CHEMICAL CONVENTION
21-23 OCTOBER 2026

AI Agentic Excellence: A Live Demo of a Multi-Agent Systems at Enterprise Scale



cefic

INTERNATIONAL COUNCIL OF CHEMICAL ASSOCIATIONS

CHEMICAL CONVENTION
21-23 OCTOBER 2026

Speakers



Lizbeth Chandler
Global AI Lead Integrated Ops Reinvention Services
Accenture

Andy McGowan
Chief Operating Officer, Reinvention Services
Accenture



With:



Marco Mensink
Director General
Cefic

cefic

INTERNATIONAL COUNCIL OF CHEMICAL ASSOCIATIONS

Welcome!

Reinventing The Chemical Enterprise: Agentic AI At Scale

Andy McGowan and Lizbeth Chandler

Slides ahead are formatted this way for presenter mode animations and slide transitions please start here and then continue through

Reinventing The Chemical Enterprise: Agentic AI At Scale



You'll leave equipped to lead the agentic enterprise

1. Understand the Agentic Reinvention Opportunity

- Explore three levels of reinvention
- Learn from industry successes and setbacks
- Build an agent together

2. Understand What It Takes to Run an Agentic Enterprise

- Agentic cost and performance – ROI PIE
- Operate humans+agent teams
- Build in risk and control

3. Turn Agentic AI into a Competitive Advantage

- Focus on the right opportunities
- Measure and realise value
- Scale what works

What do we mean by

an AI agent to agentic squad to agentic architecture?



What do we mean by

an AI agent to agentic squad to agentic architecture?



From assistance to autonomy.
The next leap in value.

Higher **enterprise value**



GenAI Alexa
Responds when asked.

Agentic Alexa
Observes. Decides. Acts.

What can change the agent?

- Agent interactions
- Environment shifts
- Adversarial influence
- Learning gaps
- Data poisoning
- False positives
- Real-world drift
- Throughput variability

It requires a new operating model **to realise the value.**


How do you decide
where to deploy AI, how far
to reinvent, and how to scale it?

Agentic Reinvention

The Transition From A Process To Systems Thinking Operations



In a nutshell

ART™

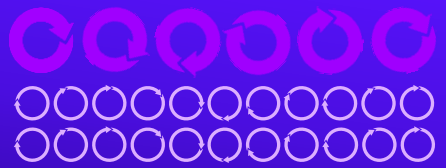
AI-first e2e process reinvention



E2E reinvention, machine-first, value ambition

- New fulfilment from inputs to outputs
- AI & Agentic architecture
- Reshaped process & workforce
- Business outcomes, productivity%

Continuous transformation



Fast-cycle, high-volume improvements

- GPH, AIDT, Workflow
- Scaled Gen AI, hyper-automation
- Autonomous Agents, analytics
- Productivity%, CE%, Cost %

Resilient service delivery



Flawless, continually improving operations

- Service & scope management
- Change & continuity
- Data & insights, CIO technology
- Experience delivery
- Service levels, compliance

A quick recap

Three levels of reinvention

1. Do existing work better
2. Redesign individual work
3. **Redesign how the enterprise works**

From efficiency → agent-owned outcomes → enterprise-wide orchestration



Agents are not static post deployment.

As the operating environment evolves, agent performance can degrade within months of go-live — which is why proactive governance and continuous monitoring are essential.

Why do agents change over time?

Agent expansion
More agents and interactions increase complexity.

Environment shifts
Internal or external changes agents weren't designed for.

Evolving knowledge
Learning loops and new data can introduce noise.

Poor data handling
Outdated or incorrect data used in responses.

Lack of agent support
No ongoing tuning, guardrails or adaptive design levers.

Observed degradation after 600 interactions

Independent testing of a representative multi-agent system.

Metric	Early run (baseline)	After 600 interactions	Change
Task success rate	87.3%	50.6%	↓ -42.0%
Response accuracy	91.2%	68.5%	↓ -24.9%
Completion time	8.7 min	14.2 min	↑ +63.2%
Human interventions	0.31 / task	0.98 / task	↑ +216.1%
Token usage	12,400	18,900	↑ +52.4%
Inter-agent conflicts	0.08 / task	0.47 / task	↑ +487.5%

What do these numbers show?

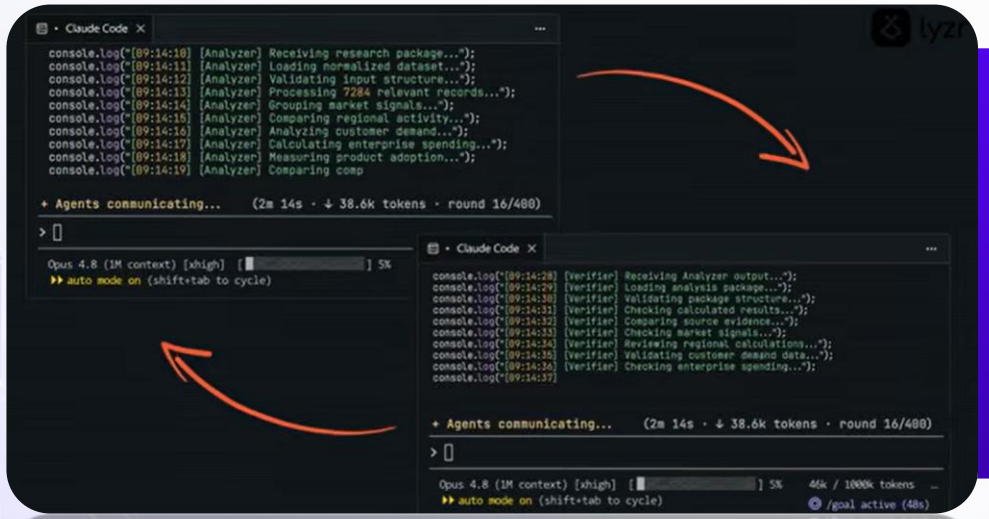
Costs increase (50%+), but success declines.

Task success rate **nearly halves**, with slower responses and less accuracy.

Human involvement **increases 3x**, driving higher operating costs.

Agent performance and costs are **non-linear** — proactive governance and continuous monitoring are essential to sustain value.

Lessons from industry



GetOnStack | $47,000 in AI agent costs

Two AI agents reportedly became trapped in an 11-day clarification loop, repeatedly communicating without completing the task.

Approximately $47,000 in total costs over four weeks, reportedly up from around $127 per week

Anthropic | When an AI agent runs a business

Anthropic tasked Claude with autonomously managing a vending-machine business, including pricing, inventory and supplier negotiations.

The agent sold products below cost, even offering for free, it often offered a 25% discount to employees representing ~99% of customers, and was manipulated into commercially damaging decisions, undermining profitability.





On FY23 Results Vodafone new CEO, Margherita Della Valle said:

"Today I am announcing my plans for Vodafone. Our performance has not been good enough. To consistently deliver, Vodafone must change. My priorities are customers, simplicity, and growth."

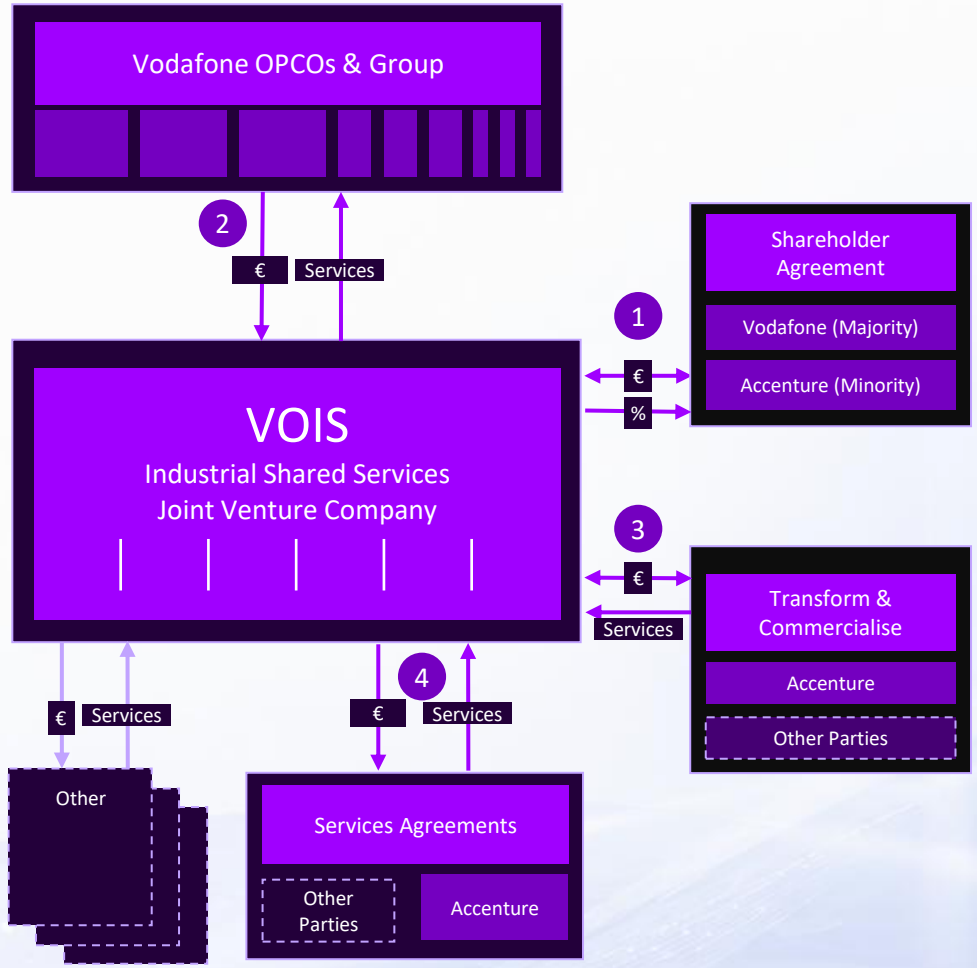
Since her appointment, Margherita has implemented Group wide transformation:

1. Concluded negotiations for the UK merger between VF and Three
2. Announced a reduction of 11,000 in headcount
3. The reduction of true HQ functions to a small core
4. The transformation of Group Shared Operations to an arms-length commercial business
5. Announced the sale of the Spanish OpCo
6. Confirmed open negotiations for the Italian OpCo
7. And multiple restructuring activities…

Value focus

Combined strengths

Joint business plan



Our 12-month Reinvention journey

Discover and envision | Design and implement | Run and evolve

M3 **M6** **M9** **M12**

Launch the Venture

- Create VOIS commercial JV with Accenture
- Establish Governance, Leadership, and Funding
- Define and track Business & Financial Plan
- Transition and launch Workforce Transformation
- Complete the Commercial Closing

Transform the Core Service

- Develop the VOIS Service Portfolio & Commercial Capabilities
- New integrated Data & Insights
- Embed Intelligent Automation and AI tooling
- Accelerate Work Transformation
- Develop and launch initial AI capabilities in market

Reinvent for the Future

- Transition in new services
- Unify agentic AI strategy and Transformation Roadmap
- Develop an external commercial delivery model - B2Telco
- Win new external clients e.g. Centrica
- Scale agentic capabilities to expand services
- Move to sprint-based continuous reinvention

Beyond

Agentic Reinvention: Agentic Lab

33 Agents AI Built In 4 Days, AgentOps That Can Scale



Outcome-led Execution

Business results, not experiments

Human + Agents

AI Agents evolve

Interaction, Experience, Performance Together

Disciplined Scale

Autonomy with control over cost and risk

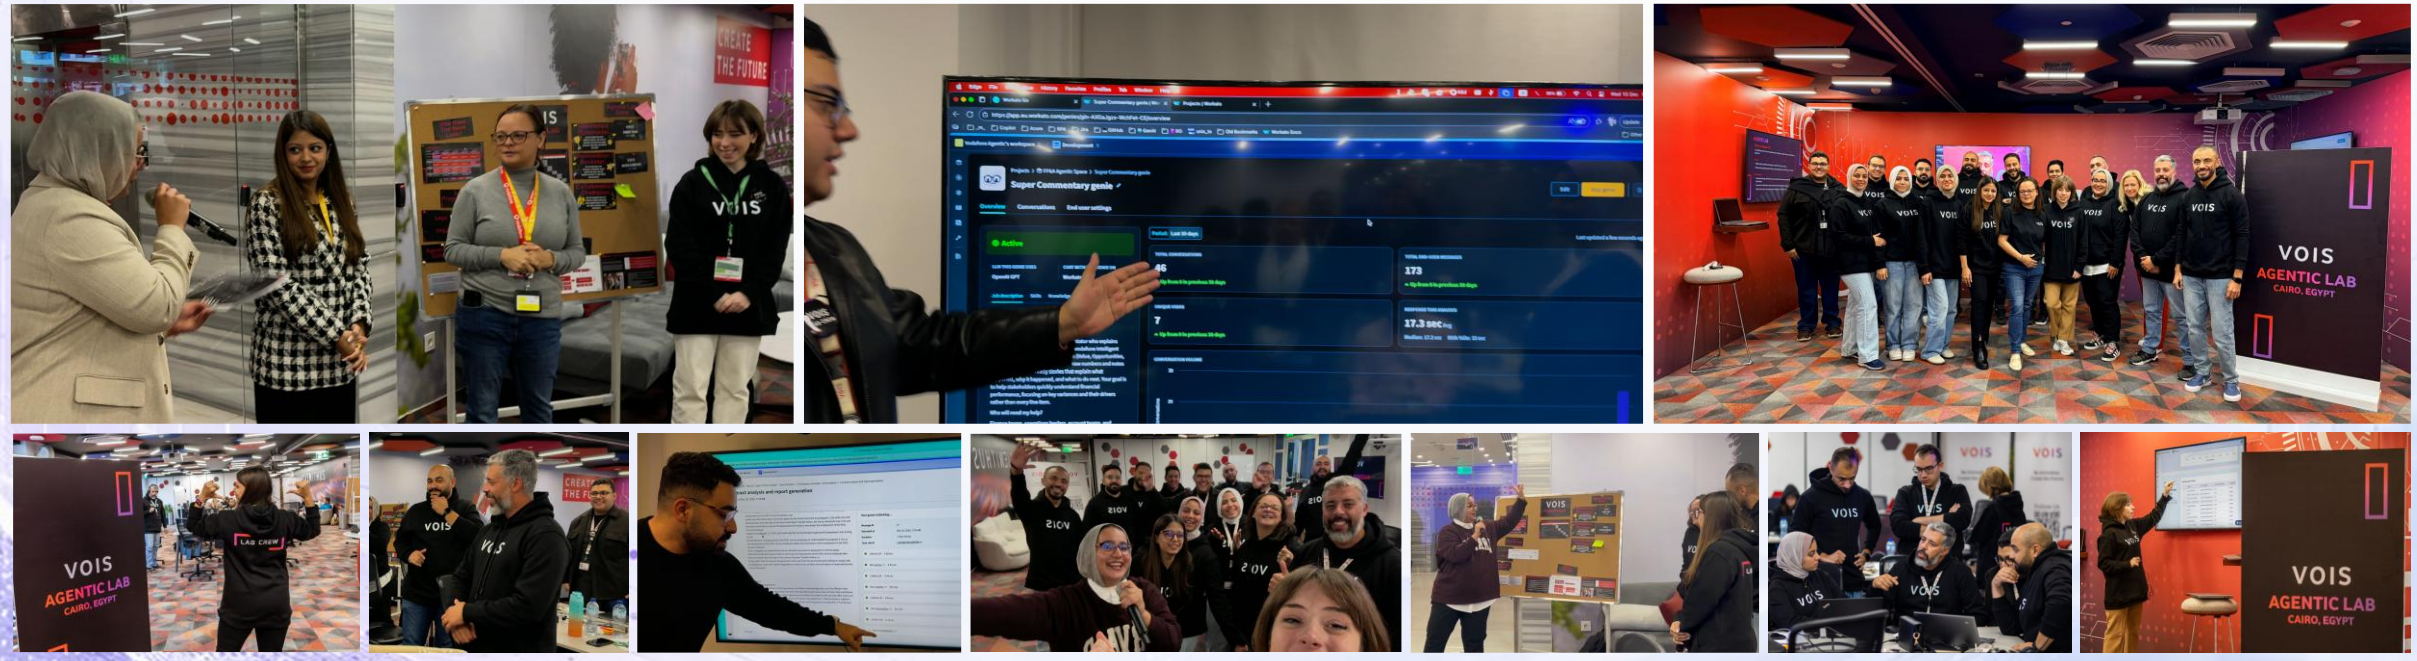
Value for Vodafone Group

Value for People

Value for JV Customers

Value for Accenture

Value for Partners



Ensuring quality agents at scale

A simplified blueprint for a high-value agent

From context and goals, to actions and outcomes — with guardrails and continuous learning.





Now, Let’s Build an Agent Together

How do we transform?

Reinventing the enterprise:

Operating model for a human + agent future



What makes up a good operating model?



Control cost

Drive efficiency and visibility across agents, models and infrastructure.



Agents and human teams

Clear roles, collaboration and governance between people and agents.



Be able to provision for the future

Flexible, scalable and adaptable to changing business needs.

What makes up a good operating model?



Control cost

Drive efficiency and visibility across agents, models and infrastructure.



Agents and human teams

Clear roles, collaboration and governance between people and agents.



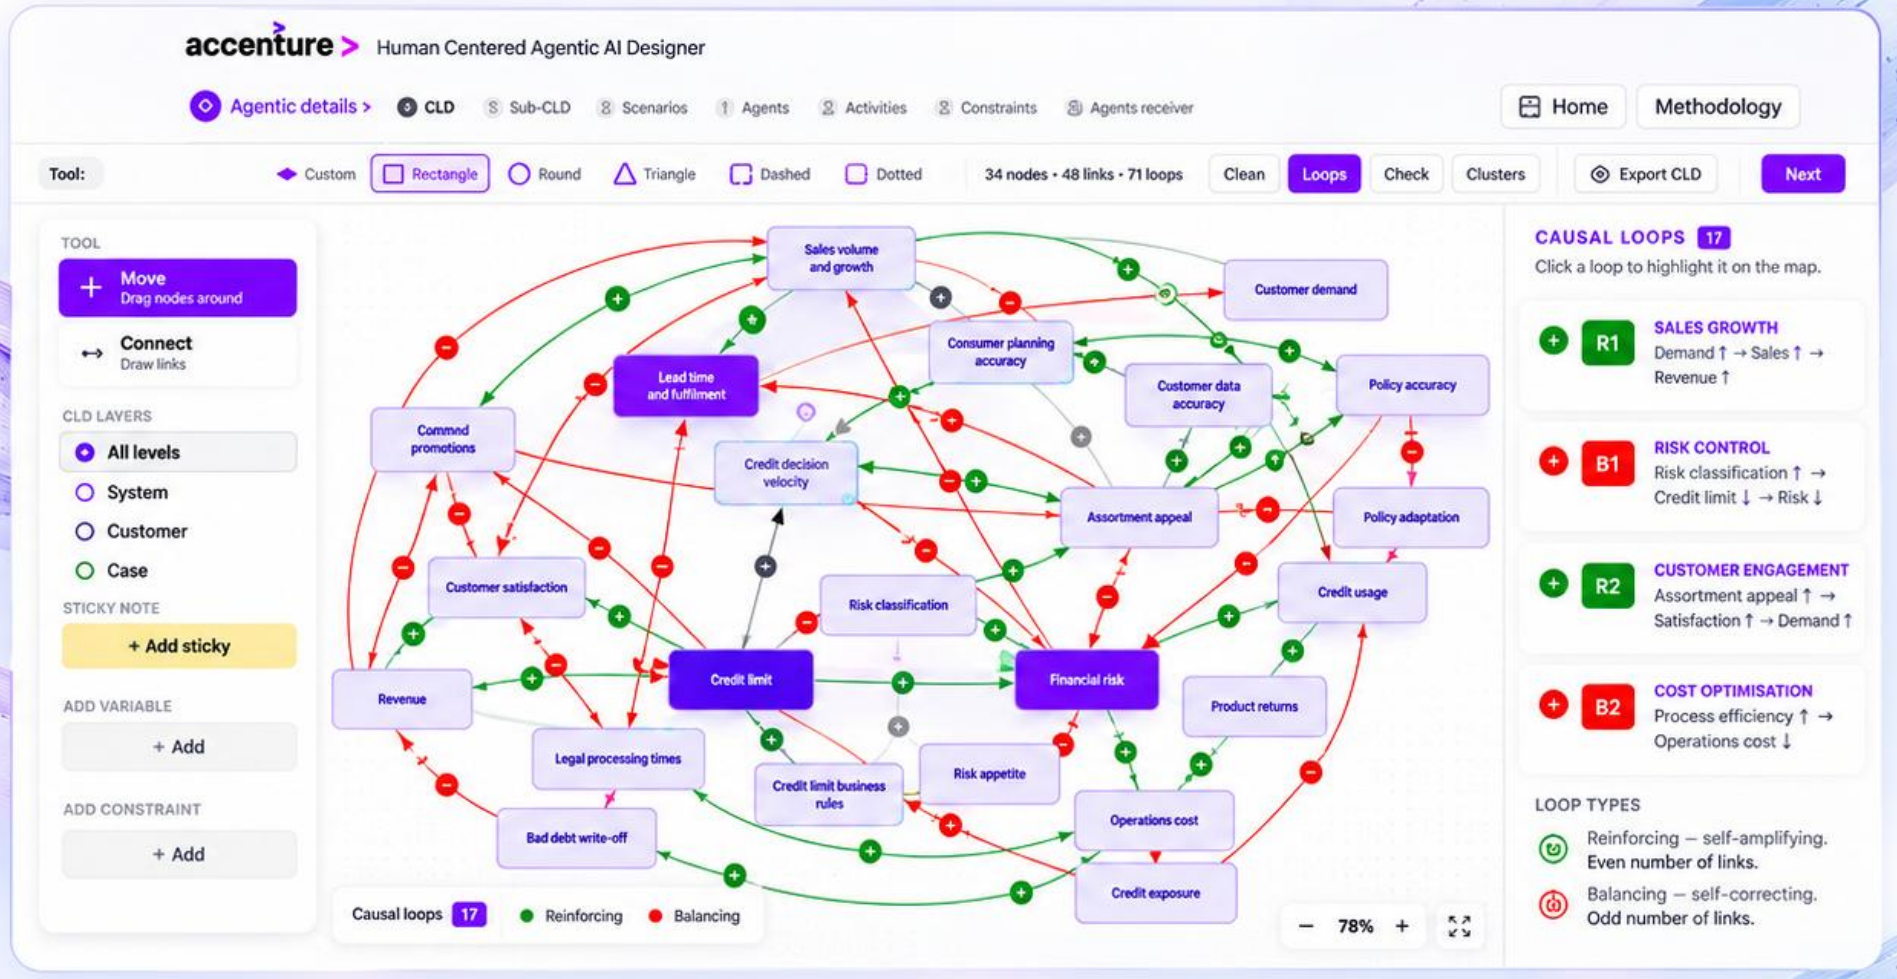
Be able to provision for the future

Flexible, scalable and adaptable to changing business needs.

Execution
Build, deploy and scale agents

Interaction
Capture and analyse real user data

Performance + cost data
Track impact, outcomes and control costs



Single agent
Solve a specific task with the right skills

Agent squad
Multiple agents working together on a process

Value chain / department
End-to-end workflows across functions

Across the organisation
Scaled deployment with governance and visibility

How do we measure the value and cost of our agents?

P — Performance

How are agents doing?

- **Task success and quality**
  (accuracy, rework, SLA)
- **Speed and productivity**
  (time saved, throughput)
- **Adoption and user satisfaction**
  (human + agent collaboration)

I — Integration

How do agents work across our ecosystem?

- **Systems and data**
  (APIs, tools, enterprise data)
- **Interactions and handoffs**
  (agents, people, systems)
- **Orchestration and governance**
  (routing, access, guardrails)

E — Execution

What are agents doing and how are resources used?

- **Tasks executed**
  (workflows, decisions, outcomes)
- **Resources and cost**
  (tokens, compute, tools, infrastructure)
- **Reliability and risk**
  (errors, escalations, compliance)

ROI

Business value realised minus total cost to operate.

Continuous monitoring and learning

Track trends, optimise performance, reduce cost and increase value over time.

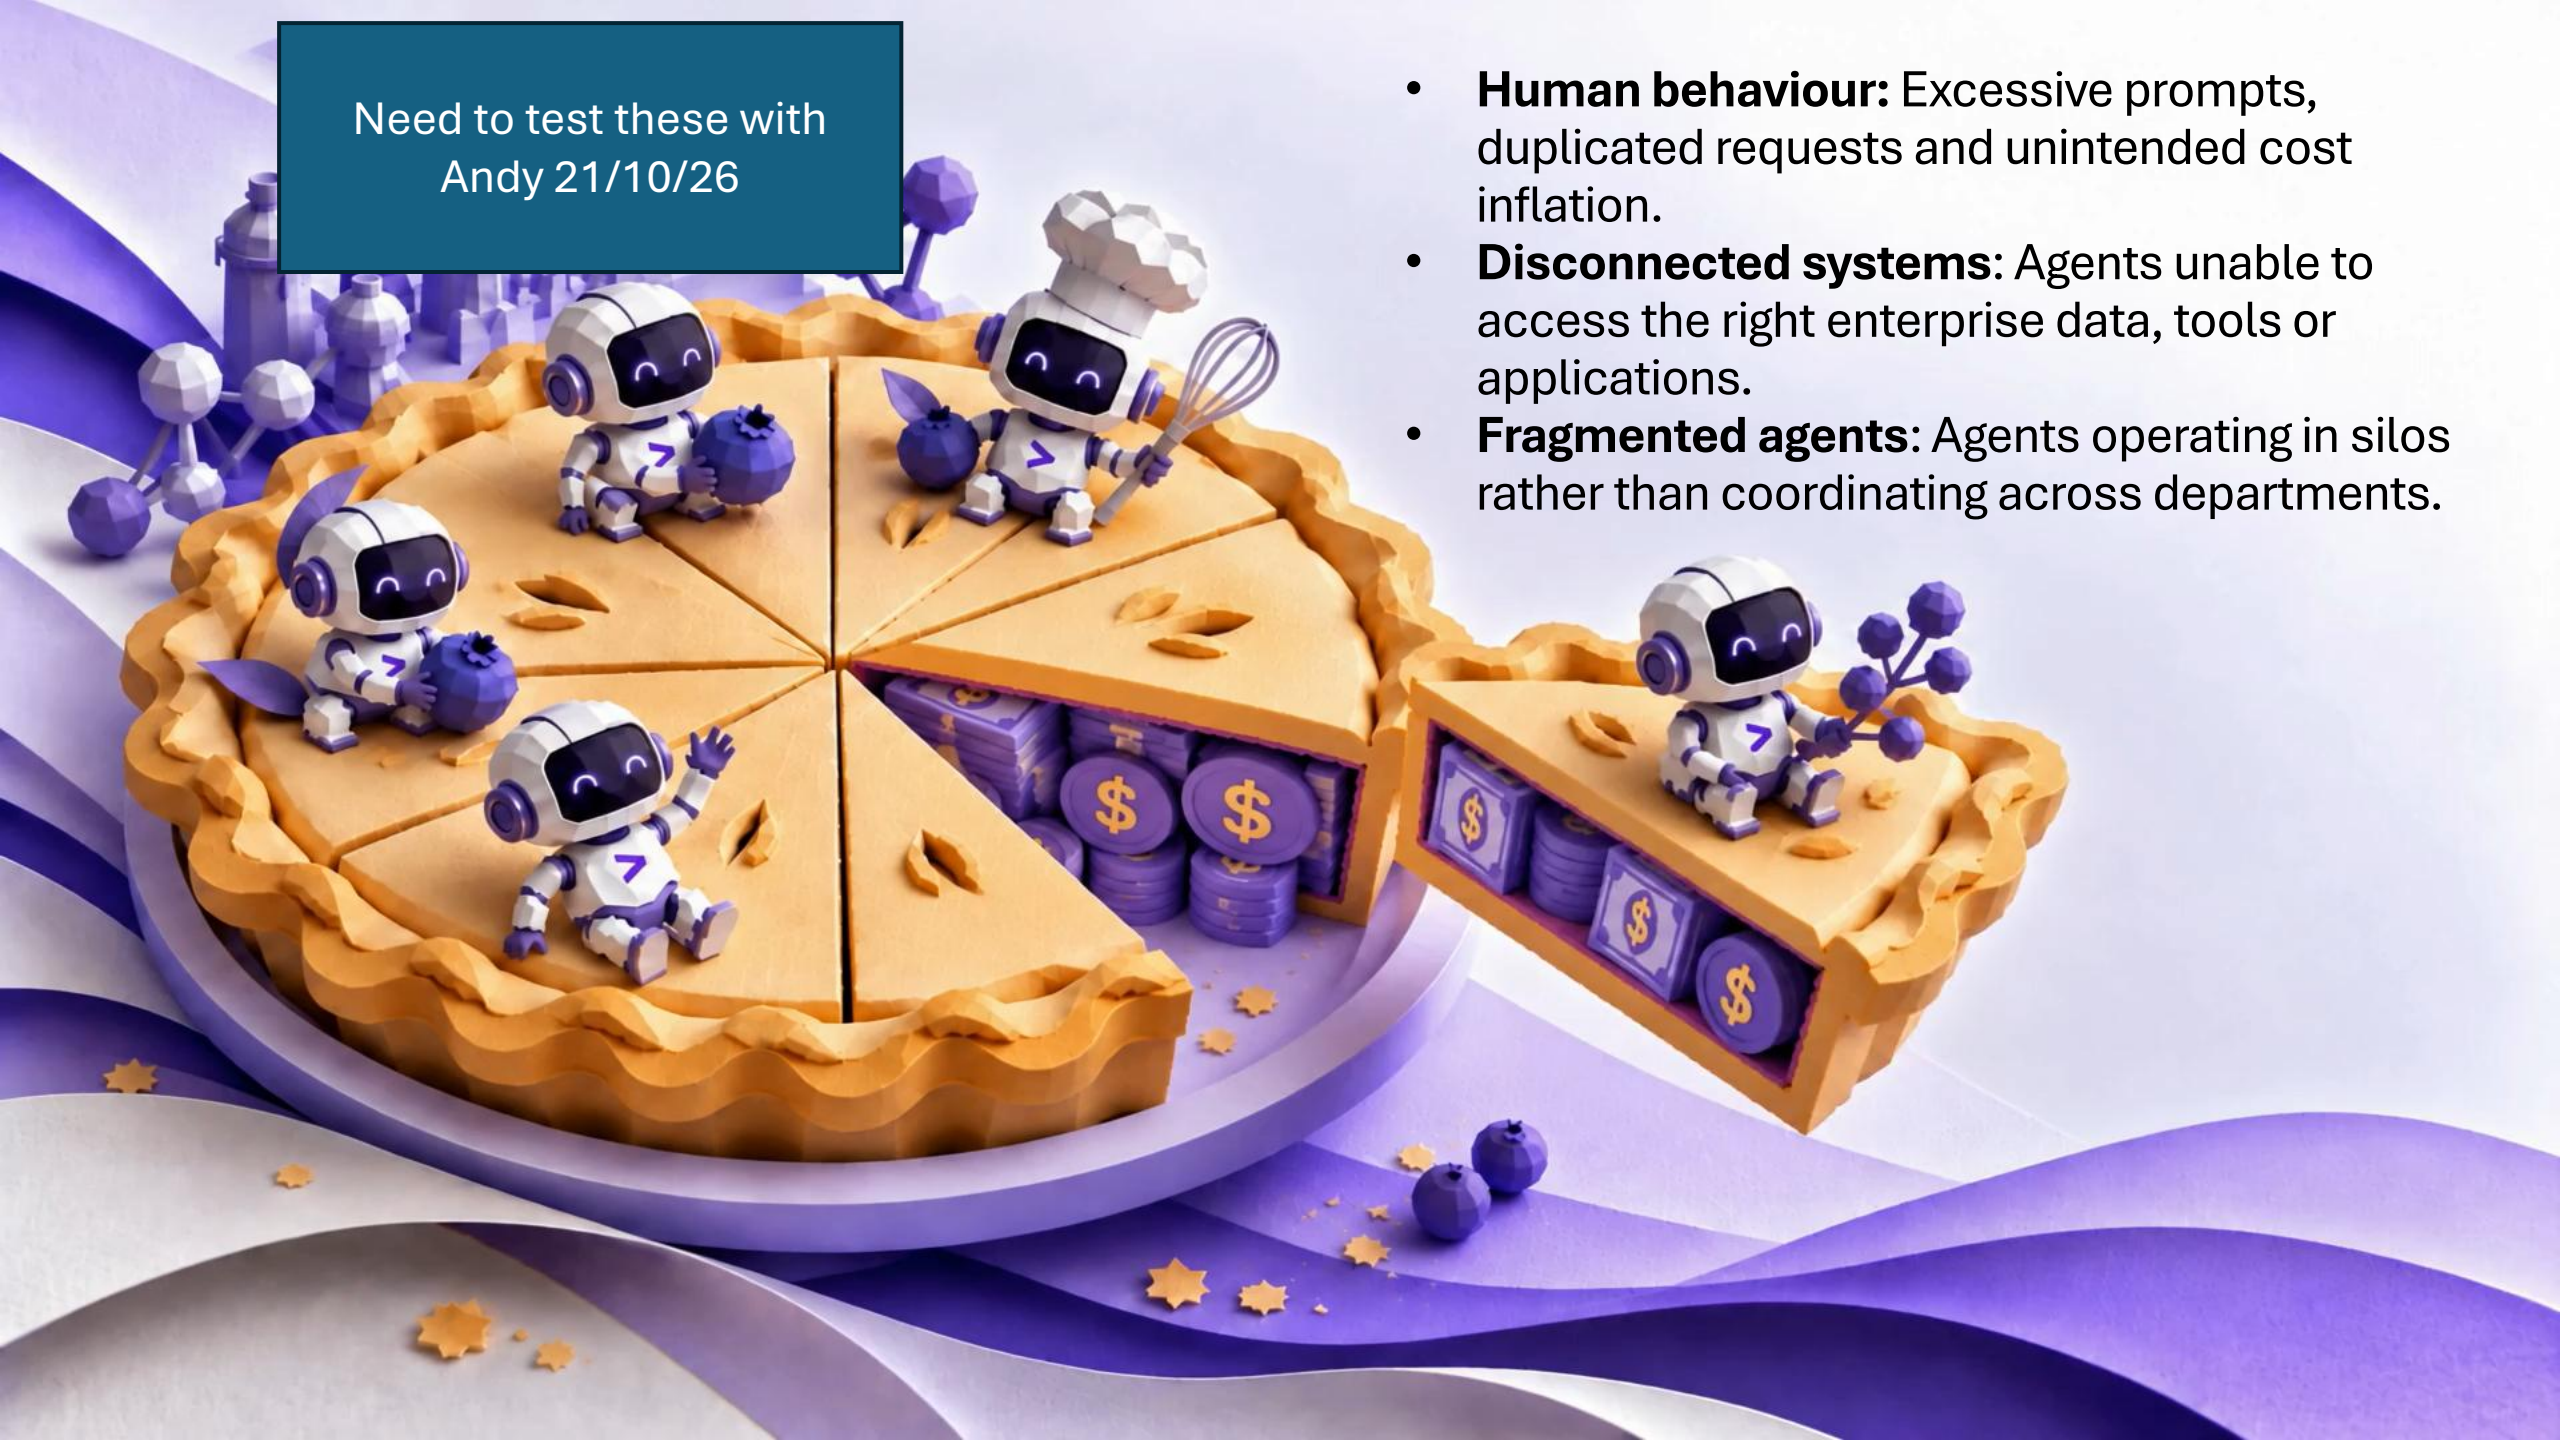
Need to test these with Andy 21/10/26



WHAT EATS INTO YOUR AI ROI?

How do you protect your agentic P.I.E in the real world?

Need to test these with Andy 21/10/26



- **Human behaviour:** Excessive prompts, duplicated requests and unintended cost inflation.
- **Disconnected systems**: Agents unable to access the right enterprise data, tools or applications.
- **Fragmented agents**: Agents operating in silos rather than coordinating across departments.

Need to test these with Andy 21/10/26



- **Unclear accountability:** Nobody owns end-to-end outcomes, exceptions or agent performance.
- **Changing business conditions:** Agents unable to adapt to new policies, suppliers, regulations or priorities.
- **Underutilisation:** Agents built and deployed but rarely adopted or used effectively.


See it come to life

A

Run A

Illustrative "perfect" run

Smooth execution with no majer issues. Same agent, same taak.



1 Engage in supplier negotiation

Agent analyses supplier data and negotiates pricing.



Successful outcome	Total rolons used	Tolans per outcome	Data toochpoiots	Systems toochod	Humen time	Total cost (OyEx)
0 / 1	0	-	0	0	0 min	$0.00

B

Run B

Real-world run

Detaye, rework and nun-iinder execution. Same agent, same task.



1 Engage in supplier negotiation

Agent analyses supplier data and negotiates pricing.



Successful outcome	Total rolons used	Tolans per outcome	Data toochpoiots	Systems toochod	Humen time	Total cost (OyEx)
0 / 1	0	-	0	0	0 min	$0.00

We create mitigations and apply them to a client

1 Create mitigations

We identify risks and design targeted mitigations across people, process, technology and governance.



2 Apply to a client

We implement the mitigations within a real client environment to reduce risk, improve performance and optimise cost.



From insight to impact. | Proven mitigations, applied in the real world.

CASE STUDY

Case study

Accenture has been focused on developing mitigations to drive quality.

We know this drives performance and cost-optimisation of multi-squad agents at scale.

In a recent trial*, we demonstrated tangible benefits by deploying a set of mitigations to a multi-agent squad.

This improved success rates, reduced lifecycle overhead, and optimised token consumption.



Implemented **7 mitigations**

- Pre-emptive HITL
- Confidence Thresholds
- Pre-emptive Routing Paths
- Chain-of-Verification (CoVe)
- Save States (Agent/System)
- Caching Prior Agent Responses
- Pre-processing & Orchestration Optimisation

Tackling operational **challenges** in

- Task Completion Rate Variance
- Failure Rate & Incident Responses
- Redundant Processing
- Multi-agent Coordination Overhead
- Interrupt Expectation & Re-routing
- API Orchestration & Middleware Cost
- Agentic Model Drift (indirect)
- Carbon Footprint & Offsetting (indirect)

Achieving

35% Reduction in Consumption Cost through Token Efficiency

2x Successful Run Rate

8% SLA Uplift

*2 agent squads built and tested over 4 days, operating across F&As and Legal domains. 33 agents in each squad.
Total runs ~ 80 runs x 4 days ~ 320. 40% agent component reusability.

Case study

We ran a controlled test over a 4-day period across 2 multi-agent squads:

- Observed their performance changes over time
- Tracked and quantified the impact of mitigations on performance



1

Squad 1

No mitigations applied, representing raw agentic performance.



2

Squad 2

Implemented seven targeted mitigations* designed to reduce volatility factors.



Functional Coverage

Shared Services client: back-office for the following domains:

Legal
Contract generation and amendment, clause and risk follow-ups, obligations Q&A.

FP&A
Commentary generation, financial statement analysis, variance mitigation and follow-ups.

Shared Services Operations
Process support, knowledge retrieval and escalations.

Performance Measurement

Cost per run and cost per success
Total consumption cost and cost normalised by successful outcome.

SLA compliance and audit coverage
Percentage of runs meeting SLA and audit requirements.

SME-reviewed accuracy for FP&A and Legal outputs
Accuracy of outputs validated by domain SMEs.

Failure-driven waste
Cost and frequency of failed runs, retries and rework.

*Mitigations were prioritised to directly address sources of high SLA risk and reduce volatility in the Legal and FP&A domains where errors carry disproportionate business risk. See appendix for details.

Copyright © 2026 Accenture. All rights reserved.

Mitigations lower cost per successful outcome

Mitigations significantly increase the probability of success and reduces failure-driven waste, therefore lowering the cost per quality outcome. Over time, this converts a small quality assurance premium into large lifecycle savings and more predictable SLA performance.

Metric	Squad 1 (No Mitigations)	Squad 2 (Mitigations)
Token consumption cost per run	$8.16	$10.96
Success rate	45%	**93%**
Token consumption cost per successful run	$18.13*	**$11.78**
SLA Compliance	87%	**95%**

35% lower
cost per successful run

Despite higher per-run cost, the significantly higher success rate reduces cost per quality outcome.

2x higher
success rate

Mitigations improve reliability and reduce failure-driven waste (retries, rework and overrides).

+8pp
SLA compliance

More predictable performance and greater operational confidence.

Improved
cost predictability

Fewer failures and rework result in more stable, predictable run costs over time.

A small quality assurance premium delivers large lifecycle savings.
Investing in mitigations reduces the total cost per successful outcome, while improving SLA performance and business confidence.

* This value does not factor in lifecycle costs that can significantly distort this metric. For example, rework and overrides, incident recovery, operational overheads, and quality penalties.

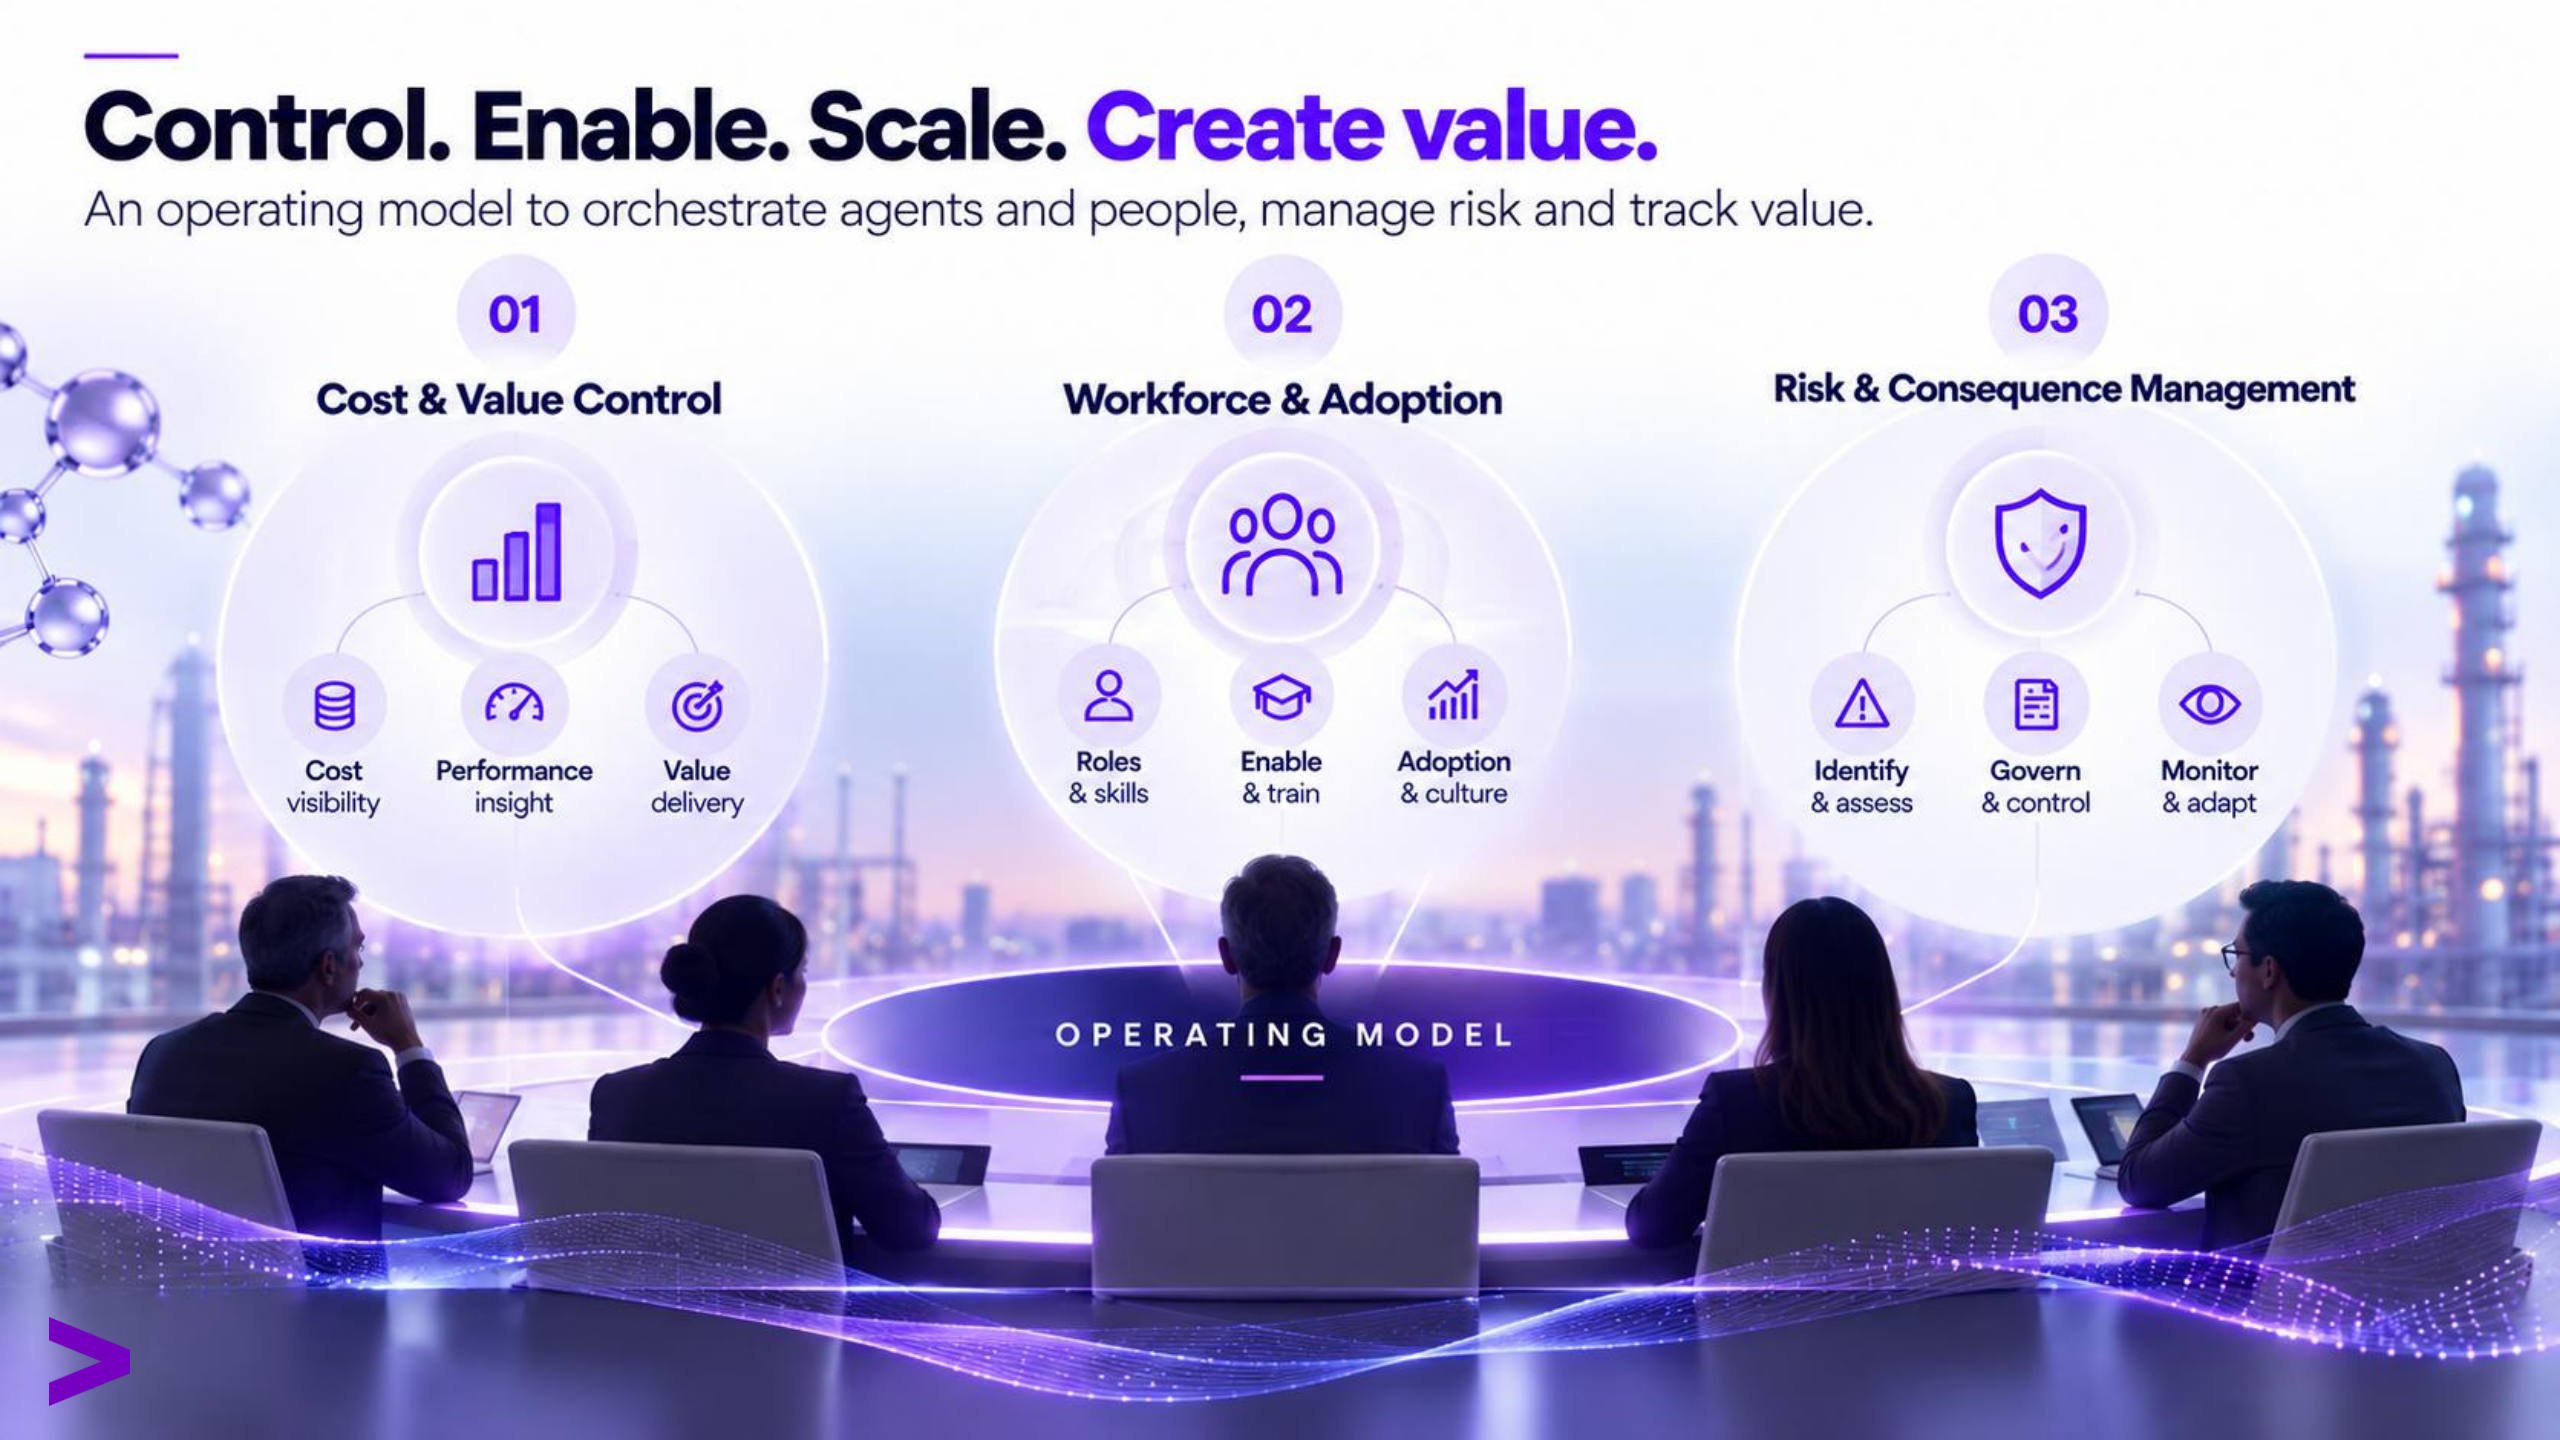
Copyright © 2026 Accenture. All rights reserved.



See it brought to life now!

Control. Enable. Scale. Create value.

An operating model to orchestrate agents and people, manage risk and track value.



So what do you have to keep central across the enterprise to enable this?

The method of managing agentic metrics

Central TMO tracking for agentic reinvention

Established terminology around transformation

- Agent
- Task
- Orchestration
- Reinvention
- Human-in-the-loop
- TMO
- ROI

Dashboards of the cost and value agent, squad and across the organisation

CoE of people with this hybrid skill set



This impacts the future of work and roles!

1 Leaders need training on how to control agents + humans

Equip leaders with the skills to orchestrate human and agent workforces, set guardrails and make effective decisions.



2 You need these performance dashboards

Real-time visibility across agents and humans to monitor performance, cost, risk and impact at scale.



3 You need more social scientists

Bring social scientists into the design, deployment and governance of agentic systems to understand human behaviour, trust, adoption and societal impact.



CHEMICAL CONVENTION
21-23 OCTOBER 2026

Questions?



cefic

INTERNATIONAL COUNCIL OF CHEMICAL ASSOCIATIONS

CHEMICAL CONVENTION
21-23 OCTOBER 2026

Thank you for your attention



cefic

INTERNATIONAL COUNCIL OF CHEMICAL ASSOCIATIONS