ECR TAG 2026 | VIENNA | 4 NOVEMBER 2026

Invisible Decisions

How Agentic AI Is Transforming Supply Chains



Maria Villablanca
Global Supply Chain Commentator
Host, Transform Talks & Voices in Transformation

The elephant in the room.

SCEPTICISM.

And it is a reasonable place to start.

77%

Cost and margin pressure

Top-of-mind concern for European grocery CEOs

47%

AI and automation

Among their top-three priorities

70%

AI's EBIT impact

Not measurable yet, or too early to tell

Source: McKinsey & EuroCommerce, The State of Grocery Retail Europe 2026, Grocery CEO Survey.

The problems don’t stop at the company boundary.

ECR Austria’s own survey makes two priorities clear.

91% **COLLABORATION**

of respondents see stronger supply-chain collaboration as relevant today.

65% **EU REGULATIONS**

want ECR Austria to give greater attention to EU regulatory issues.

RETAILERS — **MANUFACTURERS** — **LOGISTICS PARTNERS**

Source: ECR Austria survey, reported by REGAL (8 Jun 2026) and PRODUKT (7 Jul 2026). Questions differ; sample/methodology not disclosed.

Autonomy is advancing.
Governance must keep pace.

Gartner identifies 8 supply-chain technology trends across 3 themes in 2026.

01
AUTONOMY + AGENCY

ACT

Agentic AI

Collaborative multiagent systems

02
SPECIALISATION + INTELLIGENCE

UNDERSTAND

Intelligent simulation

Domain-specific language models

03
TRUST + GOVERNANCE

CONTROL

Decision governance

Product provenance

MY TAKE **The question isn’t whether an agent can act — it’s when we should allow it to.**

Source: Gartner, “Top Supply Chain Technology Trends for 2026”, 30 Jun 2026. Six selected trends shown; grouping is Gartner’s.

Seeing the problem is no longer the end of the process.

Imagine a retail promotion selling faster than forecast.

01

DETECT

Detect the deviation

Demand is exceeding forecast.
A stockout may follow.

02

EVALUATE

Evaluate the trade-offs

Check available stock, capacity, cost and customer priorities.

03

ACT & VERIFY

Initiate permitted action

Request a stock transfer, notify teams and track results.

AUTHORISED — NOT UNLIMITED

Good data • Decision rules • Approval limits • Human escalation

Illustrative FMCG scenario — not a documented live deployment. Concept informed by Gartner (30 Jun 2026) and public SAP Agent Shift discussions.

Unilever: 15–20 minutes down to 90 seconds

Five agents. One tightly defined workflow. Human review when needed.

FIRST STAGE OF ORDER PROCESSING

15–20 min BEFORE

~90 sec WITH NOVA

Reported across 11 pilot markets

95% TOUCHLESS ORDERS in covered workflow

FROM MY INTERVIEW | WILLEM UIJEN & JENNIFER HAN



Source: Microsoft UK Stories / Unilever, 26 October 2026. Company-reported results; not a measure of full supply-chain autonomy.

Nestlé USA: exceptions, not endless chasing

An agent contacts carriers automatically — and escalates the cases that need people.

CARRIER OUTREACH

~95%

of external customer shipments

23,000+

automated carrier emails in one month

+10 pp

ON-TIME PICKUP

Laporte pilot site

+8 pp

ON-TIME DELIVERY

Jonesboro pilot site

A BOUNDED DECISION LOOP

DETECT — **CONTACT** — **LOG** — **ESCALATE**

Human coordinators focus on the exceptions that remain unresolved.

Source: FourKites / Nestlé USA, 23 October 2026. Site-specific, company-reported pilot results; pp = percentage points.

Walmart: many decisions, one customer promise

A coordinated system — not a claim that one AI agent runs the entire supply chain.

01

PREDICT

Demand and inventory placement

Predictive AI / ML

02

SELECT

Best fulfilment node

AI agents + decision logic

03

COORDINATE

Routes, drivers and dispatch

Multi-agent workflow

04

UPDATE

Delivery promise and ETA

Real-time ML + location

A SINGLE OUTCOME: THE RIGHT ORDER ARRIVES WHEN PROMISED

<30 MIN some delivery windows (not agent-only impact)

Source: Walmart Global Tech, 8 Dec 2025. Walmart describes its architecture; no agent-specific financial ROI or causally isolated speed gain disclosed.

The real prize is preventing the problem

In my Agent Shift conversation with SAP and McKinsey, one example changed the question.

THE OLD WAY

MONDAY MORNING

Hundreds of exceptions

Review the queue – Investigate consequences – Mitigate the damage

REACTIVE / EXCEPTION HANDLING

THE OPPORTUNITY

BEFORE MONDAY

Prevent avoidable exceptions

Detect the risk – Evaluate options – Intervene earlier

PROACTIVE / EXCEPTION PREVENTION

MY QUESTION: ARE WE AUTOMATING FIRE-FIGHTING — OR PREVENTING THE FIRE?

Source: Maria Villablanca, Voices in Transformation / Agent Shift interview with Maximiliano Schlichter (McKinsey) and Matthias Vogel (SAP), 2026. Guest-described example; no quantified outcome.

Progress is real. Scaling is another matter.

A few strong deployments do not mean that agentic AI is already mainstream.

WHERE CPG COMPANIES STAND ON AI

75% In pilots or exploration

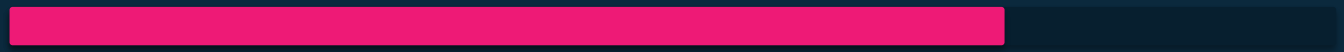

18% Scaling significant impact

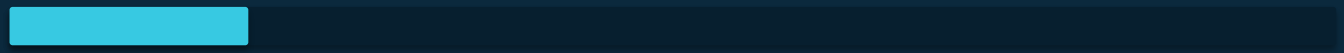

Figures describe AI maturity generally, not agentic supply-chain deployments specifically.

THE MEASUREMENT PROBLEM

>50%

of all surveyed CPG and retail leaders do not formally measure AI ROI.

BCG + Consumer Goods Forum (24 Jun 2026), survey of 39 CPG and retail executives worldwide. Small directional sample; percentages are for distinct cohorts/questions.

Whose decision is the right decision?

An agent can optimise a company's KPI — and still make the whole value chain worse.

RETAILER

Availability + stock turns

Protect the shelf without tying up cash

MANUFACTURER

Margin + production efficiency

Protect service without costly disruption

LOGISTICS PARTNER

Capacity + transport economics

Deliver reliably within real constraints

ONE SHARED RESULT: AVAILABILITY + COST + SERVICE + TRUST

WHO SETS THE PRIORITIES? • WHO SHARES THE DATA? • WHO TAKES RESPONSIBILITY?

Illustrative cross-company decision tensions, not a documented production agent network. Context: ECR Austria's Joint Forecasting programme; BCG 2026 research.

Where would you permit autonomy?

The consequences of a mistake matter as much as the ability to automate a task.

	LESS REPEATABLE / UNCLEAR RULES	MORE REPEATABLE / CLEAR RULES
HIGHER STAKES	**HUMAN DECISION** Unusual shortage allocation; unverified supplier changes	**APPROVAL + CONTROLS** High-value inventory transfers; contractual changes
LOWER STAKES	**ASSIST + LEARN** Unfamiliar, reversible exceptions; recommendations reviewed	**AUTOMATE WITH GUARDRAILS** Carrier follow-ups; routine order validation

Autonomy is not binary: permission can expand as reliability, evidence and safeguards improve.

Illustrative editorial matrix; informed by Gartner (30 Jun 2026) and published Agent Shift interviews.

Trust has to be engineered.

Reliable context. Clear rules. Bounded authority. Evidence.

01 DATA + CONTEXT
Accurate, current, commercially meaningful information

02 DECISION RULES
Explicit business priorities, constraints and trade-offs

03 AUTHORITY
Clear permissions and human approval thresholds

04 EVIDENCE
Traceable actions, review and measurable outcomes



VOICES IN TRANSFORMATION

Maria with Maximiliano Schlichter (McKinsey) and Matthias Vogel (SAP)

THE LEADERSHIP TEST: What evidence justifies delegating this decision?

Source: Voices in Transformation, Agent Shift — The Trust Imperative; Gartner, 30 June 2026. Framework is Maria's synthesis.

Are we transforming — or automating the past?

A smarter process is not necessarily a better operating model.

THE DEFAULT

“How can AI make each function faster?”

PLAN | MAKE | MOVE

Risk: optimising separate KPIs

THE ALTERNATIVE

“Which decisions create the right business result?”

SERVICE + MARGIN + RESILIENCE

Then design the decisions, permissions and workflows around it.

THE TEST: Are we redesigning the decision — or merely accelerating the old hand-offs?

Original editorial comparison. Supporting context: McKinsey (22 Jul 2026); BCG (2026); published SAP/McKinsey Agent Shift interview.

Start with a decision. Prove the economics.

Five questions before any agentic AI investment.

01 OUTCOME

What must change in the business?

02 DECISION

Which recurring choice stands in the way?

03 READINESS

Do data, rules and systems support action?

04 AUTHORITY

What may an agent do without approval?

05 EVIDENCE

What baseline and return will we measure?

NOT READY? Fix the basics first. No agent is sometimes the right decision.

Maria's original investment checklist; synthesised from published SAP Agent Shift interviews and BCG/McKinsey 2026 research.

Resilience is the real test of connected decisions.

A disruption quickly becomes a problem across planning, inventory, logistics and service.

THE SHOCK A critical inbound shipment is delayed or a trade route closes

01
SENSE
Detect the disruption and forecast impact

02
ASSESS
Map affected orders, stock and customers

03
COMPARE
Test inventory, source and transport options

04
ACT
Execute permitted steps; escalate trade-offs

Every response is a business choice: SERVICE • COST • CAPACITY • RISK • COMPLIANCE

Illustrative 2030 scenario from Maria's SAP/LeverX interview — not a live autonomous deployment.

Invisible decisions.

Visible responsibility.

If we were designing the relationships between manufacturers, retailers and their supply chains today — knowing what AI agents might make possible —

would we design them the same way?

CONNECT ON LINKEDIN

The goal is not the most autonomous supply chain. It is the most effective one.