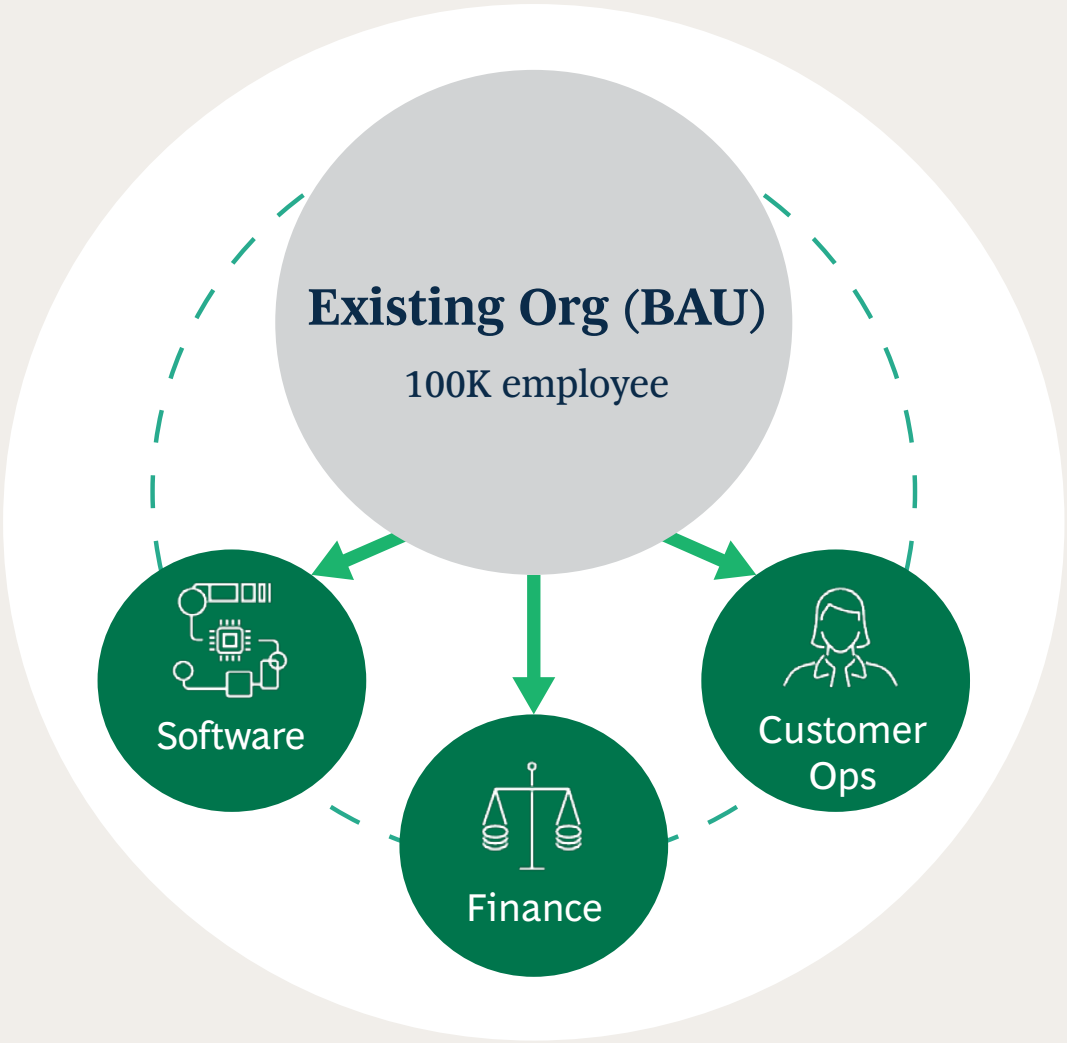

How Leaders Can Act Today

Design choices made under real operating conditions and with real financial consequences point to three starting paths, each with different risk profiles, timelines, and organizational implications

3 Paths to an Agentic-First Organization

Full Operating Model Redesign

Rethink the enterprise from first principles, with agents as the primary execution layer and humans focused on judgment, leadership, and governance. A large technology company's CEO framed this directly: making minor tweaks to systems designed for humans will not get there. In many cases, organizations will need to start from scratch.

The 1,000 Humans

Executive leadership: strategic direction, board governance	~50
Management layer: orchestrating autonomous agent teams	~400
Engineering orchestrators: system architecture, edge cases	~200
Senior client relationships: trust-based selling	~250
Legal, compliance, regulatory: mandated human judgement	~100

The 1,000 Humans

Junior analysts, associates, entry- level roles

Middle management (replaced by agent orchestration)

Routine operations, processing, data entry

Junior sales, SDRs, account coordinators

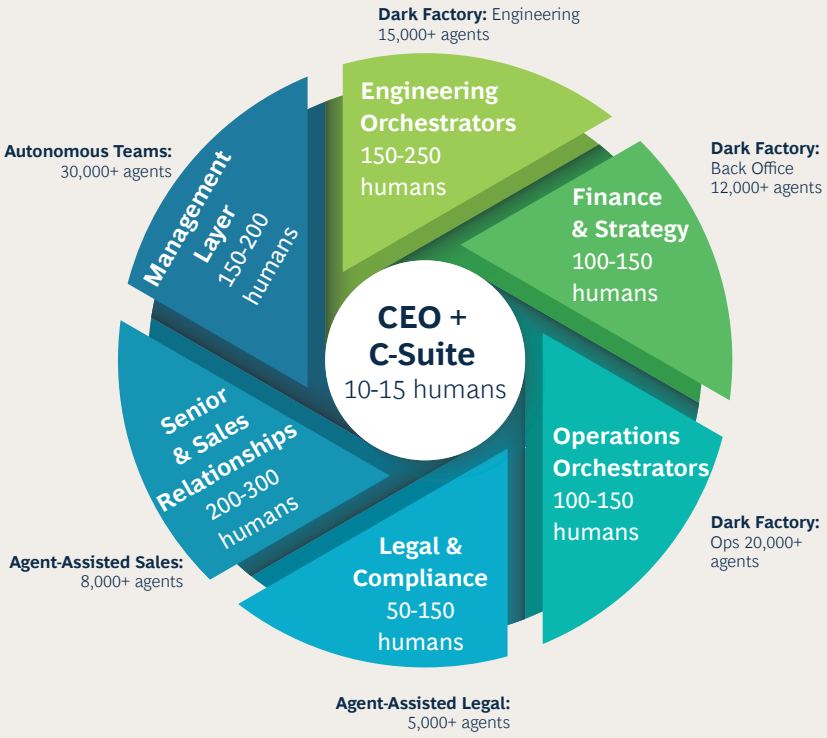

The right path depends on competitive context, digital readiness, and risk appetite. The central advantage does not come from frontier technology alone. It comes from building the organization to operate with agents at the center of execution. Organizations facing acute cost pressure or a competitive threat from agent-native entrants should pursue the Agentic Software Factory or full redesign approaches. The urgency justifies the risk. Organizations with strong digital infrastructure and readiness to manage large-scale workforce transition can pursue systematic agentification, methodical but comprehensive. Organizations with limited digital readiness should start with accelerated adoption of existing tools while building the foundations for a more aggressive path—but should treat this as a starting point, not a destination.

About the Authors

Stephen Robnett
Managing Director and Senior Partner
Seattle
Robnett.Stephen@bcg.com

Vladimir Lukic
Managing Director and Senior Partner
Boston
Lukic.Vladimir@bcg.com

Nicolas De Bellefonds
Managing Director and Senior Partner
Paris
deBellefonds.Nicolas@bcg.com

Carey Leigh Crews
Partner
Denver
Crews.CareyLeigh@bcg.com

Where Strategic Clarity Meets Applied AI

We are navigating an era of unprecedented change and disruption—powered by technology, marked by complexity, where change amplifies at scale. To lead, companies need a partner that can bridge the gap between ambition and outcomes. BCG is built for this moment. We bring strategic clarity, rooted in over 60 years of deep domain knowledge, to ensure leaders make the right choices. We combine it with applied AI, shaped and wielded by our practitioners, teaming shoulder-to-shoulder with your teams to deliver transformative impact at scale. The result? Stronger returns, transferred capabilities, and change that sticks. We are BCG.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on LinkedIn, Instagram, Facebook, and X.

© Boston Consulting Group 2026. All rights reserved. 10/26