



Executive  
Perspectives

# From Copilots to Enterprise Transformation: AI's New Mandate in Supply Chains

Supply Chain

*May 2026*

# Introduction

The story for AI and GenAI in supply chains in 2026 is no longer about potential but instead execution.

Over the past 18 months since the previous BCG Executive Perspective on supply chains, AI capabilities have evolved materially from copilots and chatbots to systems that can reason and execute multistep workflows. At the same time, supply chains have increasingly become a board-level topic, with higher geopolitical volatility, margin pressure, and a need for sharper working-capital discipline.

Yet despite investment, many organizations are struggling to translate AI into impact:

- Less than 25% of companies demonstrate AI maturity at scale in supply chain
- Most remain focused on master data, forecasting, and limited automation
- Only about 30% report measurable AI value in planning use cases

This creates a new mandate for chief supply chain officers and COOs. In the short term, it requires leveraging AI to improve high-impact, high-frequency decisions that are constrained by analytical capacity and to reorganize workflows to break down organizational silos.


In the medium term, it requires laying the foundation for a human-agentic supply chain operating model; one where agents autonomously orchestrate and balance revenue, cost, and service, while human teams manage exceptions and continuously improve the system.

This BCG Executive Perspective provides a guide for CEOs and operations leaders on how to:

- 1 Focus on key decisions to move beyond LLM-on-top experimentation
- 2 Set up an agentic supply chain platform with the right foundations to scale and evolve
- 3 Fund the journey through productivity in parallel to elevating decision quality
- 4 Incrementally build toward an agent-enabled supply chain of the future

Note: LLM = large language model.

**In this BCG  
Executive Perspective,  
we provide an  
updated vision of the  
future of supply chain  
with AI**



# Executive summary | Agentic AI is positioned to drive near-term top- and bottom-line impact by transforming supply chain decisions

Supply chains have realized little value after adopting simple GenAI apps

- AI capabilities have shifted materially over the past 18 months, from conversational models to **full systems that can reason, use tools, and execute multistep workflows** with limited input
- However, in supply chains, **most deployments are still stuck using copilots and focused automation**, with only 30% of companies reporting measurable AI value even as investments scale and expectations from boards and investors accelerate

Agentic AI has the potential to break traditional supply chain constraints

- AI agents can eliminate the historic tradeoff between effectiveness and efficiency; supply chains can now **simultaneously make faster, better decisions at scale**, something human-led planning could not consistently deliver
- Leaders are seeing **AI expand the performance envelope** beyond what any team could reach before with more frequent, more granular optimization that is unlocking new revenue, cost levels, and working capital
- The question is no longer whether AI can expand the performance envelope, but **how to set up an organization** to capture the value from it and move faster than the competition

A human-agentic operating model is emerging as the future framework for supply chain

- In 2030, we expect SCM organizations to look very different, with a **human-agentic operating model in place**, where AI autonomously runs many processes, and a small team of experts handles critical decisions and continuous system improvements
- To get there is a journey of incrementally **upgrading and connecting different parts of the supply chain** through an evolution from task-specific enhancements to process transformations and, finally, automation
- Leaders must begin redesigning workflows and decision rights now, as the organizations that move first can not just improve their KPIs but also **operate with fundamentally different economics vs. those who don't**

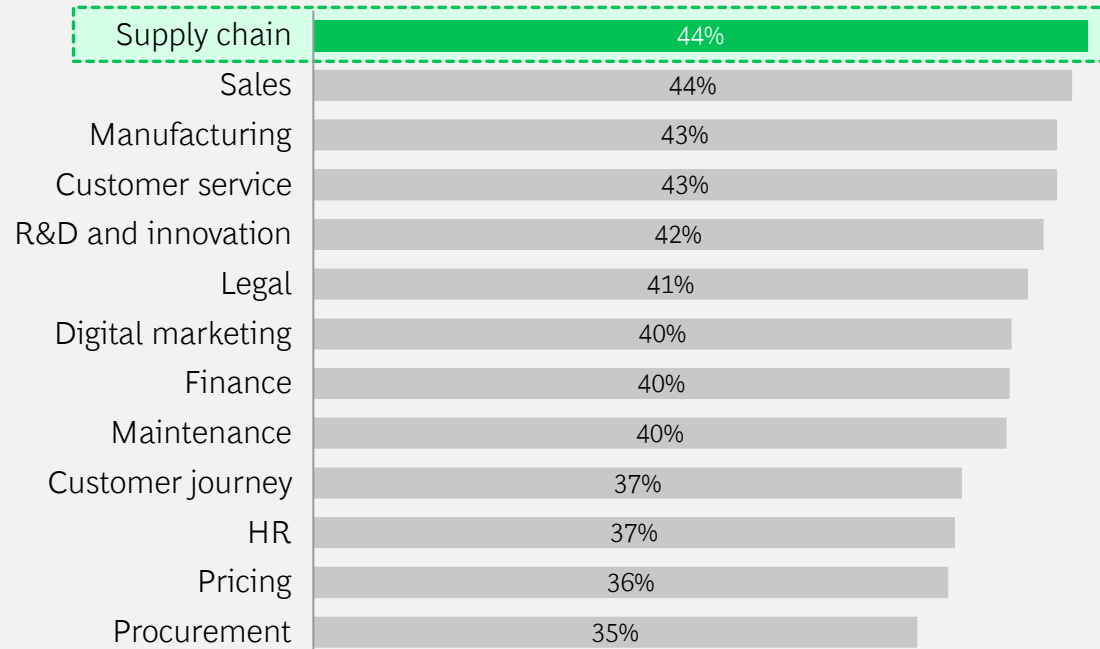
Today, leaders are making five strategic moves to ensure future success

1. **Invest in a robust data foundation**
2. **Start where decision density and value intersect**
3. **Rebuild workflows around AI-led enterprise optimization vs. functional negotiations**
4. **Adopt a hybrid build-and-buy approach to platform strategies**
5. **Make AI decisions transparent, auditable, and explainable**

# AI adoption in supply chains is widespread, but many companies are not finding value

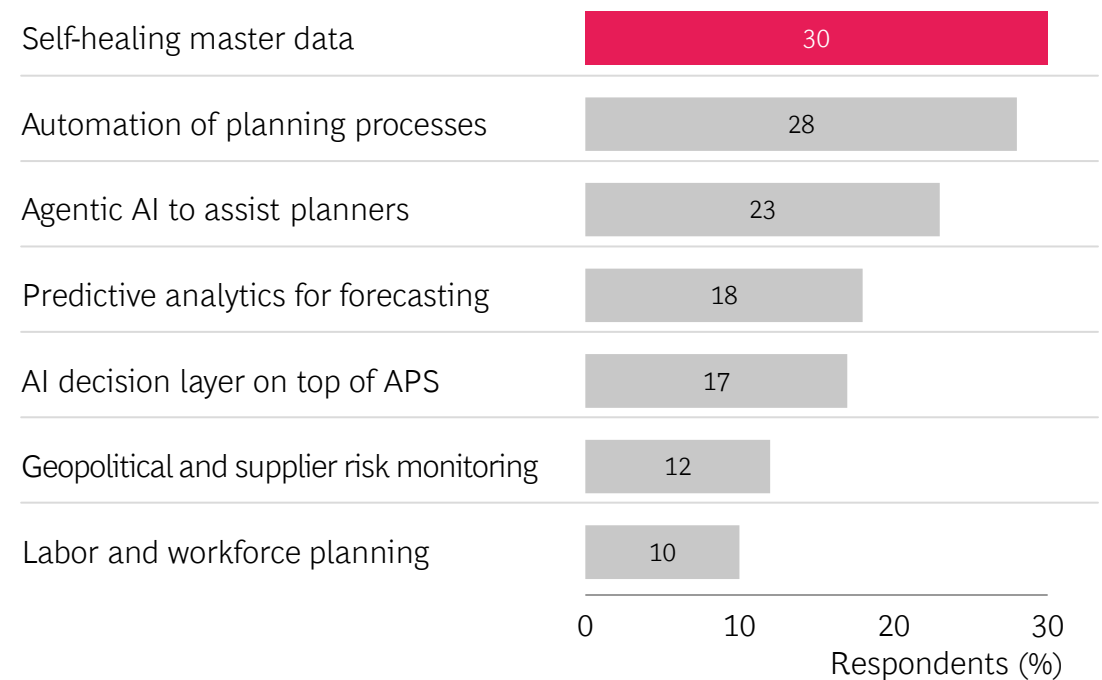
There is significant adoption of AI in supply chains (~44%)

## Adoption rate of AI across functions



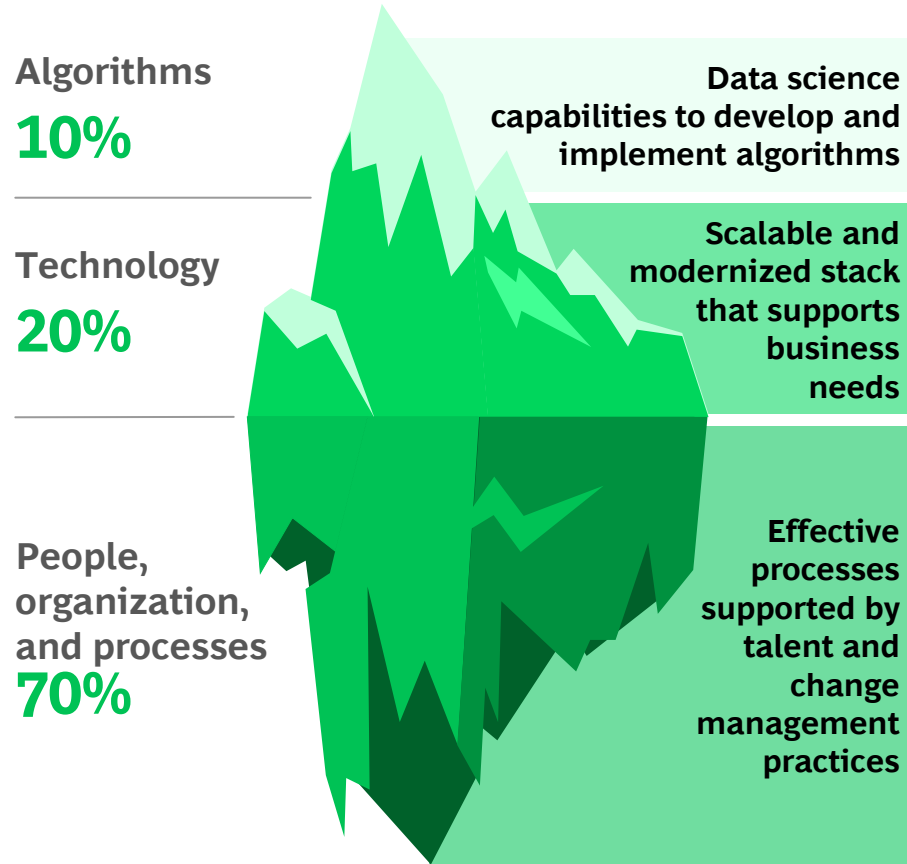
Yet many companies are seeing limited value, with the biggest gains in master data (~30%)

## Degree of impact of AI in supply chain planning use cases for the organization

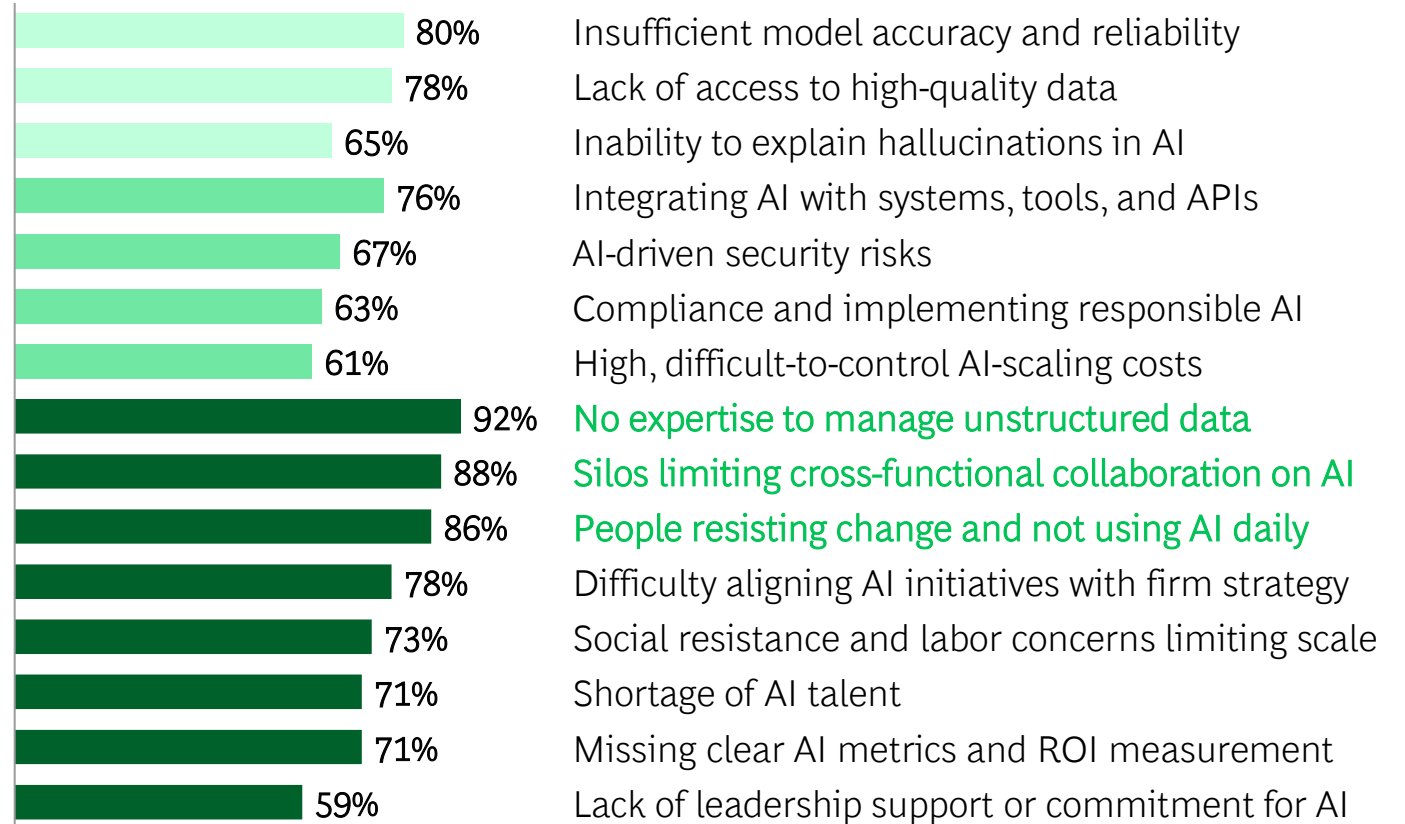


Note: APS = advanced planning system. Respondents were asked to select as many options as were applicable to them. Percentages in the bar chart were rounded to the nearest whole number. Sources: BCG Build for the Future 2025 Global Study (merged with BCG's Digital Acceleration Index), n = 1,250; BCG Annual State of Supply Chain Planning Survey 2025 (n = 181); BCG analysis.

# Even with AI, a supply chain transformation faces the same fundamental challenges of adoption, functional silos, and having the right expertise



## Key challenges<sup>1</sup> (respondents, %)<sup>2</sup>



1. Which of these challenges hinder adoption and scale of AI in your company? ; 2. Percentage of respondents who agree and strongly agree.

Note: API = application programming interface; ROI = return on investment.


Source: BCG Build for the Future 2025 Global Study (n = 1,250).

# What is an agent?



**An agent is AI that uses tools to accomplish goals**

What an agent **IS**



## Memory

Remembers across tasks and changing states



## AI models

Uses one or more AI models, usually an LLM or SLM



## Systems

Accesses external systems on your behalf

What an agent **DOES**



## Observes

Collects and processes data from environment



## Plans

Evaluates possible actions and prioritizes toward a goal



## Acts

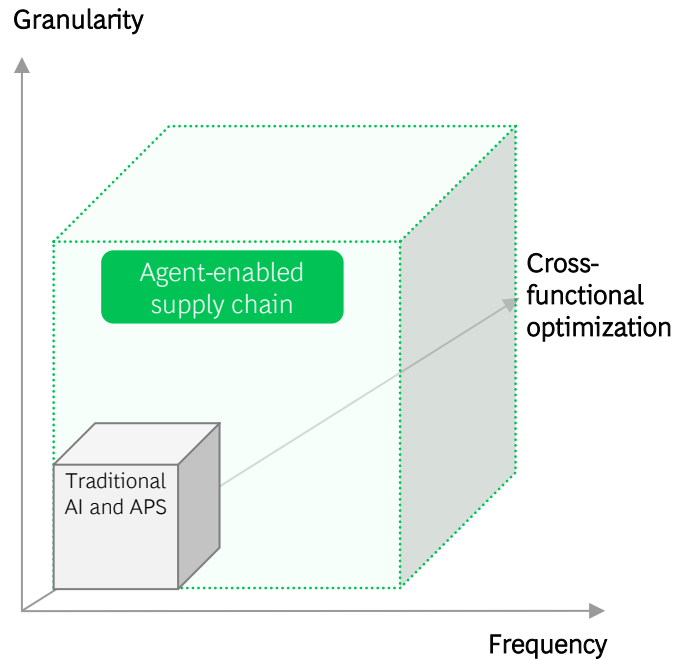
Executes by leveraging internal or external tools and systems

Note: SLM = small language model.

# Agents are expanding the performance envelope in ways that advanced planning systems and human-led processes could not

## Agents are removing the constraints that traditional AI and tools could not overcome

### Supply chain decision space



### Agents enable high decision density

#### 1 High frequency

Always-on execution that is not constrained by planning cycles or analyst availability

#### 2 High granularity

No longer forced to aggregate data to make problems manageable for humans

#### 3 Integrated cross-functional optimization

One pass across price, service, cost, and risk versus sequential team negotiations

## Agentic AI has the potential to unlock massive value creation for supply chains



Revenue upside

**+2%-5%**

Revenue uplift



EBITDA

**+2-4 pp**

Profitability increase



Working capital

**15%-30%**

Inventory reduction



Service and satisfaction

**+5-15pp**

Service rate improvement



Throughput

**+5-10 pp**

OEE uplift



Costs

**10%-20%**

Reduction in manufacturing, warehousing, and distribution costs



CO<sub>2</sub> emissions

**20%-50%**

Average near-term CO<sub>2</sub> reduction



Resilience

**Divide by 10**

Time to understand upstream scenarios and actions needed vs. suppliers



Flexibility

**Divide by 5**

Time needed to make plan and execute

Note: OEE = overall equipment effectiveness.  
Source: BCG supply chain case experience.

# An agentic-enabled supply chain optimizes inventory daily and at the most granular level, reducing out-of-stocks and increasing sales



## Inventory optimization reimaged in an agentic-enabled supply chain

From

Weekly inventory allocation based on DOH targets by product group

Target inventory (all locations):

- A SKUs: 10 DOH
- B SKUs: 15 DOH
- C SKUs: 20 DOH



To

Daily inventory optimization based on SKU economics and forecast confidence

SKU 1, location 1:  
Optimal inventory  
14 units

SKU 2, location 2:  
Optimal inventory  
18 units

Illustrative



Note: DOH = days on hand. DOH is equivalent to average daily sales times number of days.  
Source: BCG case experience and benchmarks of proprietary and nonproprietary AI forecasting tools.



## Impact



**Higher revenue** from better alignment between short-term forecast and stock deployment



**Lower out-of-stocks** via de-averaged, granular replenishment logic



**Inventory reduction** through targeted turn-rate improvements enabled by de-averaged analytics



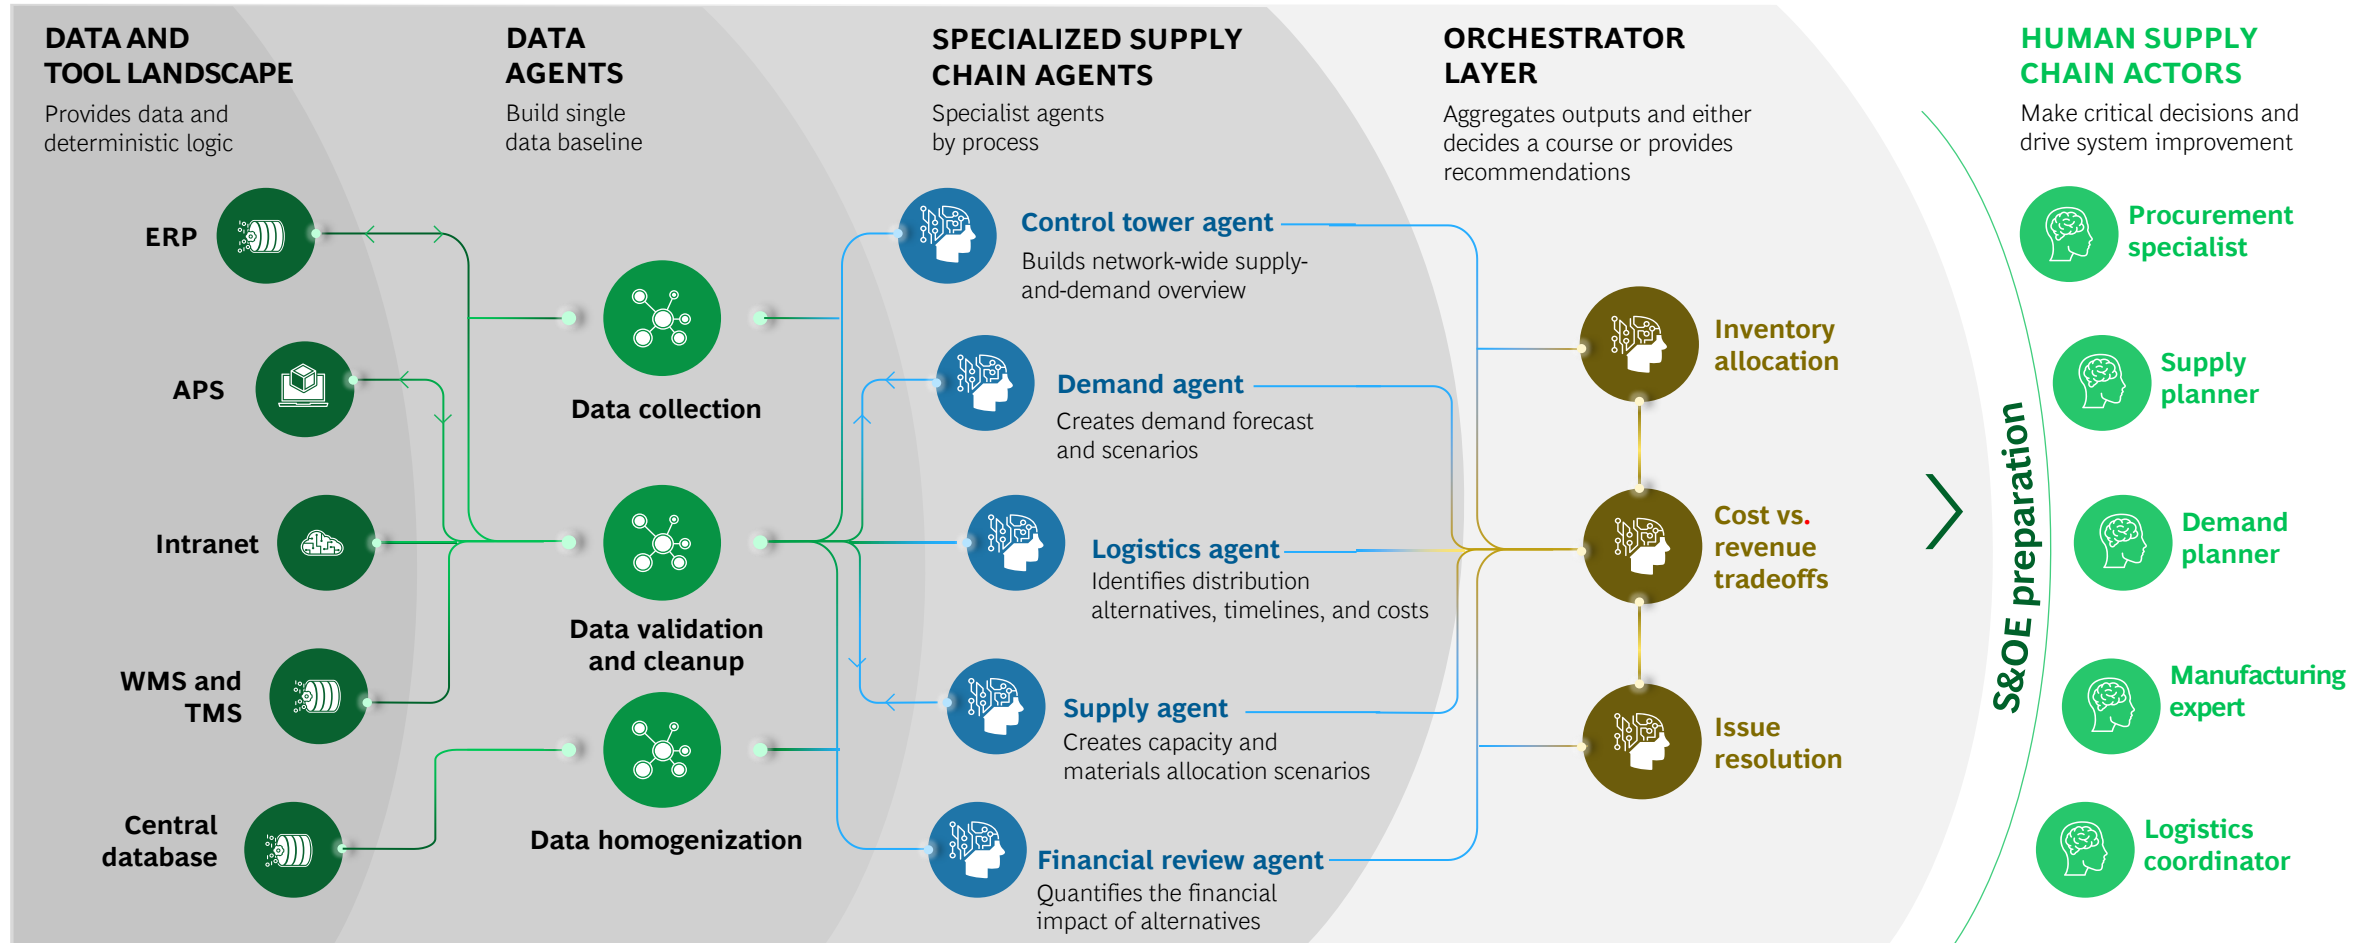
**Lower logistics costs** by avoiding unnecessary replenishments and transfers

# A human-agentic operating model is emerging as the future supply chain framework

Example

## Advanced E2E planning workflow vision

Meta-agent orchestrates subtasks and coordinates data sources and tools



Note: E2E = end to end; ERP = enterprise resource planning; WMS = warehouse management system; TMS = transportation management system; S&OE = sales and operations execution.

# By 2030, SCM organizations are expected to have agents planning and executing and humans managing exceptions

## Planning activity types

### Design: Improve the system

**Continuously improve** planning rules, thresholds, and systems

### Resolve: Manage exceptions

**Prioritize and resolve exceptions** through cross-functional tradeoffs


### Run: Plan and execute

**Execute plans** through routines and rule-based activities, including manual analysis

Limited digital enablement;  
focus on systems

Today

Most effort is spent  
running the machine



Planners keep plans  
alive with little time to  
improve them

**Design: 5%**


**Resolve: 25%**

**Run: 70%**

Early agent adoption; high  
human oversight

2030

Agents plan and execute;  
humans focus on resolving



Automation reduces  
noise; planners shift to  
problem-solving

**Design: 20%**


**Resolve: 50%**

**Run: 30%**

Scaled adoption; increasingly  
autonomous agents

2035

Design-led, exception-  
based planning



Planning becomes a  
strategic capability focused  
on shaping outcomes

**Design: 60%**

**Resolve: 35%**

**Run: 5%**

# A human-agentic supply chain shifts from siloed optimization and firefighting to a predictive, unified view that facilitates enterprise decisions



## Maria, vice president of supply chain

It's Monday morning and a key supplier has just missed a critical delivery. Maria is now facing a 20% shortfall on her highest-margin product line and needs an answer as soon as possible on where to go next.



### Maria's week **without** agents

- **Monday: Commercial scrambles**  
Team finds out late; customer commitments already at risk, no time to reprioritize
- **Wednesday: Operations reallocates blindly**  
Capacity shifted on cost alone, blind to downstream revenue impact
- **Next week: Finance takes stock of the impact**  
Forecast revised downward only after decisions are already made

Maria gets a **suboptimal answer** after spending **two to three weeks firefighting**



### Maria's Monday **with** agents

- **Overnight: Agent flags the shortfall**  
Models customer impact by segment, flags commitments to protect based on margin and relationship value
- **6:00 AM: Agent evaluates all options**  
Alternative supply, production resequencing, and partial shipments all assessed with revenue and service impact, not just cost
- **7:00 AM: P&L tradeoffs quantified**  
Full scenario comparison with recommendation, data sources, and step-by-step rationale

Maria has an **agent-orchestrated response** that balances revenue, cost, and service **before the first meeting is scheduled**

# Agentic planning evolution | A CPG company redesigned replenishment to harmonize cross-functional processes and proactively make tradeoffs

## Context

Global CPG leader is facing volume and service pressure from top retailers

**Goal:** Strengthen E2E planning capabilities

### Key challenges:

- Inconsistent customer data
- Reactive analytics approach
- Fulfilment seen as execution, not strategic enabler
- Limited team capabilities
- Team spends most time firefighting
- Limited cross-functional input

## Execution

### Shifted replenishment from reactive to proactive

Developed **GenAI-enabled agents** to deliver in-stock alerts and **smart replenishment recommendations** (e.g., DC-to-store transfers and expedited orders)

Validated recommendations via **human-in-the-loop feedback**

Built modular agent back end for future expansion

Implemented an **enhanced collaboration workflow** with retailers facilitated by agentic AI

## Business value captured

### Impact from GenAI agent

**+2-4%**

In-stock percentage driving revenue lift

**+4-10%**

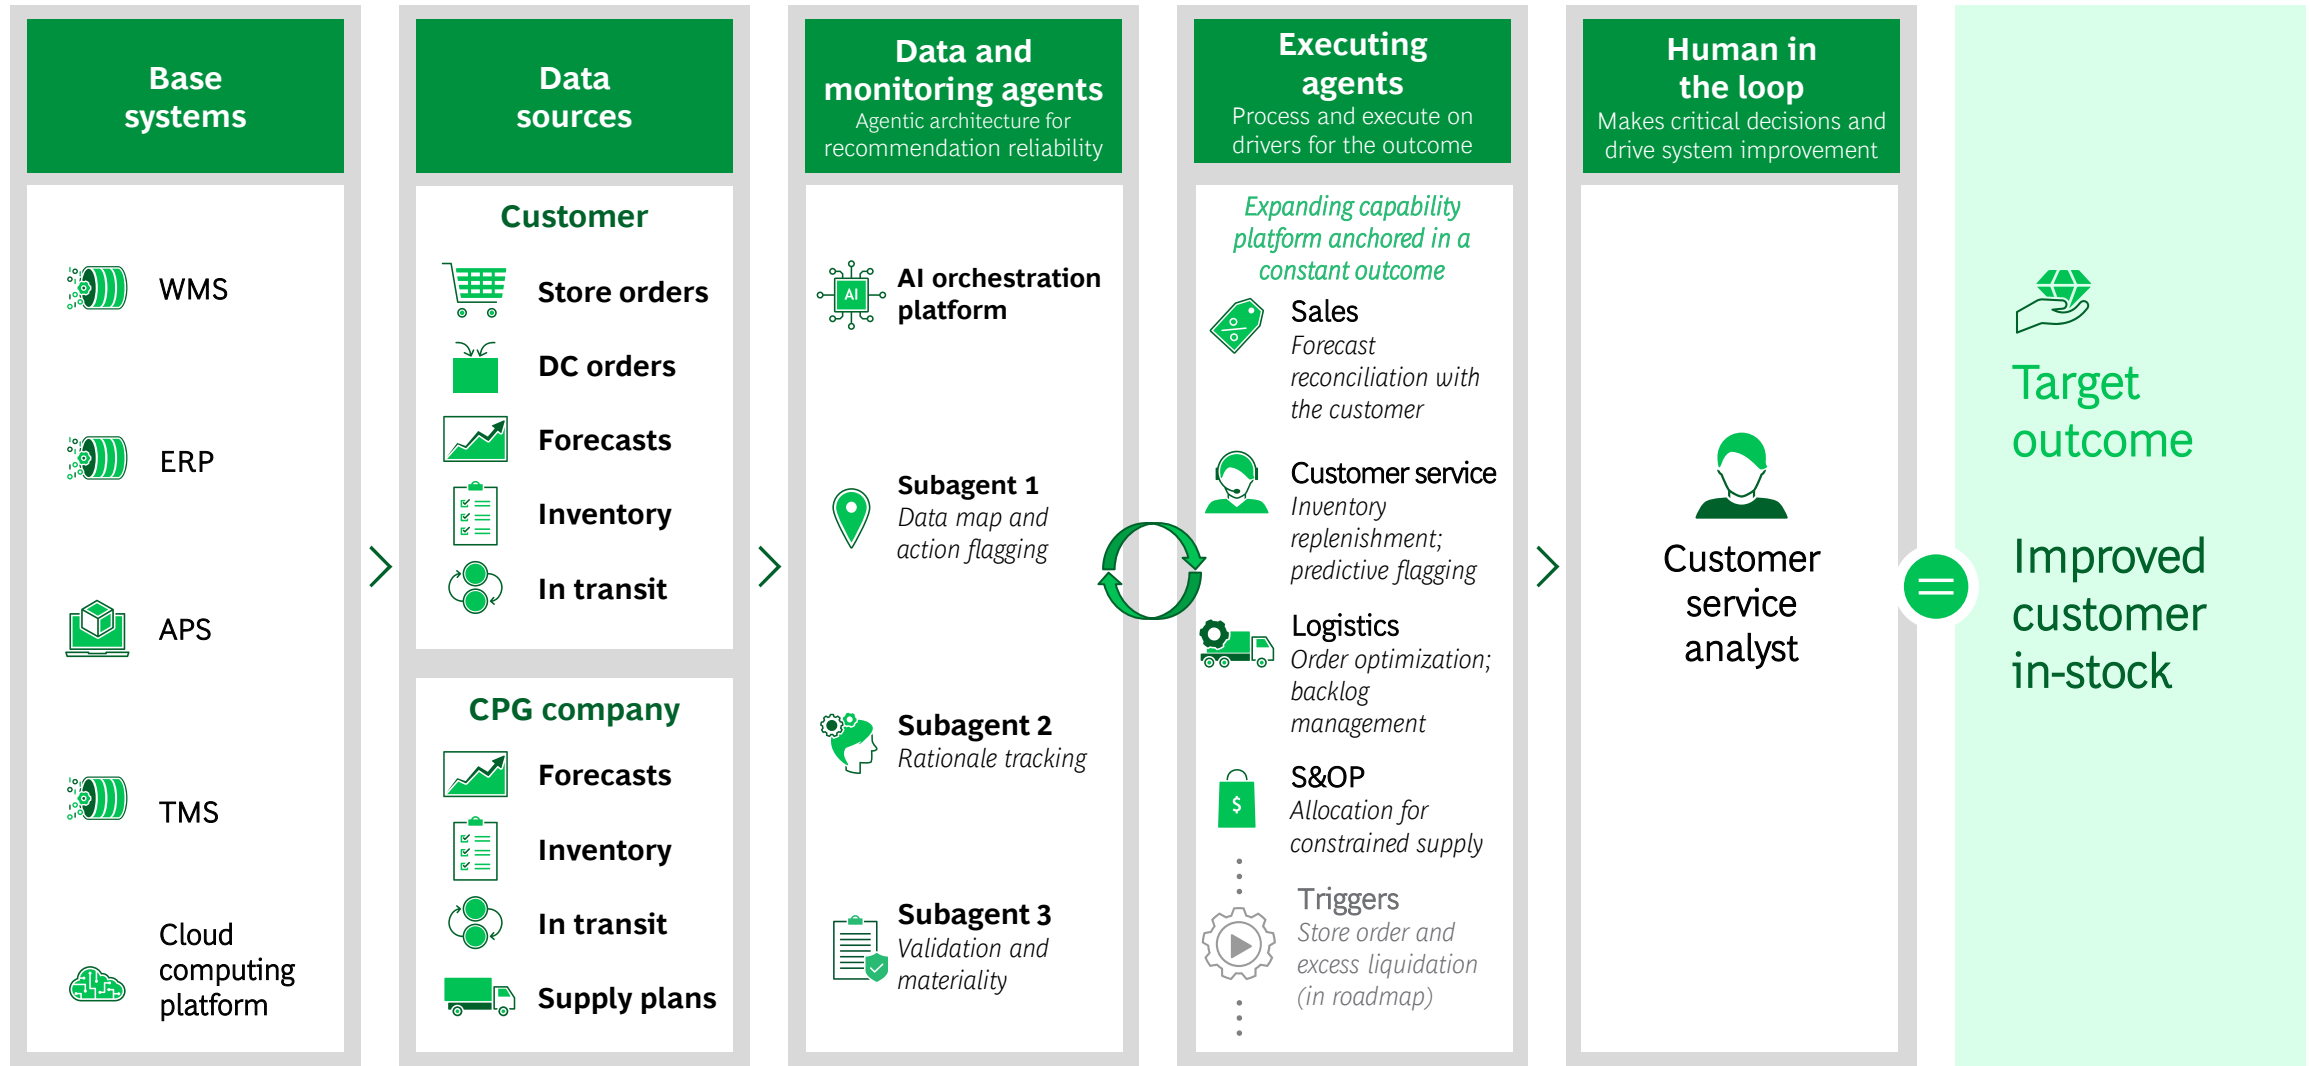
Fill rate

**40-60%**

Administrative savings

Date	Warehouse	SKU	Action	Rationale	Priority	Status
Mar 21, 05:00 PM	O20	963345570	Req_excessive_inventory	Overstock: 71 DOCS (optimal 9 DOCS). Check promotion status for "Tal End SKU."	Medium	new
Mar 23, 05:00 PM	S06	4254199919	Req_excessive_inventory	108 cases overstock, 33 days supply vs 10 optimal. Tal End SKU, low 4.71 daily demand.	Medium	new
Mar 23, 05:00 PM	F17	4254199919	Req_potential_order_delay	Critically low SOH (2 cases), pending order for 40 cases delayed 14 days.	Medium	new
Mar 23, 05:00 PM	S06	7572886071	Req_excessive_inventory	Excess 300 cases (22 DOCS over target); systemic supply (207%) vastly exceeds demand (15%).	Medium	new
Mar 23, 05:00 PM	A39	5204372292	Req_constrained	No resupply planned for warehouse or distributor DC. SKU on promotion.	Medium	new
Mar 23, 05:00 PM	S06	5204372292	Req_constrained	SKU not constrained. DGS 78.91 well above optimal/max. Healthy stock, no immediate supply concerns.	Medium	new
Mar 21, 05:00 PM	O09	5204372292	Req_constrained	Capacity is 0.55, only 56% demand met. Persistent low fulfillment despite 24.6 days lead supply.	Medium	new
Mar 23, 05:00 PM	F17	363057044	Req_potential_order_delay	Order delayed. Warehouse 0 cases, DC 12 cases; insufficient for order. 8350 cases in network.	Medium	new
Mar 23, 05:00 PM	A39	5476848836	Req_constrained	Inventory DGS (14.98) is above optimal (10.55), not constrained, no stocks.	Medium	new
Mar 23, 05:00 PM	S06	5476848836	Req_constrained	40 Days of Supply (near max recommended). Total supply significantly exceeds total demand.	Medium	new

# Agentic planning evolution | The company's agentic platform provides validated recommendations with human analysts making critical decisions



Note: S&OP = sales and operations planning.

# Agentic planning evolution | The journey consisted of funding AI builds through productivity gains while elevating decision quality

## Journey to make agentic planning operational


### Pillar 1 – Launch pilots and capture productivity

1. Agent process capability increases
2. Head count productivity improves from automating operational tasks

### Pillar 2 – Redesign work to leverage smarter decisions

1. Agent reliability grows, building trust and autonomy
2. Focus on higher-quality decision making and competitive advantages

### Management priority



Fund AI builds with more than 50% FTE productivity savings



Extract value at scale with consistent execution

Smooth transition, with **humans retaining majority control through pilot phase**

Agents start taking direct action in the system via **continuous tuning and learning**

Pilots serve as a **template for an organization-wide model shift**

# GenAI planning | An industrial goods firm reshaped supply chain processes with GenAI tools running complex scenarios and root-cause analysis

## Context

Leading Europe-based industrial goods company was making **thousands of supply chain decisions daily**

Goal: Supercharge existing supply chain simulation capability

Aimed to **optimize operations and sharpen decision making** by:

- Identifying bottlenecks
- Testing different strategies
- Running complex scenarios

## Execution

Integrated two BCG X assets via a natural language interface:

- AgentKit: **GenAI agent toolkit** (open-source code framework)
- E2E Plan by BCG X: **Comprehensive planning suite**

Live solution streamlines S&OP and enables planners to independently:

- Create simulation scenarios
- Analyze root causes
- Summarize KPIs
- Run sensitivity analyses
- Share key simulations outputs

## Business value captured

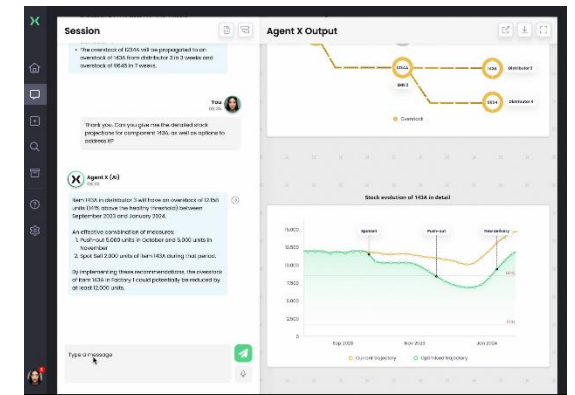
**Impact from underlying AI capability**

**+2pp** | EBITDA increase in year two and after

**Impact from GenAI agent**

**+25** | Planning professionals trained

**3x** | Process cycle time reduction



# Smart inventory management | A specialty materials leader reshaped its spare parts strategy by using AI's comparison and substitution capabilities

## Context

Global specialty materials leader has eight manufacturing sites across the US and Europe

Goal: Reduce MRO spending without impacting service

Key challenges:

- Excessive MRO inventory
- Fragmented sourcing and local spare-parts decisions
- Inconsistent MRO master data; no optimization tools

## Execution

AI embedded throughout sourcing to enable smarter purchasing

- Developed AI algorithm to identify part similarity using unstructured data and catalogues
- **Embedded into sourcing: checks for similar existing parts before new purchases; compares manufacturers for best price**

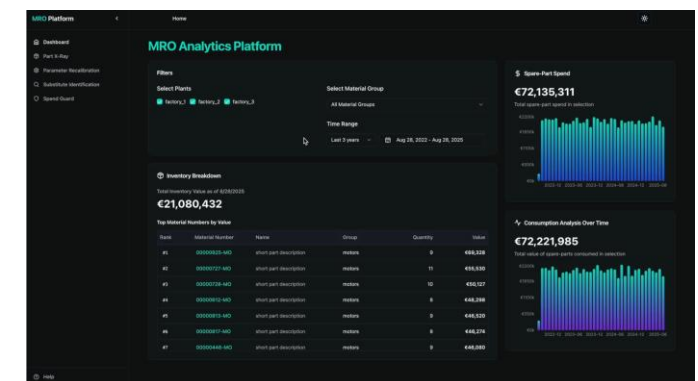
Recalibrated stock parameters in parallel to optimize inventory

## Business value captured

Impact from AI and advanced analytics

**+2%-5%** | Savings in annual spare-parts sourcing spending

**15%-20%** | Inventory reduction with minimal impact to service



# Five moves that leading supply chains are making now to embed AI and shift from functional optimization to enterprise decision orchestration

1

## **Invest in a robust data foundation**

Agents require clean, connected data; a modern cloud-native platform delivering trusted data at business speed is essential

2

## **Start where decision density and value intersect**

Prioritize planning and/or trade negotiations—areas with frequent tradeoffs and clear financial impact—rather than spreading AI thin

3

## **Rebuild workflows around AI-led enterprise optimization vs. functional negotiations**

AI generates an E2E optimized plan as the starting point, freeing teams to focus on structural decisions, not tactical tradeoffs

4

## **Adopt a hybrid build-and-buy approach to platform strategies**

Balance strategic value, customization needs, long-term cost, and pace of innovation

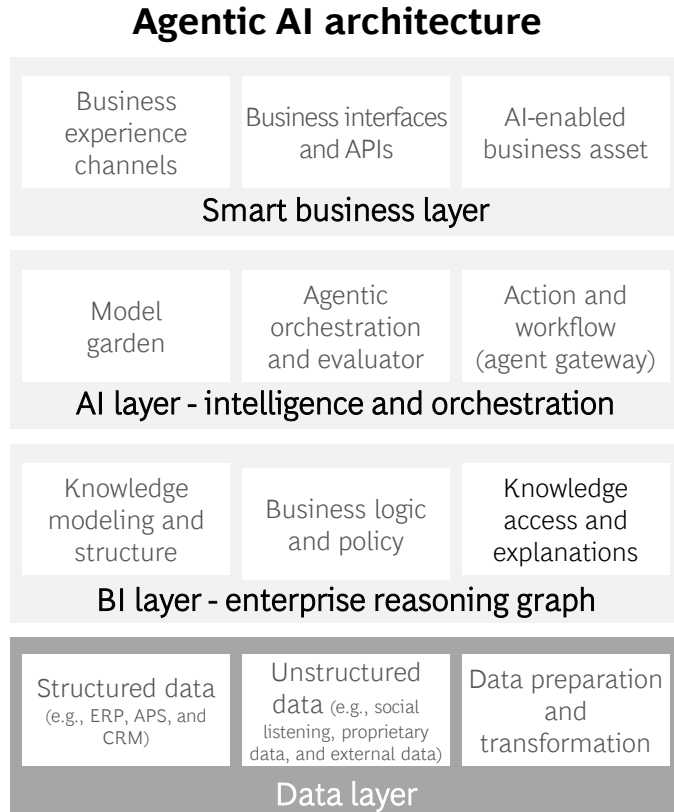
5

## **Make AI decisions transparent, auditable, and explainable**

AI plans must show data sources, assumptions, and tradeoff logic; this is essential to building shared trust across commercial, operations, and finance

# Invest in a robust data foundation

## Leaders are using AI to build the foundation as they go



### 1. Start with AI on the data itself

- AI recommends data fixes (e.g., merge records and fill gaps)
- Stewards review AI recommendations instead of hunting for problems

### 2. Integrate and connect data across silos as you go

- Migrate incrementally, not via multiyear programs
- Build a metadata layer that enables cross-system visibility (OMS, TMS, APS, and ERP) without requiring full migration upfront

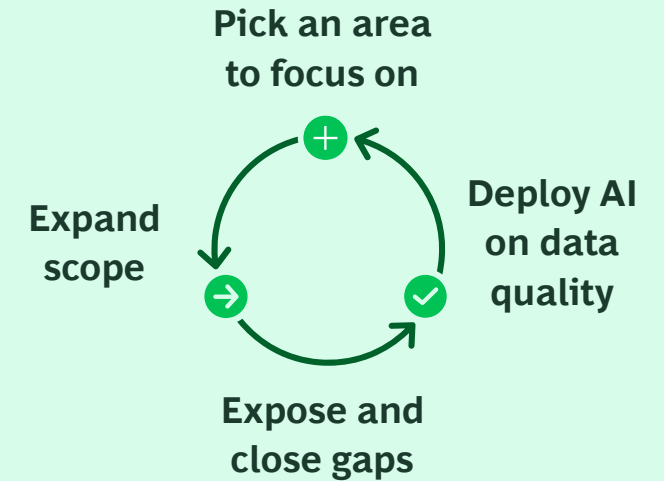
### 3. Let data governance emerge from usage

- Assign data owners with sponsors, but don't let governance delay deployment
- Let agent activity surface issues and highlight highest-value gaps

Best-in-class data strategies use AI to build iteratively

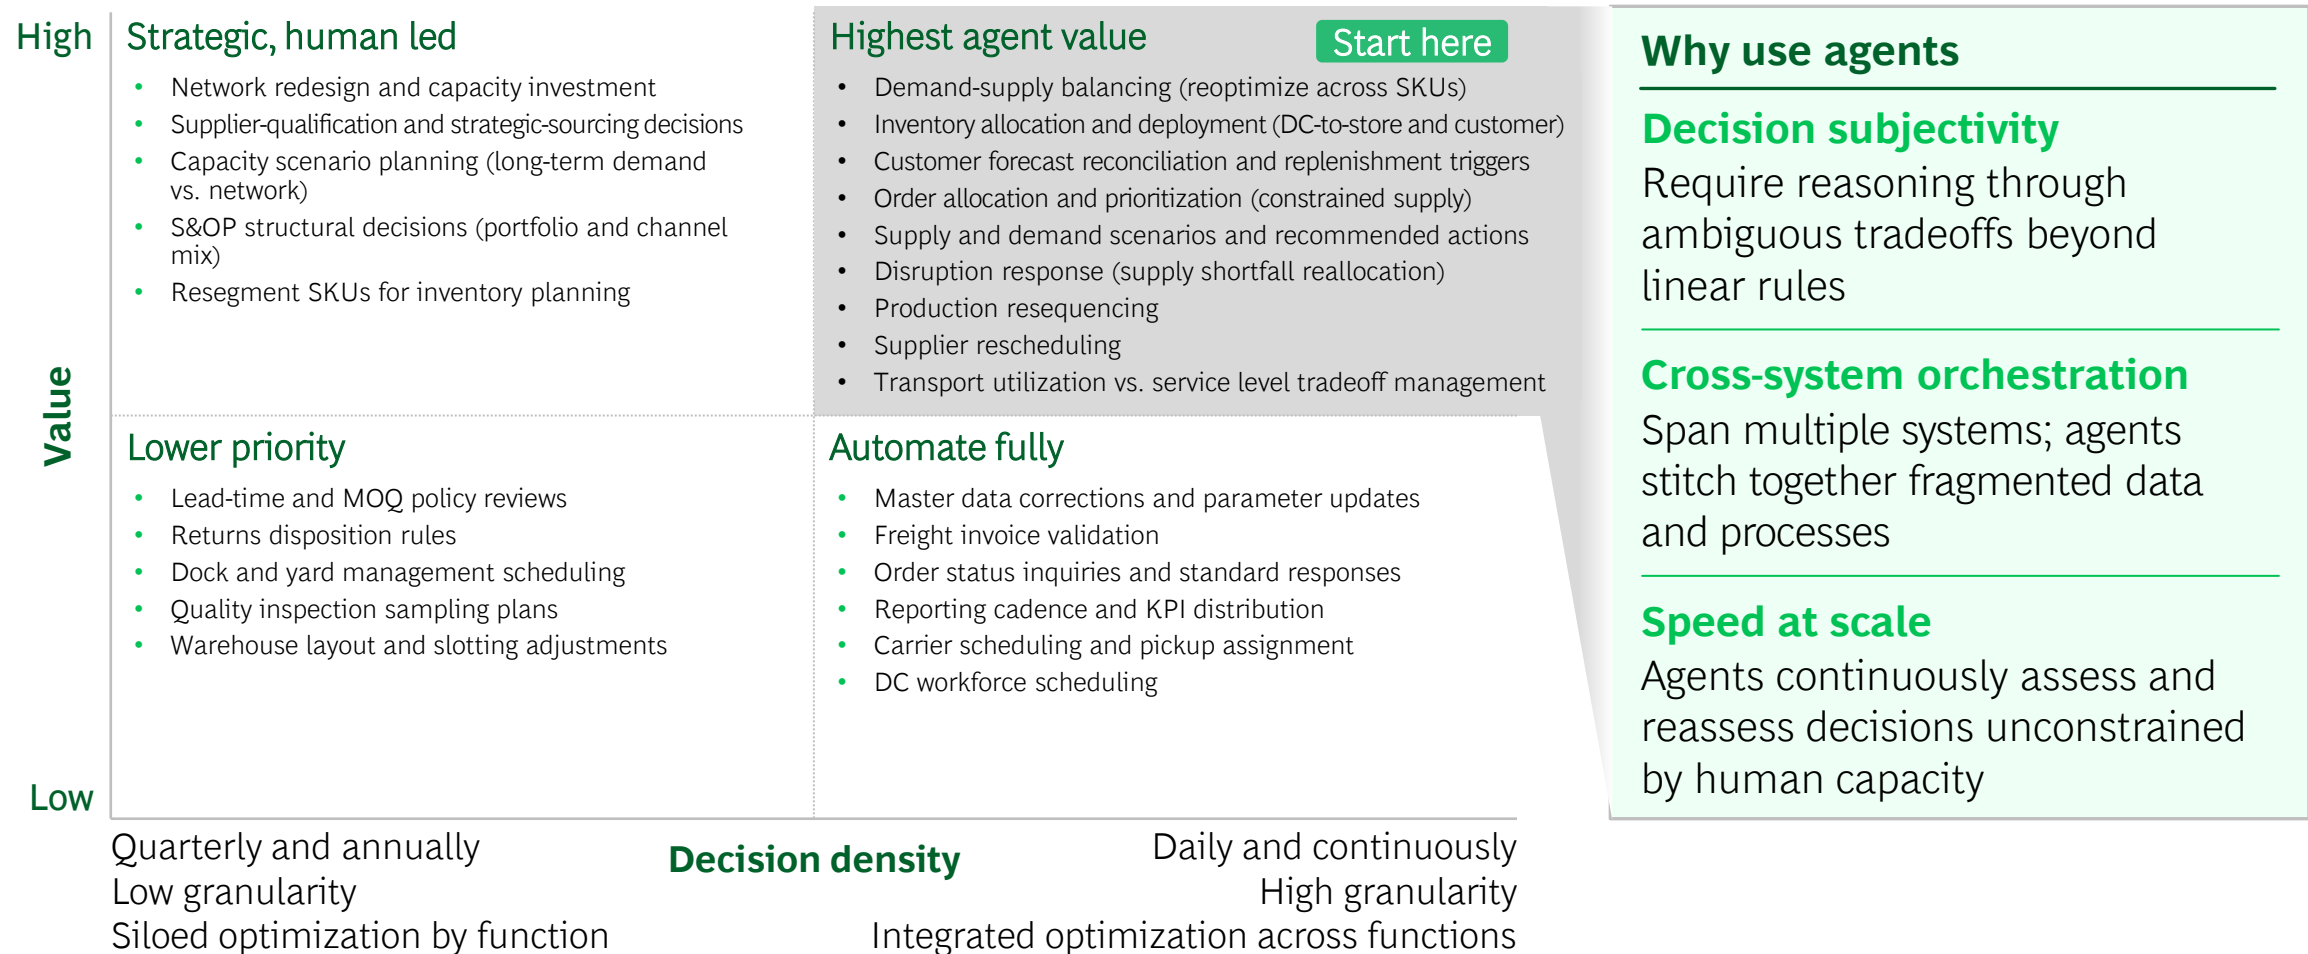
Note: BI= Business Intelligence; CRM = customer relationship management; OMS = order management system.

## Don't wait for a perfect data lake



- ✓ Invest incrementally and use each cycle to expose and close data gaps
- ✓ Ensure the core data foundation improves with every step
- ✓ Deploy AI agents to flag and auto-correct master and transactional data gaps

# Focus first on improving high-value, high-density decisions that are constrained by analytical capacity



Note: MOQ = minimum order quantity.  
Source: BCG analysis.

# Rebuild workflows to be AI first and break down organizational silos across functions

Illustrative

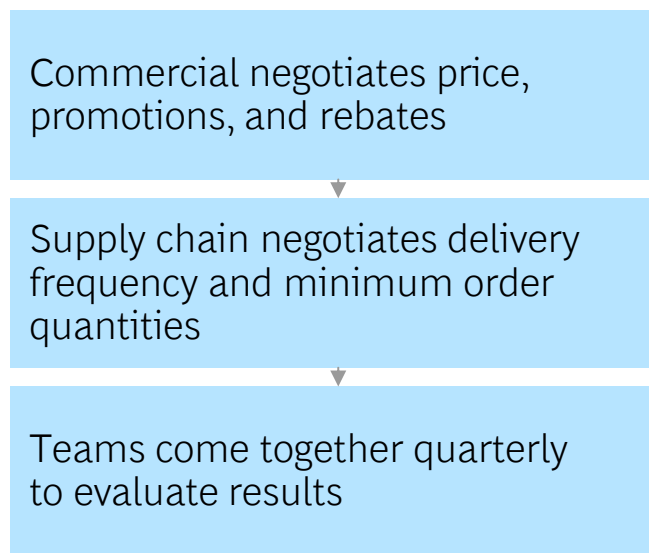
**Currently: Commercial and operations negotiate tactically,** each pushing their own targets

**AI first: Agents provide impartial E2E optimization** and execute low-risk decisions

**Teams shift to driving structural transformation** with shared goals

## Customer negotiations reimaged as an integrated workflow

**From**  
Siloed optimization



Typical process

**To**  
E2E integrated negotiation



AI reimaged

# Implement a hybrid build-and-buy strategy to unlock greater tech value

## Buy the foundational layer and customize the agentic layers

Buy

### APS and foundational planning systems

- Use off-the-shelf planning systems as your system of record for demand, supply, capacity, and inventory
- Use vendor agents for standard APS functions, but own the configuration and deeply understand the decision logic
- Retain in-house control over data, architecture, and governance even when the software is bought

Build

### Custom workflow agents and orchestration layer

- Build agents for cross-system orchestration that no APS can deliver alone
- Focus on highest-value use cases with direct integration into critical systems
- Design for portability using popular frameworks (e.g., ADK and LangGraph) to avoid vendor lock-in

## ⚠ Why neither pure buy nor pure build works

1. APS vendors are embedding agents fast, making it difficult to compete on core planning logic
  2. System-of-record migration risk outweighs the benefit of building one from scratch
  3. Vendors benefit from cross-industry learning that no single organization can replicate
- 
1. Vendors rarely cover cross-process orchestration across OMS, TMS, APS, and CRM
  2. Agent logic must reflect your specific business rules, tradeoffs, and decision thresholds
  3. Custom agents are where supply chains encode structural competitive advantage

# Establish the right transparency and access into decision rationale to build trust in new processes and tools

Traditionally, every team optimizes for their own targets and skills

**Everyone brings their own view:**

Operations on OEE, commercial on revenue, and finance on margin

**Siloed data** and tools **erode cross-functional trust and visibility**

Different team lingo and specializations make **day-to-day collaboration difficult**

Agentic AI can break the silos through transparency and access

Agents provide an E2E optimized plan as a base; **everyone can query and challenge in natural language**

**Agents connect across systems** and make data sources and calculations fully transparent

**Agents act as translators and sounding boards** to bridge functional gaps

## Agent explanation scorecard

Reallocate capacity from Plant B to Plant A

Summary	Shifting to Plant A saves \$30,000 net and puts top customers back on track two days earlier
Recommendation	Shift 12,000 units of SKU-4420 production from Plant B to Plant A in two weeks
Why	Plant A: 18% unutilized capacity Plant B: 97% utilization, three delayed shipments Shift cuts lead time by 2.1 days for the top five affected customers
Tradeoffs considered	Four key alternatives to consider: <ul style="list-style-type: none"> <li>Overtime at Plant B (+\$38,000)</li> <li>Partial customer allocation (-\$120,000 revenue risk)</li> <li>Third-party copacker (+\$52,000 and five-day lead)</li> <li>Reallocation to Plant A (+\$8,000 freight)</li> </ul>
Confidence	Similar reallocations: six times last quarter; 94% average on-time delivery
Data sources	APS (capacity), OMS (order backlog), TMS (freight rates), and CRM (customer tier)

This level of reasoning and transparency is not possible with traditional and rule-based automation

# Scaling agents requires a strong oversight model to secure value

## Risk of ungoverned agent proliferation

As organizations scale, the barrier to creating new agents drops significantly

This is a feature—it enables team-level experimentation and innovation

But as more control shifts to agents, challenges emerge:

- AI-related incidents rose 21% from 2024 to 2025
- AI decision making is expected to rise from 10% to 35% in three years
- 69% of executives say agentic AI requires new management approaches

The challenge is to establish governance without slowing experimentation

## What practical agent governance looks like



### Agent registry and life cycle management

Shared registry so teams know what agents exist, what they do, and where they overlap



### Outcome-based steering

Measure agent outcomes on service, cost, and inventory—not just whether they run



### Defined escalation and decision rights

Define which decisions agents make autonomously vs. which require human sign-off, tiered by risk



### Graduation pathways from pilot to production

Clear process from sandbox to production so experimentation doesn't become permanent tech debt

# Get started on the agentic supply chain journey with these foundational steps

## Strategic Moves

## Foundational Steps

1

**Invest in a robust data foundation**

### **Lay the groundwork with AI, not before AI**

Don't wait for a perfect data lake; use AI to clean, enrich, and connect data across systems, closing the highest-value gaps during each cycle

2

**Start where decision density and value intersect**

### **Identify key decisions and pick one or two use cases to go deep**

Tie each agent to a measurable outcome (revenue, service, cost, working capital); test in a controlled setting, validate, and scale

3

**Rebuild workflows around AI-led enterprise optimization vs. functional negotiations**

### **Define future-state processes with cross functional stakeholders**

Establish target-state vision for service, cost, automation; set AI automation and enhancement as the new starting point; launch the change management journey

4

**Adopt a hybrid build-and-buy approach to platform strategies**

### **Run a build-or-buy diagnostic across core capabilities**

Score each platform on strategic value, customization needs, total cost, and pace of innovation; buy commodity, build source of advantage

5

**Make AI decisions transparent, auditable, and explainable**

### **Stand up an oversight model that scales with agents**

Maintain a shared agent registry with clear ownership, decision rights, and escalation protocols; require every agent to expose reasoning, data sources, and tradeoffs

# BCG experts | Key contacts for the agentic supply chain



**Dustin Burke**  
Managing Director  
and Senior Partner  
*Chicago*



**Olivier Bouffault**  
Managing Director  
and Senior Partner  
*Paris*



**Dan Sack**  
Managing Director  
and Partner  
*Stockholm*



**Tristan Mallet**  
Managing Director  
and Partner  
*Paris*



**Andres Garro**  
Managing Director  
and Partner  
*Miami*



**Aneesh Saxena**  
Managing Director  
and Partner  
*Miami*



**Camille Engel**  
Managing Director  
and Partner  
*New Jersey*



**Stefan Gstettner**  
Partner and Director  
*Frankfurt*



**James Boudreau**  
Partner  
*Toronto*



**Samantha Dunn**  
Director – Supply  
Chain, BCG Vantage  
*London*



**Erick Wesche**  
Associate Director  
*Amsterdam*



**Ashish Pathak**  
Sr. Director – Supply  
Chain and Operations  
*Gurugram*



**BCG**