EXELIS

EXELIS INSIGHTS

# When AI Agents Join the Operating Model

*Why AI raises the stakes for execution architecture*

The economics of acceleration, distributed authority, and collective behavior in the agentic operating model.

**TODD O. GLASSMAN**

EXELIS GROUP / SEPTEMBER 2026

ARCHITECTS OF ADVANTAGE

THE ARGUMENT IN BRIEF

# Speed compounds the economics of the system.

AI agents bring information, authority and action closer together. That creates an extraordinary operating opportunity—and raises the cost of weak design. Execution architecture must enable adaptation while keeping distributed action economically coherent.

### 01 Acceleration does not repair weak economics.

Just as growth below the cost of capital destroys value, faster execution can compound the losses embedded in an operating model.

### 02 Distributed authority depends on designed boundaries.

Agents can act closer to the work when permissions, shared information, validation and stopping conditions support their authority.

### 03 Local success can produce enterprise failure.

Independent agents may meet their own objectives while their combined actions weaken margins, consume cash or reduce flexibility.

### 04 Simple rules must keep action economically coherent.

Shared economic rules, feedback and timely intervention connect local decisions to enterprise outcomes—and must evolve as conditions change.

> The ability to execute faster is a capability. The ability to make faster execution economically coherent is an advantage.

AI & EXECUTION ARCHITECTURE

# When AI Agents Join the Operating Model

*Why AI raises the stakes for execution architecture*

By Todd O. Glassman

With all the talk about AI agents and what they might do inside a business, I've been thinking a lot about the organization around them. What happens when decisions can move closer to the work, actions can follow immediately, and fewer things need to wait for someone further up the hierarchy? And what does that mean for the way a company creates—or destroys—value?

There is an extraordinary operating opportunity here. Agents could help companies achieve something leaders have pursued for decades: distributed authority, a shorter distance between information and action, and adaptation without routing every decision through the center.

But the same capability could expose weaknesses that organizations have learned to live with. Conflicting priorities, ambiguous authority, and disconnected measures become more consequential when they can generate repeated action before anyone pauses to reconcile them.

To understand that tension, I keep returning to three connected ideas: the economics of growth, the behavior of complex adaptive systems, and the power of simple rules. Together, they offer a useful way to think about what changes when agents become participants in the operating model.

They also bring the importance of execution architecture—the mechanisms that translate strategy into coordinated action—into sharper focus. As agents take on more work, that architecture must enable adaptation while keeping distributed action economically coherent.

The central question is what the organization will do with its new speed.

## The economics come before the acceleration

One of the most useful principles in corporate finance is also one of the easiest to overlook: growth is valuable only when the economics of the additional investment justify it.

Tim Koller and his colleagues at McKinsey have long emphasized the relationship between growth, returns on invested capital, and the cost of capital. When additional investment earns more than its cost, growth creates value. When a company continues investing below that threshold, faster growth compounds value destruction. The principle is central to McKinsey's *Strategy and Value: A Timeless Framework for Enduring Value Creation*.

A company can therefore grow revenue and earnings while making its owners economically worse off.

That principle deserves a prominent place in the discussion of agentic AI. The analogy is straightforward: **speed amplifies the consequences of the system being accelerated.**

If the system directs resources toward attractive outcomes, greater speed can expand the opportunity. If it repeatedly underprices commitments, misallocates resources, or rewards activity that undermines enterprise economics, greater speed can deepen the damage.

Imagine an agent authorized to win business against a revenue target, using pricing rules that fail to capture the full cost of serving the customer. Improving its conversion rate may compound the underlying mistake. Expanding its authority may enlarge the loss.

The relevant performance question is whether its decisions improve the economics of the business after their downstream consequences are included.

Before asking how much faster an agent can execute, leadership needs to establish what successful execution actually means.

## Clockspeed is moving from communication into commitment

The internet accelerated the movement of information and expanded the reach of commerce. Businesses could communicate, coordinate, and transact across distance with far less friction.

Agentic systems create the possibility of compressing more of the sequence between recognizing a need and acting on it. An authorized agent can evaluate options, invoke a service, update a system, or initiate a transaction without waiting for a person at every step.

Programmable payments extend that possibility. For example, the x402 protocol supports automatic stablecoin payments for access to digital services, including purchases initiated by agents. A request for information can therefore include the mechanism to pay for it.

Agents operate through software and interfaces; blockchains can provide a settlement layer for some transactions. Agentic execution does not depend on every action occurring on a blockchain. The important operating development is that finding a service, selecting it, and paying for it can become part of a connected software process.

That brings clockspeed—the pace at which a business senses, decides, acts, and adjusts—into sharper focus.

Physical delivery, customer behavior, and many other constraints will still move at their own pace. But digital commitments can accumulate faster than the operating system beneath them can respond. A company can promise more quickly than it can fulfill, purchase more quickly than it can consume, or change prices more quickly than it can understand the resulting customer behavior.

Koller's principle becomes especially relevant in that interval. Faster commitment to weak economics can increase the cost of discovering the mistake.

## A more explicit form of delegation

The opportunity is equally significant.

Effective delegation requires an objective, relevant information, authority to act, boundaries, and feedback. In a human organization, these elements are often distributed across job descriptions, relationships, experience, and unwritten expectations. The employee must interpret how they fit together.

Agents offer the possibility of making more of that arrangement explicit. Access to information and systems can be specified. Spending and transaction permissions can be bounded. Required checks can be built into the action path. Decisions can be recorded and evaluated against outcomes.

This makes agents a particularly revealing test of execution architecture: leadership must confront what it is actually delegating.

Within a well-defined domain, an agent may be able to act where the information originates, with only exceptions moving upward. That could make decentralized decision-making more practical and reduce the delay between recognizing a condition and responding to it.

However, language-model agents remain capable of mistakes and unpredictable behavior. Instructions alone do not enforce an operating boundary. Permissions, validation, and stopping conditions must be supported by the surrounding software—an engineering distinction reflected in Anthropic's guidance on *Building Effective Agents*.

Human judgment also supplies something essential: the ability to question whether an objective, assumption, or rule still makes sense.

The opportunity is to build a stronger relationship between that judgment and distributed execution. Leaders establish intent and resolve fundamental trade-offs; agents act within domains where authority can be made explicit and outcomes can be observed.

## What an ant colony can teach an operating model

Early in my career, in the late 1990s and early 2000s, the internet was changing how we thought about strategy and business models. Ideas about complexity and adaptation were entering the management conversation through works such as Richard Pascale and his coauthors' *Surfing the Edge of Chaos*. The central insight feels especially relevant today: coordinated behavior can emerge from many participants responding to local information, shared signals, and simple rules. As AI agents join the operating model, leaders have a new opportunity—and responsibility—to design those conditions.

An ant colony or beehive makes the idea easier to see.

Foraging ants can leave pheromone trails that influence which routes other ants follow. Useful routes attract more traffic and reinforcement, while fading signals leave room for the colony to shift. Honeybee colonies also adapt the allocation of work as conditions change. Coordination develops through local responses and shared signals, without a central dispatcher assigning every movement. Eric Bonabeau and Christopher Meyer explore these mechanisms and their business applications in *Swarm Intelligence: A Whole New Way to Think About Business*, published in *Harvard Business Review* in May 2001.

For leaders, the implication is that local autonomy needs a means of producing collective coherence.

A business cannot assume that capable individuals—or capable agents—will automatically arrive at an attractive enterprise outcome. The signals they respond to and the rules connecting their actions matter.

## Digital pheromones make the analogy concrete

Researchers applied ant-inspired behavior to telecommunications routing well before today's generative AI systems.

In one early approach, software agents explored a simulated network and updated routing information that functioned like digital pheromones. Their experience of congestion influenced those updates, helping redirect calls toward more favorable paths as conditions changed. Calls affected congestion; congestion affected the exploring agents; the agents influenced subsequent routing. Ruud Schoonderwoerd and colleagues described this work in *Ants for Load Balancing in Telecommunications Networks*, a 1996 Hewlett-Packard Laboratories technical report.

The significance is the feedback mechanism. Information about current conditions became part of how the system selected its next action.

A company's equivalents could include available capacity, the cost of serving an order, customer service failures, inventory exposure, or commitments already made by other agents.

Applied to an operating model, the analogy suggests that those signals must influence decisions where they occur. If one agent makes a commitment, another needs to see the resulting reduction in available capacity or budget. If a course of action produces poor outcomes, the system needs a way to reduce its continuation.

Otherwise, the organization may generate increasingly sophisticated reports while its agents continue acting on outdated or incompatible assumptions.

## Local success can still produce enterprise failure

Consider an illustrative company with three agents.

A commercial agent adjusts discounts to improve conversion. A procurement agent increases purchasing quantities to secure better unit costs. A scheduling agent fills available capacity to improve utilization.

Each can perform well against its assigned metric. Together, however, they may produce lower margins, higher inventory, and less flexibility to serve the customers the company values most.

The problem can deepen as they respond to one another. Excess inventory prompts further discounting. Increased orders trigger more purchasing. Capacity pressure produces delivery changes or more expensive fulfillment.

No individual agent has to malfunction. The failure can sit in the relationships between their objectives.

This is where the finance and complexity perspectives meet. Economically unattractive decisions can be reinforced through interaction, then repeated at greater speed. The organization becomes more effective at reproducing a result it does not want.

Leadership must therefore evaluate the combined behavior of the system. A collection of improving functional metrics is insufficient evidence that enterprise value is increasing.

**EXHIBIT 1**

### Local success is not enterprise success

Illustrative — shared economic rules change what independent agents optimize.

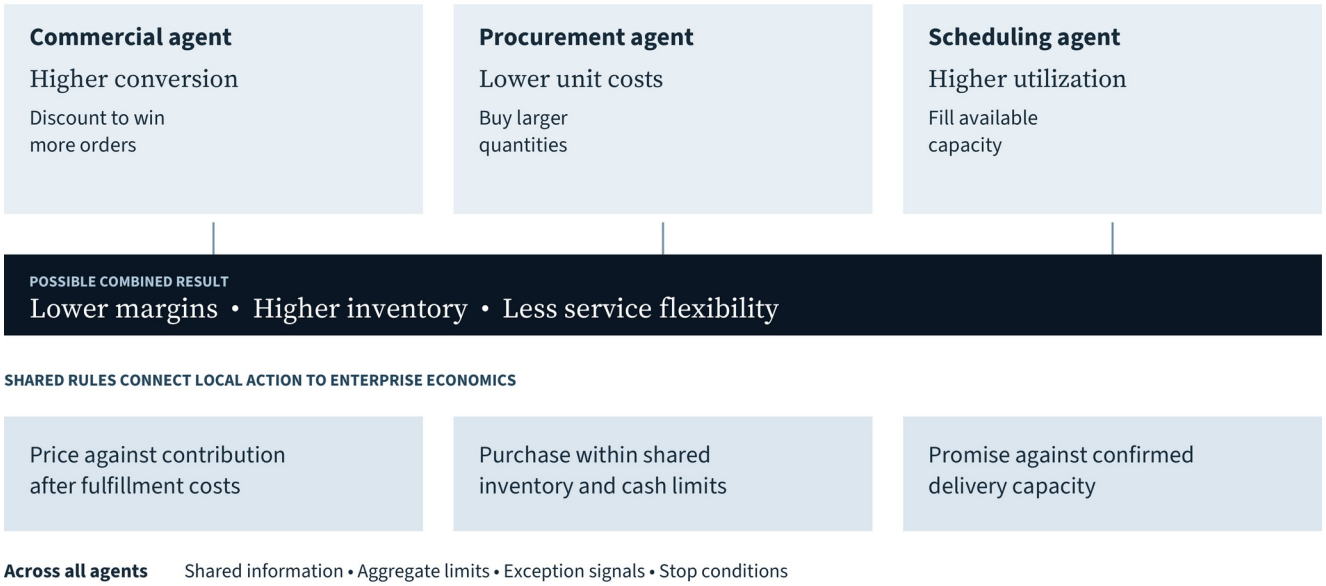


Design for the combined outcome. Test interactions before expanding authority.

*Illustrative framework. Source: Exelis Group.*

# Simple rules make autonomy usable

The response begins with making the critical trade-offs explicit.

Kathleen Eisenhardt and Donald Sull's *Strategy as Simple Rules*, published in *Harvard Business Review* in January 2001, offers a useful foundation: focused rules can guide decentralized action amid complexity.

For an agentic operating model, the value of this approach is the clarity it brings to independent decisions.

In the illustrative company, discount authority could depend on contribution after fulfillment costs. Purchasing could remain flexible within shared inventory and cash limits. Customer commitments could require confirmed capacity. A rise in service failures could suspend further commitments until the cause is understood.

These rules connect local action to enterprise consequences.

They also need to operate cumulatively. A thousand individually permitted purchases can exceed an intended budget. A series of acceptable customer promises can collectively exhaust capacity. Limits must reflect the exposure of the system, with commitments visible across participants.

Crucially, the rules themselves require review. A rule that once protected value may become restrictive or misleading as demand, costs, and capabilities change.

Execution architecture must therefore enable adaptation while keeping distributed action economically coherent. It establishes what must remain protected, where discretion is useful, and how experience changes subsequent behavior.

# The architecture must keep pace with the action

As delegation expands, management's traditional review intervals may cease to match the speed at which consequences accumulate.

The practical question is how much the company can commit before it detects an adverse pattern and responds.

A bounded analytical task may need periodic review. Authority to spend, price, purchase, or promise may require immediate exception signals, aggregate limits, and mechanisms that suspend action when conditions deteriorate.

Leadership remains accountable for the system. That accountability must be supported by visibility and intervention that operate within the available time.

This also makes gradual expansion of autonomy economically sensible. A company can observe how agents behave together, test whether limits hold, and establish that corrective action works before extending authority. The evidence should include outcomes across the workflow, not only successful completion of individual tasks.

The aim is to increase the speed at which the organization can act coherently—and learn when its assumptions are wrong.

# What leadership is really accelerating

AI agents could bring companies closer to an operating aspiration that has often proved difficult to sustain: decisions made near the work, within clear authority, informed by current conditions, and connected to enterprise priorities.

That potential makes execution architecture more valuable. It also makes weak architecture more expensive.

The enduring lesson from corporate finance is that expansion compounds the economics beneath it. The lesson from adaptive systems is that collective outcomes depend on interactions and feedback. Simple rules help connect those insights to the decisions a business makes every day.

All that said, this is still an emerging area. I do not mean to suggest that the design choices are straightforward, or that a few well-written rules will resolve the complexity. Defining appropriate authority, reconciling competing objectives, and understanding how agents behave together will require experimentation, judgment, and revision. Some of the most consequential interactions may become visible only through carefully bounded use.

What these ideas offer is a frame of reference as agents take on a greater share of the work inside companies. Leaders need to be thoughtful about where they delegate, what they measure, how they learn, and how quickly they expand authority. The architecture will need to evolve alongside the capability.

**The ability to execute faster is a capability. The ability to make faster execution economically coherent is an advantage.**

Before increasing the clockspeed of the business, leaders need to understand what their architecture will cause it to do—and what will happen when those decisions begin reinforcing one another.

That is what they are about to accelerate.

ABOUT THE AUTHOR

## Todd O. Glassman

Todd is the Founder and Principal of Exelis Group. Over the past 25+ years, he has worked across executive leadership, investing and advisory roles, helping owners, investors, boards and management teams build, scale, transform and institutionalize businesses, capabilities and strategic initiatives. His work sits at the intersection of strategy, finance, operations, investing and execution.

ABOUT EXELIS GROUP

## Architects of Advantage

Exelis Group is a principal-led C-suite advisory and interim leadership firm helping owners, investors, boards and management teams improve performance, navigate critical inflection points and build more durable enterprise value. We combine investor perspective with operator accountability, translating strategic intent into the mechanisms, capabilities and management systems required for durable execution.

START A PRIVATE CONVERSATION

info@exelis.co
www.exelis.co