



**INFORMATION
FACTORY**

Agentic AI in Context

Executive Brief August 2025

Contents

Agentic AI in context



03 Agentic - the next leap in enterprise IT

07 Reference architecture

05 Reimagine operations

08 ROI model

06 Governance & strategic challenges

Agentic AI represents the next leap in enterprise IT






From doing the work to driving the outcome



What is Agentic AI?

Agentic AI is a goal-driven artificial intelligence that doesn't just answer questions.

It takes action by:

-  Understanding goals
-  Breaking them into steps
-  Using tools, APIs, and systems
-  Adapting and making decisions
-  Completing work autonomously

Think of it as a digital colleague, not just a tool.

Why it matters now

- Paradigm shift: From tool usage → **goal execution**
- Enterprise impact: Automates the *Cognitive Execution Layer* — knowledge work that applies rules, executes workflows, and makes decisions using ERP, CRM, ITSM, and analytics systems.



Expected productivity gains

- 2–10× in knowledge work
- 50–100× in fully automatable tasks


Agentic role

- Cognitive Executor
- Workforce Multiplier
- System Orchestrator
- Next-Gen Interface
- Economic Transformer

What it does

- Transforms intent into execution
- Eliminates repetitive cognitive tasks
- Connects apps, APIs, workflows & people
- Replaces forms with natural conversation
- Drives new value & scale

Agentic models with high impact require autonomy and trust



The window for advantage is closing fast

Agentic will become rapidly commoditized. The real advantage lies in how early and deep it's embedded.



Agentic AI is the **industrialization of cognition** and the next layer of enterprise evolution after SaaS and APIs.

Just as machines transformed physical labor, it **will reshape knowledge work**, team structures, and value chains — creating **intent-driven** economies, agent-augmented teams, and **adaptive** enterprise **ecosystems**



Key actions for leaders

- A** Assess: Identify agent-ready workflows and processes.
- G** Govern: Enable APIs, orchestration, and governance.
- E** Experiment: Test in high-impact areas (IT ops, reporting, customer service).
- N** Navigate: Scale with monitoring and security controls
- T** Tune : Use ROI tracking to optimize and drive further automation.

Proof Points

- Incident triage: 42% MTTR reduction, 1,800 hrs/year saved.
- Regression testing: 75% manual QA reduction, release cycles halved.



Risks & governance priorities

1 Access control & least-privilege principles

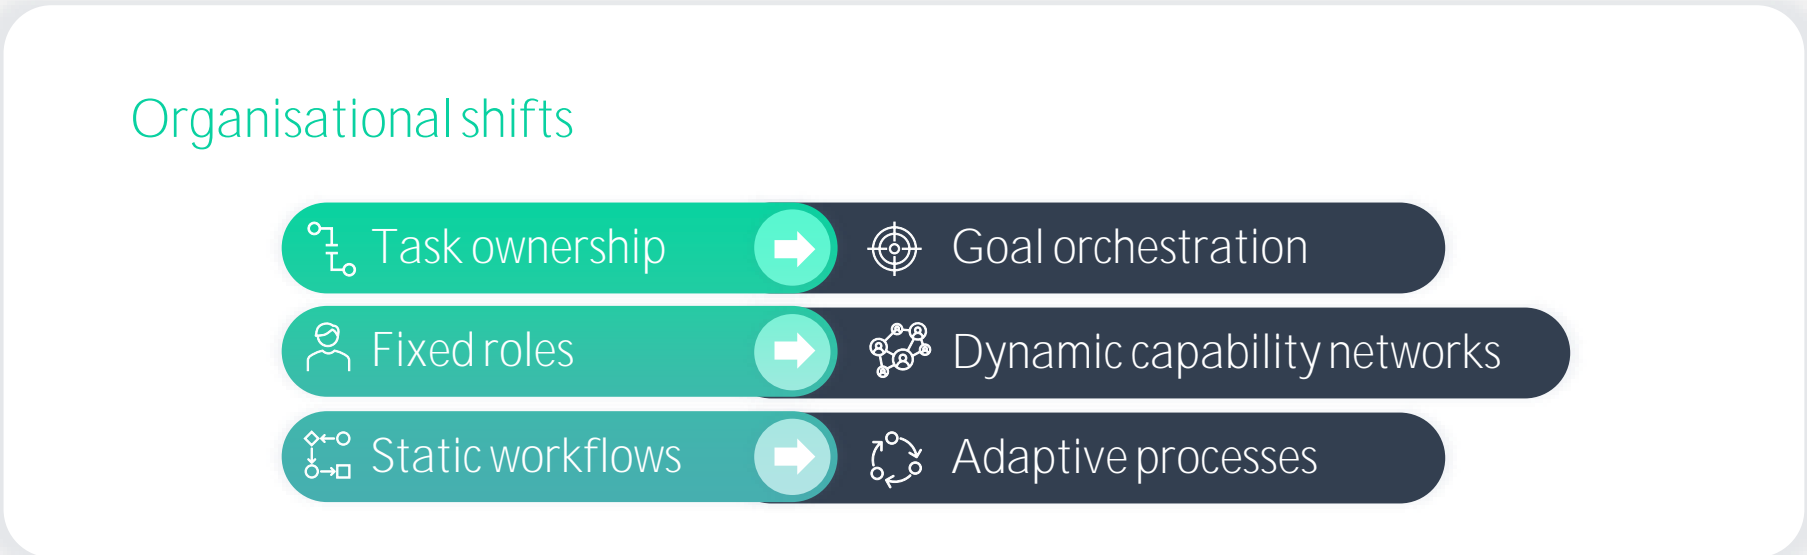
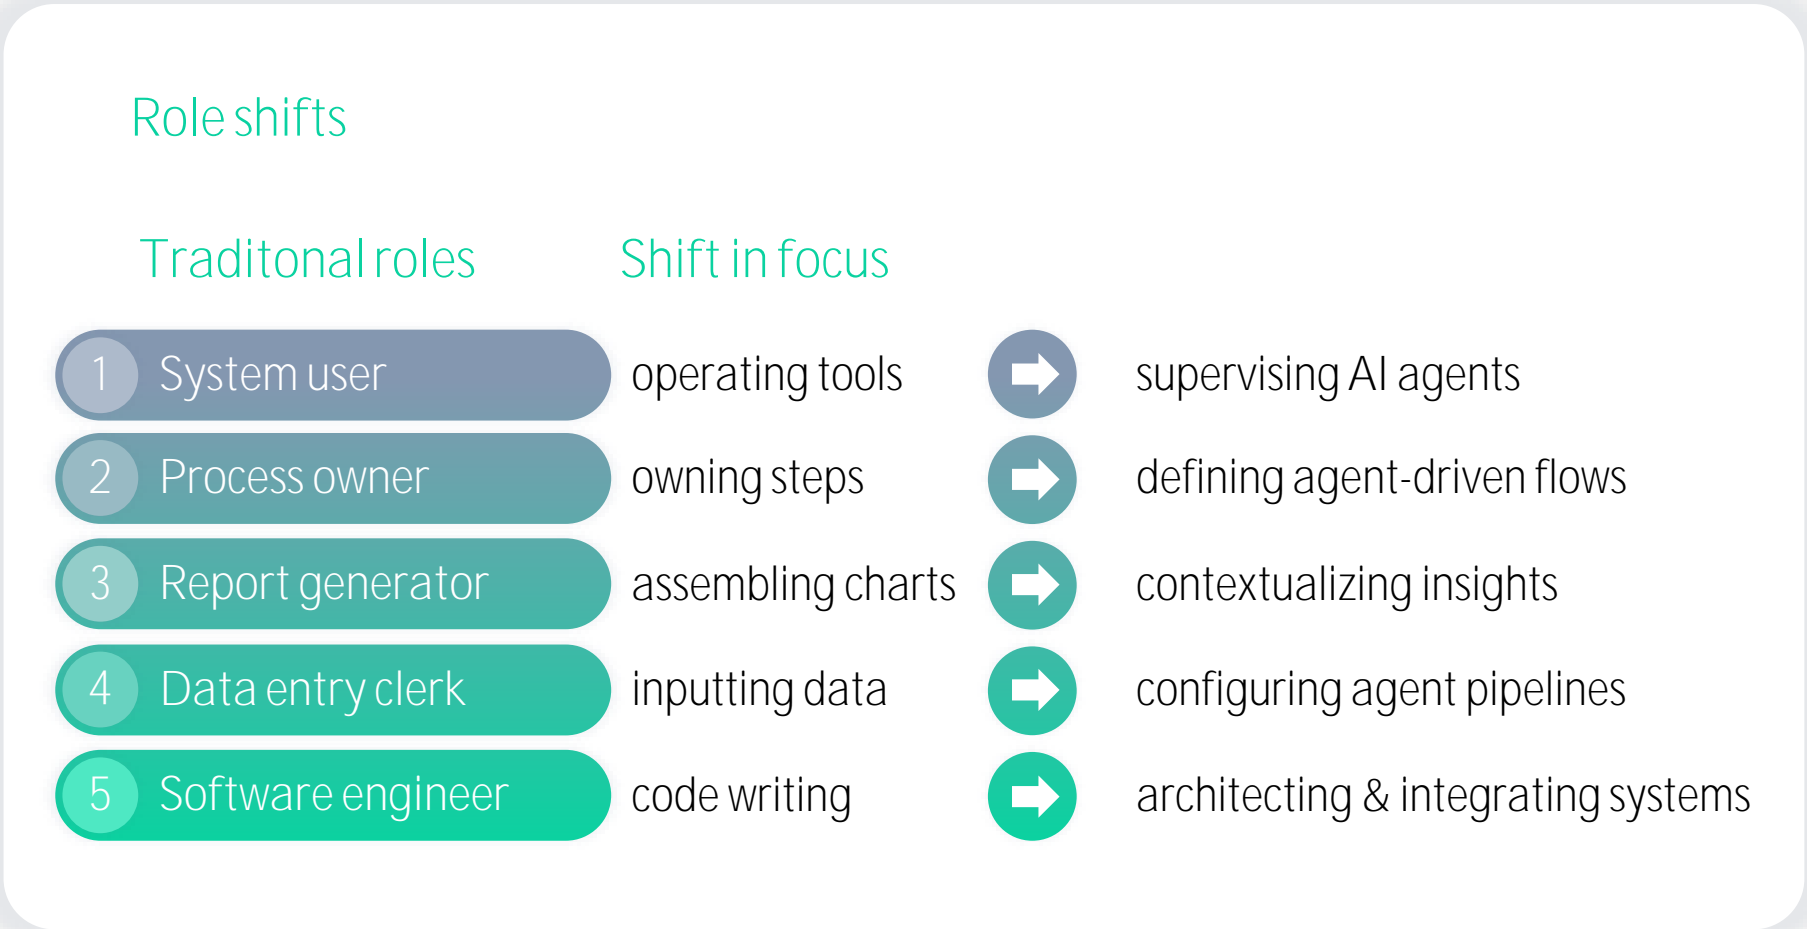
2 End-to-end audit logging

3 Sandboxed execution environments

4 Workforce reskilling for AI oversight

To leverage Agentic AI, we will have to reimagine operations

Shifts in paradigms, roles, and organizational design



...and put the right governance in place to address strategic challenges

Making AI safe, trusted, and human-centered

Governance

Who governs AI actions and outcomes?

How do we enforce ethics and risk limits at scale?

How do we guarantee auditability and transparency?

Trust

How do we ensure traceability and alignment with values?

How do we prove reliability and fairness?

How do we prevent reputational risk from errors or bias?

Regulation

Who holds liability for autonomous actions?

How do we stay ahead of evolving AI regulations?

How do we prove cross-border compliance in real time?

Workforce

How do we reskill displaced workers?

How do we prevent skill erosion as agents scale?

How do we enable productive human/agent collaboration?



Technical Governance

Access Control

Role-based agent permissions with least privilege

Auditability

Comprehensive logs for all decisions and actions

Security

Sandboxing, vaulting, anomaly detection

Data Provenance

Lineage tracking for agent-managed data



KPIs for CIO/CTO Measurement

1

Mean Time to Resolution (MTTR)

2

% of automated ticket closures

3

Developer hours saved per month

4

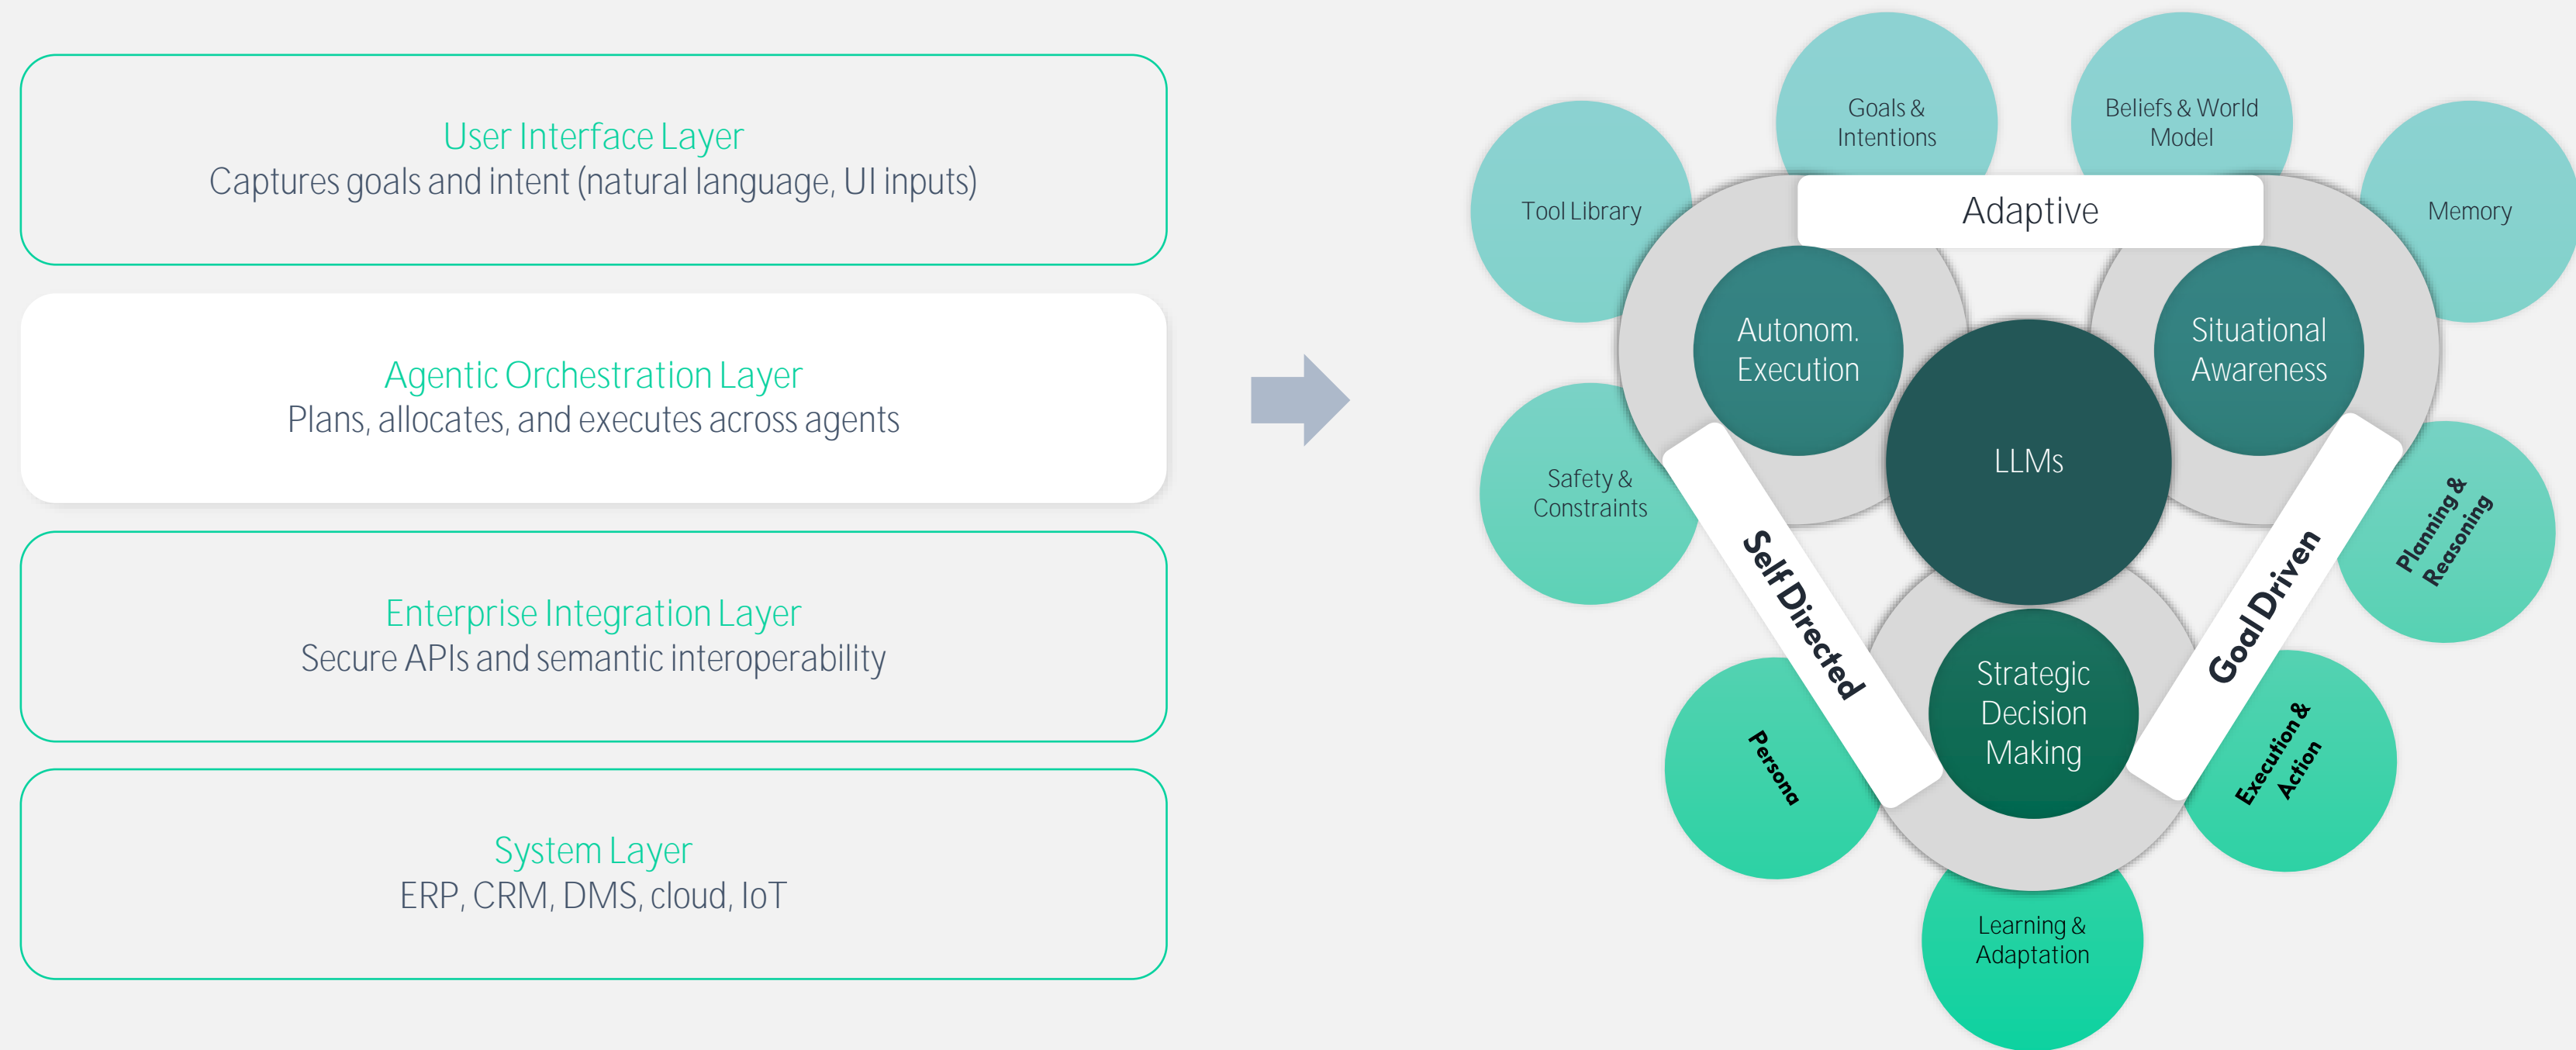
SLA compliance improvement (%)

5

Reduction in manual data-entry errors

Reference architecture for Agentic AI deployment

Agentic AI sits above core systems — orchestrating tasks, cutting latency, and freeing talent for higher-value work.



ROI Model for Agentic AI

Even in their current, limited scope, Agentic deployments deliver impact — an early signal of the exponential value ahead



Formula

Monthly Savings = (Hours automated × Hourly rate) × Automation %
Payback Period = Deployment cost ÷ Monthly savings

Example : IT Operations

- ➡ Manual workload: 800 hrs/month × \$60/hr
- ➡ 50% automated = 400 hrs saved → **\$24K/month**
- ➡ Deployment cost: \$120K → **Payback in 5 months**

Indirect Benefits

- ➡ Faster SLA resolution → **higher customer satisfaction**
- ➡ Less burnout from repetitive work
- ➡ More **innovation capacity** for skilled staff

IT operations: incident triage

Background
A global financial services firm averaged 6.5 hours MTTR for incidents, draining engineering resources.

Agentic AI Deployment

- Integrated orchestration into ServiceNow ITSM.
- Automated incident classification, log retrieval, and runbook triggering.

Results

- MTTR reduced by 42% (6.5 → **3.8** hrs)
- 55% of P3/P4 tickets auto-resolved
- 1,800 hours/year freed (~\$216K cost avoidance)
- SLA compliance up 15%

Software development: regression testing

Background
Manual nightly regression tests consumed 200+ dev-hours/month and delayed releases.

Agentic AI Deployment

- Used LangChain integrated with CI/CD.
- Automated test script generation, execution, bug logging, and priority tagging.

Results

- QA effort cut by 75% (150 hrs/month saved)
- Regression cycle time dropped from 2 days → 6 hours
- Shifted from monthly to bi-weekly releases
- Increased developer satisfaction

Let's stay connected



Cosima Geyer
Principal | Head of Management Consulting

cosima.geyer@information-factory.com
Phone: +41 (0)79 705 12 01

Information Factory AG
Foerrlibuckstrasse 10 • CH-8005 Zurich