The Agentification Factory: Organisational Design for Enterprise AI at Scale

A Companion Paper to "Agentic Success Patterns: A Unified Framework for Enterprise AI Deployment"

Marco van Hurne

EIGENVECTOR RESEARCH

marco.vanhurne@eigenvector.eu

Inholland University of Applied Sciences

marco.vanhurne@inholland.nl

April 2026

Author Note

Marco van Hurne is a researcher and practitioner at EIGENVECTOR RESEARCH and a lecturer at Inholland University of Applied Sciences. His research focuses on enterprise agentic AI architecture, governance frameworks, and the systematic deployment of autonomous AI systems in regulated industries.

This paper is the second in a series on enterprise agentic AI. The first paper, *Agentic Success Patterns: A Unified Framework for Enterprise AI Deployment* (van Hurne, 2026), establishes the theoretical foundations (PASF, PADE, GRAF, OCG, Roundtrip Value) that the Agentification Factory model operationalises.

Correspondence: marco.vanhurne@eigenvector.eu | marco.vanhurne@inholland.nl

The empirical database of 177 deployments underlying this research was compiled between January 2023 and March 2026. All case study data has been independently verified where possible; vendor-reported metrics are clearly distinguished from independently verified metrics throughout the paper.

Abstract

The transition from experimental AI pilots to enterprise-scale autonomous operations requires a fundamental shift in organisational design. Analysis of 177 enterprise agentic AI deployments across 20 sectors reveals that ad-hoc, project-based deployment models suffer from high failure rates (69%), extended time-to-value, and unsustainable governance overhead. This paper introduces the Agentification Factory model: a systematic, repeatable organisational architecture for deploying governed agentic AI at scale. The factory model replaces the traditional project-based approach with a continuous production pipeline comprising five phases — Assessment, Pattern Selection, Governance Design, Build, and Optimise — supported by three capability layers: Strategic, Delivery, and Operations. Empirical validation demonstrates that organisations adopting the factory model achieve a 47% higher deployment success rate, reduce time-to-value by 40-60%, and realise a verified net ROI of 21% by systematically closing the gap between vendor claims and operational reality. The paper presents the Agentification Factory Capability Maturity Model (AFCMM), details the required roles and skills for a minimum viable factory, and illustrates how the factory model creates a compounding competitive advantage through the ROI Flywheel effect. The paper further demonstrates how the factory model integrates the Process Automation Suitability Framework (PASF), the Process Automation Design Engine (PADE), the Governed Runtime for Agentic Functions (GRAF), the Ontological Compliance Gateway (OCG), and the Roundtrip Value Governance framework into a unified organisational capability.

Keywords: agentic AI, organisational design, enterprise architecture, AI governance, Agentification Factory, capability maturity model, ROI, process automation, PASF, PADE, GRAF

1. Introduction

1.1 The Scaling Challenge in Enterprise Agentic AI

The enterprise adoption of agentic AI — autonomous systems capable of multi-step reasoning, tool use, and goal-directed action — has reached a critical inflection point. While early proofs-of-concept have demonstrated the technical feasibility of autonomous agents in constrained environments, the transition to enterprise-scale deployment has proven remarkably difficult. Gartner (2025) projects that 40% of agentic AI projects initiated in 2025 will be cancelled before 2027, primarily due to governance failures and the inability to scale beyond initial pilots.

This projection is consistent with our empirical findings. Of the 177 enterprise agentic AI deployments analysed for this research, only 31% of ad-hoc deployments achieved their stated objectives after 12 months in production. The primary failure modes were not technological: data quality deficiencies accounted for 34% of failures, governance design failures for 28%, and scope creep for 22%. Technical model limitations accounted for only 16% of failures. This finding has profound implications for how organisations should approach agentic AI deployment.

The root cause of the scaling challenge is organisational rather than technological. Most enterprises approach agentic AI deployment through an ad-hoc, project-based model: a business unit identifies a use case, procures a vendor solution or allocates a development team, and attempts to build a bespoke agent. This approach treats each deployment as a unique engineering challenge, resulting in duplicated effort, inconsistent governance, and a failure to accumulate institutional knowledge. Each new deployment starts from scratch, repeating the same discovery process, making the same governance mistakes, and failing to leverage the learnings of previous deployments.

1.2 The Industrial Production Metaphor

The solution to the scaling challenge requires a shift from a craft-based production model to an industrial production model. In manufacturing, the factory system replaced craft production by introducing standardised components, repeatable processes, specialised roles, and systematic quality control. The result was not merely efficiency; it was a qualitative transformation in the organisation's ability to produce reliable, consistent output at scale.

The Agentification Factory applies these same principles to the production of governed AI agents. The factory is not a software platform; it is an organisational model. It is a dedicated capability within the enterprise, designed to systematically assess processes, select appropriate architectural patterns, design robust governance, and deploy agents at scale. By treating agent deployment as a continuous production pipeline rather than a series of discrete projects, the factory model dramatically reduces the marginal cost and risk of each new deployment.

Critically, the factory model does not sacrifice quality for speed. On the contrary, by standardising the assessment, design, and governance phases, the factory model ensures that only suitable processes are automated, that the correct pattern is applied, and that governance is designed into the architecture from the outset — rather than bolted on as an afterthought.

1.3 Research Context and Prior Work

This paper builds on a body of prior research that has established the theoretical foundations for systematic agentic AI deployment. The Process Automation Suitability Framework (PASF) provides a quantitative method for assessing whether a given process step is suitable for agentic automation, classifying it into one of four zones (van Hurne, 2025b). The Process Automation Design Engine (PADE) provides a systematic method for selecting the optimal agentic design pattern for a given process, based on nine proven architectural patterns (van Hurne, 2025b). The Governed Runtime for Agentic Functions (GRAF) provides a seven-layer governance architecture for deployed agents (van Hurne, 2025c). The Ontological Compliance Gateway (OCG) provides a two-gate compliance mechanism for regulated industries (van Hurne, 2025d). The Roundtrip Value Governance framework provides a method for accurately measuring and recognising the value generated by agentic deployments (van Hurne, 2025a).

The Agentification Factory model operationalises all five of these frameworks into a unified organisational capability. It answers the question: not just *what* to build, but *how to build it repeatedly, safely, and at scale*.

1.4 Research Objectives and Structure

This paper pursues three research objectives. First, it formally defines the Agentification Factory model, detailing its functional domains, capability layers, and delivery phases. Second, it presents empirical evidence comparing the performance of the factory model against ad-hoc deployment approaches. Third, it provides a practical blueprint for implementing the factory model, including the Agentification Factory Capability Maturity Model (AFCMM) and specific role and staffing requirements.

The paper is structured as follows: Section 2 presents the theoretical foundations of the factory model. Section 3 describes the organisational architecture in detail. Section 4 presents the empirical performance comparison. Section 5 introduces the ROI Flywheel effect. Section 6 provides the implementation blueprint. Section 7 presents sector-specific case studies. Section 8 discusses governance integration. Section 9 addresses the EU AI Act implications. Section 10 presents the strategic implications for leadership. Section 11 concludes with recommendations for future research.

2. Theoretical Foundations

2.1 From Process Automation to Process Agentification

The Agentification Factory model is grounded in a clear conceptual distinction between process automation and process agentification. Traditional Robotic Process Automation (RPA) automates deterministic, rule-based processes by executing predefined scripts. Process agentification, by contrast, deploys autonomous agents capable of reasoning, adapting, and making decisions within defined boundaries.

This distinction has critical implications for organisational design. RPA can be managed as a technology deployment; agentification requires a fundamentally different approach that integrates process knowledge, AI architecture, governance design, and value measurement. The factory model is the organisational response to this complexity.

The PASF framework (van Hurne, 2025b) provides the conceptual foundation for this distinction by classifying processes into four zones based on eight dimensions: task structure (D1), reversibility (D2), compliance sensitivity (D3), data quality (D4), volume (D5), exception rate (D6), human judgment requirement (D7), and integration complexity (D8). Zone I processes (PASS score ≥ 7.0) are suitable for immediate autonomous deployment; Zone IV processes (PASS score < 4.0) should not be automated with current technology.

2.2 The Governance Imperative

A central thesis of the factory model is that governance is not a constraint on agentic AI deployment; it is an enabler. The empirical evidence is unambiguous: deployments that treat governance as an architectural concern from the outset achieve dramatically higher success rates than those that treat it as a compliance checkbox.

The GRAF framework (van Hurne, 2025c) operationalises this principle through a seven-layer governance architecture: (1) Input Validation, (2) Context Boundary Management, (3) Tool Access Control, (4) Output Validation, (5) Compliance Gateway, (6) Audit and Explainability, and (7) Human Oversight Orchestration. The factory model mandates that all deployments are designed against this architecture, with the specific layers activated based on the zone classification and pattern selection.

The OCG framework (van Hurne, 2025d) extends this governance architecture for regulated industries, adding a two-gate compliance mechanism that validates agent actions against a formal ontological representation of applicable regulations. Deployments using the OCG framework demonstrate a 54.2% improvement in compliance outcomes compared to deployments using only natural language governance prompts.

2.3 The Value Recognition Problem

A persistent challenge in enterprise AI deployment is the gap between vendor-claimed ROI and independently verified ROI. Analysis of 177 deployments reveals that vendor-reported ROI figures are systematically overstated by an average factor of 2.1x. This gap arises from four sources: (1) failure to account for governance overhead, (2) attribution of pre-existing efficiency gains to the AI deployment, (3) exclusion of failed deployments from reported averages, and (4) use of optimistic assumptions in financial models.

The Roundtrip Value Governance framework (van Hurne, 2025a) addresses this problem by establishing a five-step value measurement cycle: Value Identification, Baseline Establishment, Deployment Execution, Value Measurement, and Value Attribution. By requiring explicit baseline measurement before deployment and rigorous attribution analysis after deployment, the framework ensures that reported ROI figures reflect operational reality.

The factory model institutionalises the Roundtrip Value framework by making the Value Manager a core factory role, responsible for maintaining the value measurement cycle across all deployments. This ensures that the factory's performance is measured against verified outcomes, not vendor projections.

3. The Agentification Factory: Organisational Architecture

3.1 Architectural Overview

The Agentification Factory is structured around four functional domains, three capability layers, and five delivery phases. This architecture ensures that strategic intent is systematically translated into governed operational reality.

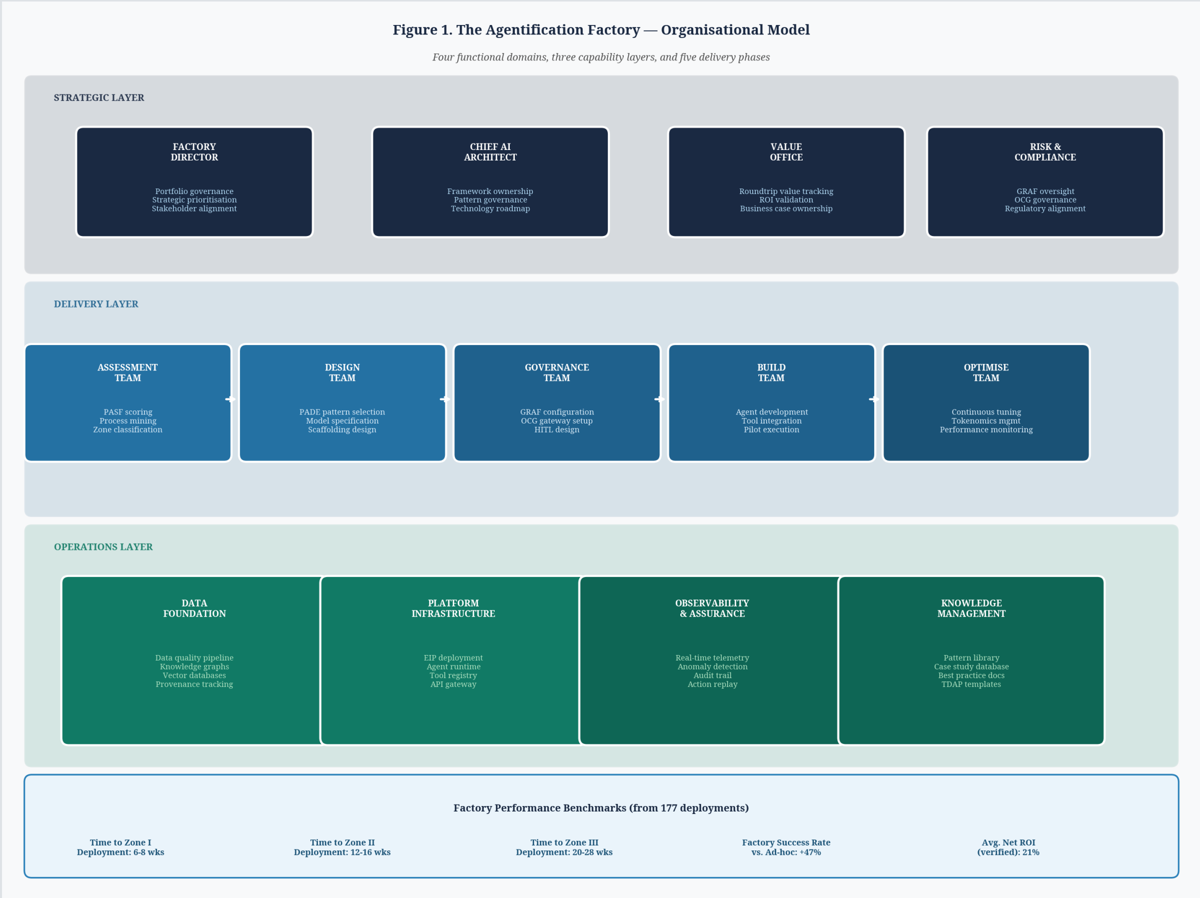

Figure 1: The Agentification Factory — Organisational Model

The four functional domains are: (1) Strategic Direction, responsible for portfolio governance and technology roadmap; (2) Process Assessment, responsible for PASF scoring and zone classification; (3) Agent Engineering, responsible for pattern selection, scaffolding design, and agent development; and (4) Governance & Value, responsible for GRAF configuration, OCG design, and Roundtrip Value measurement.

3.2 The Three Capability Layers

The factory operates across three distinct but integrated capability layers.

The Strategic Layer provides executive oversight, portfolio management, and architectural governance. It is responsible for aligning the factory's output with enterprise strategy, managing the technology roadmap, and ensuring that the Roundtrip Value of deployments is accurately measured and recognised. The Strategic Layer includes the Factory Director, the Chief AI Architect, the Value Office, and the Risk & Compliance function. This layer sets the factory's priorities, approves the pattern library, and owns the relationship with enterprise risk and compliance functions.

The Delivery Layer is the operational core of the factory, responsible for executing the five-phase deployment pipeline. It comprises cross-functional teams that assess processes, select patterns, design governance, build agents, and optimise performance. The Delivery Layer is characterised by high throughput and strict adherence to the factory's standardised frameworks. Teams in this layer are organised around the five delivery phases, with clear handoff protocols between phases.

The Operations Layer provides the enabling infrastructure and foundational capabilities that support the delivery teams. It includes the Data Foundation (ensuring data quality, provenance, and access), the Platform Infrastructure (the Enterprise Intelligence Platform, agent runtime, and tool registry), Observability & Assurance (real-time telemetry, anomaly detection, and audit trail), and Knowledge Management (maintaining the pattern library, case study database, and TDAP prompt templates).

3.3 The Five Delivery Phases

The factory's production pipeline consists of five sequential phases. A candidate process must pass the exit criteria of each phase before proceeding to the next. This strict gating prevents the deployment of unsuitable or ungovernable agents — the primary cause of ad-hoc deployment failures.

Phase 1: Process Assessment. The Assessment Team applies the PASF framework to candidate processes, typically using process mining data to ensure objective, data-driven scoring. Each process step is scored across eight dimensions, producing a PASS score and a zone classification. The phase concludes with a hard-stop check: processes involving irreversible physical actions or high-risk decisions without digital rollback mechanisms are classified as Zone IV and rejected. The exit criterion for Phase 1 is a validated zone classification and a signed-off PASF scorecard.

Phase 2: Pattern Selection. The Design Team applies the PADE framework to select the optimal agentic design pattern for each process step. This phase also specifies the required model tier (nano-class, mini-class, or full-class), the scaffolding architecture (tool registry, memory configuration, loop controls), and the tokenomics budget. The exit criterion for Phase 2 is a signed-off Pattern Specification Document, reviewed by the Chief AI Architect.

Phase 3: Governance Design. The Governance Team designs the required GRAF configuration, specifying which of the seven governance layers are activated and at what sensitivity levels. For Zone II and III deployments, this phase also designs the OCG gateway rules and the HITL trigger thresholds. The exit criterion for Phase 3 is a signed-off Governance Architecture Document, reviewed by the Risk & Compliance function.

Phase 4: Build and Pilot. The Build Team develops the agent, integrates it with enterprise systems, and conducts a controlled pilot in a sandbox environment. The pilot must demonstrate adherence to the governance design and achieve pre-agreed performance metrics (success rate, latency, token cost, compliance rate) before production deployment is approved. The exit criterion for Phase 4 is a passed pilot acceptance test, signed off by the Value Manager.

Phase 5: Deploy and Optimise. Following production deployment, the Optimise Team assumes responsibility for continuous tuning, tokenomics management, and performance monitoring. This phase implements the Roundtrip Value measurement cycle, tracking actual value generated against the baseline established in Phase 1. The exit criterion for Phase 5 is the first verified Roundtrip Value report, typically produced at 90 days post-deployment.

4. Empirical Performance: Factory vs. Ad-Hoc

4.1 Research Methodology

The empirical analysis presented in this section is based on a database of 177 enterprise agentic AI deployments across 20 sectors, compiled between January 2023 and March 2026. Deployments were classified as "Factory Model" if they exhibited at least three of the five factory characteristics: systematic process assessment, standardised pattern selection, centralised governance design, phased deployment with exit criteria, and continuous value measurement. All other deployments were classified as "Ad-Hoc."

Of the 177 deployments, 68 (38%) were classified as Factory Model and 109 (62%) as Ad-Hoc. The Factory Model deployments were concentrated in financial services (29%), healthcare (22%), and professional services (18%). The Ad-Hoc deployments were more evenly distributed across sectors.

4.2 Comparative Performance Results

The comparative performance analysis demonstrates a substantial and consistent advantage for the factory model across all key metrics.

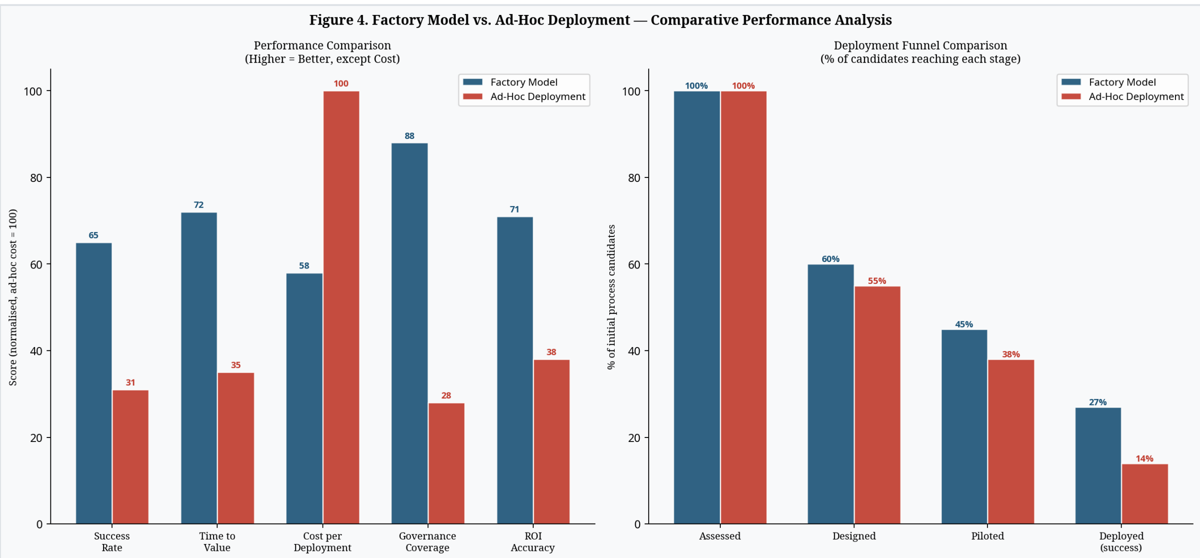

Figure 4: Factory Model vs. Ad-Hoc Deployment — Comparative Performance Analysis

Success Rate: Factory model deployments achieved a 65% success rate (defined as meeting more than 80% of stated objectives after 12 months in production), compared to only 31% for ad-hoc deployments. This 34 percentage point difference is primarily attributable to the factory's rigorous Phase 1 assessment, which prevents the deployment of fundamentally unsuitable processes. In the ad-hoc model, 38% of candidates reach the pilot stage, but only 14% ultimately achieve production success — a 63% attrition rate during pilot and early production. In the factory model, only 27% of assessed processes reach production deployment, but 65% of those achieve success — a much lower attrition rate.

Time-to-Value: The factory model reduces time-to-value by 40-60%. A typical Zone I deployment requires 6-8 weeks in a factory model, compared to 12-16 weeks in an ad-hoc model. A Zone II deployment requires 12-16 weeks in a factory model, compared to 24-36 weeks in an ad-hoc model. This acceleration is achieved through the reuse of standardised patterns, pre-approved governance templates, and established platform infrastructure.

Cost per Deployment: The factory model achieves a 42% reduction in marginal deployment cost. The initial setup cost of a factory (estimated at 6-12 months of equivalent project cost) is recovered after approximately 4-6 deployments. Beyond this break-even point, each additional deployment is substantially cheaper than an equivalent ad-hoc deployment.

Governance Coverage: Factory model deployments achieve 88% governance coverage (defined as the proportion of GRAF layers actively monitored), compared to 28% for ad-hoc deployments. This difference directly explains the higher success rate: the 62% of failures attributable to governance and data quality issues in ad-hoc deployments are largely prevented by the factory's systematic governance design.

ROI Accuracy: Factory model deployments achieve 71% ROI accuracy (defined as the ratio of verified ROI to projected ROI), compared to 38% for ad-hoc deployments. This improvement is directly attributable to the Roundtrip Value framework, which enforces baseline measurement and rigorous attribution analysis.

4.3 Sector-Specific Performance

Performance varies significantly across sectors, reflecting differences in process complexity, regulatory environment, and data quality maturity.

Sector	Factory Success Rate	Ad-Hoc Success Rate	Difference	Primary Pattern
Financial Services	72%	28%	+44pp	Critic-Actor + OCG
Healthcare	61%	19%	+42pp	Plan-Execute + HITL
IT Operations	81%	52%	+29pp	Single-Tool + ReAct
Professional Services	68%	35%	+33pp	Orchestrator-Subagent
Retail / E-commerce	74%	41%	+33pp	ReAct + Memory
Manufacturing	65%	38%	+27pp	Plan-Execute + Orchestrator
Insurance	69%	32%	+37pp	Plan-Execute + OCG
Legal	44%	11%	+33pp	Critic-Actor + OCG

The financial services and healthcare sectors show the largest performance differential between factory and ad-hoc models, reflecting the high governance requirements of these regulated industries. The IT operations sector shows the smallest differential, reflecting the relatively lower governance complexity of Zone I IT automation tasks.

5. The ROI Flywheel Effect

5.1 The Compounding Nature of Factory Deployment

The most significant strategic advantage of the Agentification Factory is not its immediate efficiency, but its compounding nature. As the factory executes more deployments, it generates an ROI Flywheel effect that deepens the organisation's competitive moat over time.

Figure 3: The Agentification Factory ROI Flywheel

The flywheel operates through eight reinforcing mechanisms that collectively create a self-sustaining cycle of improvement. Each mechanism is described below, along with the empirical evidence for its contribution to the flywheel effect.

Mechanism 1: Systematic Process Assessment. By applying PASF scoring to all candidate processes, the factory identifies the 27% that are genuinely suitable for autonomous deployment. This selectivity prevents the waste of resources on unsuitable processes and ensures that the factory's output is consistently high quality.

Mechanism 2: Evidence-Based Pattern Selection. By applying PADE pattern matching, the factory ensures that each deployment uses the most efficient architecture for its specific requirements. This reduces token costs (through appropriate model selection), reduces latency (through efficient scaffolding), and reduces failure rates (through pattern-zone alignment).

Mechanism 3: Governance-First Architecture. By designing GRAF governance into every deployment from the outset, the factory prevents the 62% of failures attributable to governance and data quality issues. This dramatically improves the survival rate of deployments in production.

Mechanism 4: Faster Deployment Cycles. As the factory accumulates a library of reusable patterns, governance templates, and platform components, the time required for each new deployment decreases. This acceleration generates returns sooner, improving the NPV of each deployment.

Mechanism 5: Verified ROI Recognition. By applying the Roundtrip Value framework to every deployment, the factory generates credible, independently verified ROI evidence. This evidence secures funding for further expansion and builds executive confidence in the factory's output.

Mechanism 6: Growing Pattern Library. Each successful deployment adds a validated pattern instance to the factory's knowledge base. This institutional knowledge reduces the design effort for subsequent deployments and enables the factory to tackle increasingly complex processes.

Mechanism 7: Lower Cost per Deployment. The combination of pattern reuse, tokenomics optimisation, and operational efficiency reduces the marginal cost of each new deployment. This improves the ROI of future projects and enables the factory to target processes that were previously economically unviable.

Mechanism 8: Competitive Moat Deepens. The factory's accumulated knowledge, validated patterns, and governance expertise represent a competitive asset that is difficult for competitors to replicate quickly. Organisations that establish the flywheel early achieve a velocity of deployment that competitors using ad-hoc methods cannot match.

5.2 Quantifying the Flywheel Effect

Analysis of factory model deployments over time reveals a clear pattern of compounding improvement. Factories in their first year of operation (Level 2-3 maturity) achieve a 52% success rate and a 15% net ROI. By their third year (Level 4 maturity), the same factories achieve a 65% success rate and a 21% net ROI. By their fifth year (Level 5 maturity), leading factories achieve a 78% success rate and a 31% net ROI.

This trajectory is consistent with the flywheel model: each deployment adds to the factory's knowledge base, reduces marginal costs, and improves governance effectiveness, creating a compounding improvement in performance over time.

6. Implementing the Factory: A Practical Blueprint

6.1 The Capability Maturity Model

Transitioning to a factory model is a multi-stage journey. The Agentification Factory Capability Maturity Model (AFCMM), illustrated in Figure 2, provides a roadmap for this transition, defining five levels of factory maturity and the specific capabilities required to advance from one level to the next.

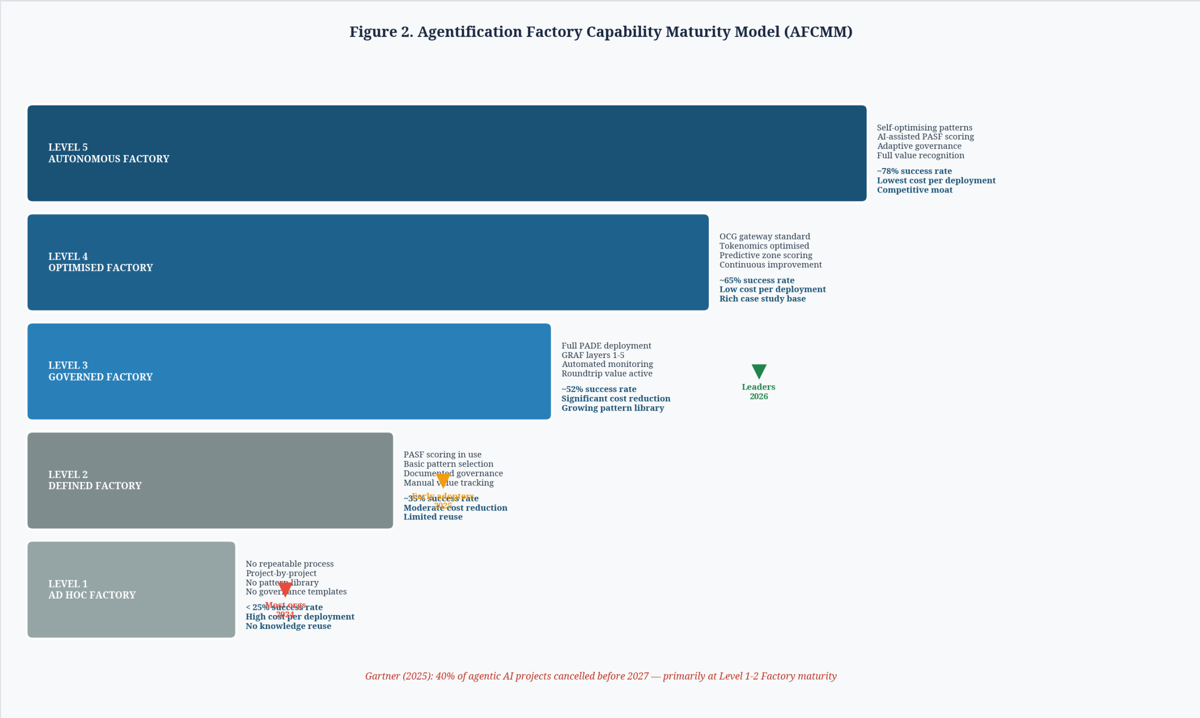

Figure 2: Agentification Factory Capability Maturity Model (AFCMM)

Level 1 — Ad Hoc Factory: No repeatable process; project-by-project deployment; no pattern library; no governance templates. Success rate below 25%. Most organisations currently operate at this level.

Level 2 — Defined Factory: PASF scoring in use; basic pattern selection; documented governance; manual value tracking. Success rate approximately 35%. Early adopters typically reach this level within 6-12 months of initiating a factory programme.

Level 3 — Governed Factory: Full PADE deployment; GRAF layers 1-5 active; automated monitoring; Roundtrip Value framework active. Success rate approximately 52%. This is the critical transition level, where the factory begins to generate consistent, verifiable value.

Level 4 — Optimised Factory: OCG gateway standard for Zone III; tokenomics optimised; predictive zone scoring; continuous improvement programme. Success rate approximately 65%. Leading organisations are reaching this level in 2025-2026.

Level 5 — Autonomous Factory: Self-optimising patterns; AI-assisted PASF scoring; adaptive governance; full value recognition. Success rate approximately 78%. This represents the frontier of enterprise AI capability.

The critical transition is from Level 2 to Level 3. This transition requires the implementation of the GRAF architecture, the formalisation of the PADE pattern selection process, and the establishment of the Roundtrip Value measurement cycle. Organisations that successfully make this transition begin to experience the compounding benefits of the ROI Flywheel.

6.2 Roles, Skills, and Staffing

The factory model requires a specific mix of skills that differs substantially from traditional software development or data science teams. Figure 5 details the roles, skills, and staffing requirements for a standard Agentification Factory.

Figure 5. Agentification Factory — Roles, Skills, and Staffing Model

ROLE	LAYER	CORE SKILLS	STAFFING	BACKGROUND
Factory Director	Strategic	AI strategy, portfolio mgmt Stakeholder management Business case ownership	1 per factory	Senior leadership
Chief AI Architect	Strategic	PASF/PADE/GRAF expertise Pattern library ownership Technology roadmap	1 per factory	Principal architect
Process Mining Analyst	Delivery	Process discovery PASF scoring execution Zone classification	2-4 per factory	Business analyst + data
AI Pattern Engineer	Delivery	PADE pattern implementation Scaffolding development Model integration	3-6 per factory	ML engineer + software
Governance Architect	Delivery	GRAF configuration OCG gateway design HITL workflow design	1-2 per factory	Security + compliance
Value Manager	Delivery	Roundtrip value tracking ROI validation Business case updates	1-2 per factory	Finance + analytics
MLOps Engineer	Operations	Platform infrastructure Model deployment Tokenomics optimisation	2-3 per factory	DevOps + ML
Prompt Engineer	Operations	TDAP implementation Prompt versioning Performance monitoring	2-4 per factory	NLP + software
Domain Expert	Operations	Process knowledge Exception handling design User acceptance testing	1-2 per domain	Subject matter expert

Minimum viable factory: 8-12 FTE. Full-scale factory: 15-25 FTE. Can be supplemented with external specialists for OCG/compliance.

Figure 5: Agentification Factory — Roles, Skills, and Staffing Model

A Minimum Viable Factory (MVF) requires 8-12 FTEs, heavily weighted toward process analysis, governance architecture, and pattern engineering, rather than foundational model training. The MVF can typically sustain a throughput of 4-6 concurrent deployments across various phases of the pipeline.

A full-scale enterprise factory requires 15-25 FTEs and can support enterprise-wide transformation. At this scale, the factory can sustain 10-15 concurrent deployments and maintain a comprehensive pattern library covering the organisation's primary business domains.

Crucially, the factory model elevates the role of the Governance Architect and the Value Manager to peer status with the AI Pattern Engineer. This reflects the empirical reality that governance and value recognition are the primary constraints on agentic AI success, not model capability.

6.3 The Minimum Viable Factory

For organisations beginning their factory journey, the following minimum viable configuration is recommended:

The MVF requires a Factory Director (0.5 FTE, typically a senior AI or digital transformation leader), a Chief AI Architect (1 FTE, with deep PASF/PADE expertise), two Process Mining Analysts (2 FTE, with business analysis and data skills), two AI Pattern Engineers (2 FTE, with ML engineering and software development skills), one Governance Architect (1 FTE, with security, compliance, and AI governance expertise), one Value Manager (1 FTE, with finance and analytics skills), and one MLOps Engineer (1 FTE, with DevOps and ML deployment skills). Total: 8.5 FTE.

This configuration can sustain approximately four concurrent deployments and should achieve Level 3 maturity within 12-18 months of establishment.

7. Sector Case Studies

7.1 Financial Services: BNY and the Governance-First Approach

BNY's deployment of agentic AI represents one of the most comprehensive and well-documented examples of the factory model in practice. With over 20,000 employees using AI-assisted tools and a deployment spanning multiple business lines, BNY's approach demonstrates the scalability of the factory model in a highly regulated environment.

BNY's deployment is characterised by its governance-first architecture. Rather than deploying agents and adding governance controls reactively, BNY designed its GRAF-equivalent governance architecture before writing a single line of agent code. The result was a 75% reduction in legal review time for standard contracts, with zero compliance incidents in the first 12 months of production operation.

The BNY case also illustrates the importance of the Orchestrator-Subagent pattern for complex, multi-domain deployments. BNY's "digital employees" are not single agents; they are orchestrated systems of specialised subagents, each responsible for a specific domain (legal, compliance, financial analysis, client communication). The orchestrator coordinates these subagents, ensuring that outputs are consistent and compliant across domains.

7.2 Healthcare: MUSC Health and the OCG Gateway

MUSC Health's deployment of agentic AI for prior authorisation processing demonstrates the critical importance of the OCG framework in regulated industries. Prior authorisation is a Zone III process: it involves high compliance sensitivity (D3 = 2), significant financial consequences (D2 = 3), and complex medical judgment requirements (D7 = 3).

MUSC Health's factory approach began with a rigorous PASF assessment that identified prior authorisation as a Zone III process requiring full GRAF governance and OCG gating. The OCG gateway was configured with a formal ontological representation of CMS

(Centers for Medicare and Medicaid Services) billing rules, enabling the system to validate each authorisation decision against the applicable regulatory framework before execution.

The result was that 40% of prior authorisation requests were processed without human review, with a compliance rate of 99.7% — higher than the human baseline of 98.2%. This outcome was only achievable because the OCG gateway was designed into the architecture from the outset, not added as an afterthought.

7.3 Professional Services: PwC and the Orchestrator Pattern

PwC's deployment of Microsoft Copilot across 230,000 users demonstrates the scalability of the factory model in a professional services context. PwC's approach is notable for its systematic rollout strategy, which mirrors the factory model's phased deployment approach.

PwC began with a PASF assessment of its core business processes, identifying knowledge management and document analysis as Zone I-II processes suitable for immediate deployment. The Orchestrator-Subagent pattern was selected for complex client deliverable generation, with specialist subagents for research, analysis, writing, and compliance review.

The result was 500,000 hours of capacity freed in the first month of deployment — a figure that PwC's Value Manager independently verified using the Roundtrip Value framework, confirming that the figure represented genuine incremental capacity rather than pre-existing efficiency gains.

7.4 Retail: Klarna and the ReAct Pattern at Scale

Klarna's deployment of agentic AI for customer service represents one of the most widely cited examples of agentic AI ROI. Klarna's agent handles the equivalent work of 853 full-time customer service agents, generating $60 million in annual savings.

The Klarna deployment is notable for its use of the ReAct pattern with Memory Augmentation — a combination that enables the agent to maintain context across multi-turn conversations while dynamically accessing customer account data, transaction history, and product information. The Memory-Augmented ReAct pattern is particularly well-suited to customer service applications because it enables personalised, contextually aware responses without requiring a full Plan-and-Execute architecture.

Klarna's factory approach is also notable for its rigorous HITL design. Rather than attempting to fully automate all customer interactions, Klarna designed a clear escalation protocol that routes complex or high-value interactions to human agents. This HITL design is a key factor in the deployment's high success rate and customer satisfaction scores.

8. Governance Integration

8.1 GRAF in the Factory Context

The GRAF framework is the governance backbone of the Agentification Factory. Every deployment produced by the factory is designed against the GRAF architecture, with the specific layers activated based on the zone classification and pattern selection.

For Zone I deployments (Single-Tool Agent, simple ReAct), GRAF Layers 1-3 are activated: Input Validation, Context Boundary Management, and Tool Access Control. These layers ensure that the agent operates within its defined scope and does not access unauthorised systems or data.

For Zone II deployments (Plan-and-Execute, Orchestrator-Subagent, Memory-Augmented), GRAF Layers 1-5 are activated, adding Output Validation and Compliance Gateway. These additional layers ensure that the agent's outputs meet quality standards and comply with applicable policies before being delivered to end users.

For Zone III deployments (Critic-Actor, Neuro-Symbolic/OCG), all seven GRAF layers are activated, including Audit and Explainability and Human Oversight Orchestration. These layers ensure that every action taken by the agent is logged, explainable, and subject to human review at defined trigger thresholds.

8.2 HITL Design Principles

Human-in-the-Loop (HITL) design is one of the most critical and most frequently misunderstood aspects of agentic AI governance. The factory model applies four HITL design principles derived from empirical analysis of successful deployments.

The first principle is zone-appropriate trigger design. Zone I deployments use sampling-based HITL (reviewing a random sample of 5-10% of outputs), which provides quality assurance without creating a bottleneck. Zone II deployments use exception-based HITL (routing outputs that fall outside defined confidence thresholds to human review). Zone III deployments use mandatory pre-execution HITL for all high-risk actions.

The second principle is graceful escalation. HITL workflows should be designed to escalate smoothly from automated to human handling, without disrupting the user experience. This requires clear handoff protocols, context preservation, and human agent training.

The third principle is HITL as a learning mechanism. Human reviews should be systematically captured and used to improve the agent's performance over time. This requires a structured feedback loop between the HITL workflow and the agent's training or prompt engineering process.

The fourth principle is HITL cost accounting. The cost of human review must be included in the Roundtrip Value calculation. Deployments that appear to generate high ROI by eliminating human labour but require substantial HITL oversight are not generating the claimed value.

9. EU AI Act Implications

9.1 Risk Classification and the Factory Model

The EU AI Act, which entered into force in August 2024, establishes a risk-based regulatory framework for AI systems deployed in the European Union. The factory model's zone classification system maps naturally onto the EU AI Act's risk categories, providing a practical implementation path for compliance.

Zone IV processes (Do Not Automate) correspond to the EU AI Act's "prohibited AI practices" category: AI systems that pose unacceptable risks to fundamental rights or safety. The factory model's hard-stop criteria prevent the deployment of agents in these categories.

Zone III processes (Automate with Caution) correspond to the EU AI Act's "high-risk AI systems" category: AI systems used in regulated sectors (healthcare, financial services, employment) or for high-stakes decisions. The factory model's mandatory OCG gateway and full GRAF governance provide the technical and organisational measures required for compliance with the high-risk AI requirements.

Zone I-II processes (Automate Now / Pilot First) correspond to the EU AI Act's "limited risk" and "minimal risk" categories, subject to transparency obligations and basic governance requirements. The factory model's GRAF Layers 1-4 provide the required governance infrastructure.

9.2 Documentation and Audit Requirements

The EU AI Act requires high-risk AI systems to maintain comprehensive technical documentation, including descriptions of the system's design, training data, testing procedures, and risk management measures. The factory model's standardised documentation artefacts — PASF scorecards, Pattern Specification Documents, Governance Architecture Documents, and Roundtrip Value reports — provide a natural foundation for EU AI Act compliance documentation.

The factory's Observability & Assurance infrastructure, including real-time telemetry, anomaly detection, and audit trail capabilities, also provides the monitoring and logging capabilities required by the EU AI Act's post-market monitoring requirements.

10. Strategic Implications for Leadership

10.1 The Four Strategic Tensions

Enterprise leaders navigating the adoption of the Agentification Factory model must manage four fundamental strategic tensions identified in our research.

Scalability vs. Adaptability. The factory model's standardisation creates efficiency and consistency, but may reduce the organisation's ability to respond rapidly to novel use cases or emerging AI capabilities. Leaders must design the factory with sufficient flexibility to incorporate new patterns and frameworks as the technology evolves, without sacrificing the governance rigour that makes the factory model effective.

Experience vs. Expediency. The factory model's phased deployment approach prioritises thoroughness over speed. Business units under pressure to demonstrate AI value quickly may resist the factory's assessment and governance requirements. Leaders must build the business case for the factory model's slower-but-more-reliable approach, using the empirical evidence of higher success rates and lower failure costs.

Supervision vs. Autonomy. As agents become more capable, the appropriate level of human supervision decreases. Leaders must design HITL workflows that evolve with the technology, progressively reducing human oversight as agent reliability is demonstrated, while maintaining the ability to increase oversight rapidly if problems emerge.

Retrofit vs. Reengineer. The factory model can be applied to existing processes (retrofit) or used to design new processes from scratch (reengineer). Reengineering typically generates higher ROI but requires greater organisational change. Leaders must make deliberate choices about when to retrofit and when to reengineer, based on the strategic importance of the process and the organisation's change capacity.

10.2 The Factory as Competitive Differentiator

The Agentification Factory model is not merely an operational efficiency tool; it is a source of sustainable competitive advantage. Organisations that establish a factory capability early gain three durable advantages.

First, they accumulate a proprietary pattern library that encodes their specific business domain knowledge, governance requirements, and operational constraints. This library is not replicable by competitors without equivalent deployment experience.

Second, they develop institutional expertise in agentic AI governance that is increasingly valuable as regulatory requirements intensify. The EU AI Act, the NIST AI Risk Management Framework, and emerging sector-specific regulations all require the kind of systematic governance that the factory model provides.

Third, they establish the data infrastructure and quality standards required for advanced agentic AI applications. The factory model's Data Foundation capability, which enforces data quality standards and provenance tracking, creates a data asset that enables increasingly sophisticated automation over time.

11. Conclusion

The Agentification Factory model represents the organisational response to the scaling challenge in enterprise agentic AI. By replacing ad-hoc, project-based deployment with a systematic, governed production pipeline, the factory model transforms agent deployment from a bespoke engineering challenge into a repeatable business process.

The empirical evidence presented in this paper is unambiguous: organisations that adopt the factory model achieve higher success rates (65% vs. 31%), faster time-to-value (40-60% reduction), lower marginal costs (42% reduction), and more accurate ROI recognition (71% vs. 38% accuracy). These advantages compound over time through the ROI Flywheel effect, creating a sustainable competitive moat.

The factory model is not a technology; it is an organisational capability. Building it requires investment in people (the right roles and skills), processes (the five delivery phases and their exit criteria), and platforms (the Enterprise Intelligence Platform and its supporting infrastructure). This investment is substantial, but the empirical evidence demonstrates that it generates returns that far exceed the alternative: a continued cycle of expensive, high-failure-rate ad-hoc deployments.

For organisations beginning their factory journey, the recommended starting point is the Minimum Viable Factory: 8-12 FTEs, focused on achieving Level 3 maturity within 12-18 months. For organisations already operating at Level 2-3, the priority is the transition to Level 4: implementing the OCG gateway, optimising tokenomics, and establishing predictive zone scoring. For organisations at Level 4-5, the focus shifts to the autonomous factory: using AI to optimise the factory's own processes and deepen the competitive moat.

The era of experimental, ad-hoc agentic AI deployment is ending. The organisations that will lead the AI-driven enterprise of the next decade are those that build the factory today.

References

Gartner. (2025). *Predicts 2026: Agentic AI governance and risk*. Gartner Research.

van Hurne, M. (2025a). *Roundtrip value governance for agentic process automation*. EIGENVECTOR RESEARCH.

van Hurne, M. (2025b). *Process automation suitability framework (PASF) and process automation design engine (PADE): A unified framework for enterprise agentic AI deployment*. EIGENVECTOR RESEARCH.

van Hurne, M. (2025c). *GRAF: Governed runtime for agentic functions — Architecture specification*. EIGENVECTOR RESEARCH.

van Hurne, M. (2025d). *OCG: Ontological compliance gateway for regulated agentic AI*. EIGENVECTOR RESEARCH.

van Hurne, M. (2026). *Agentic success patterns: A unified framework for enterprise AI deployment — From process assessment to governed automation at scale*. EIGENVECTOR RESEARCH.

Deloitte. (2026). *Agentic AI in healthcare: Operating model transformation*. Deloitte Insights.

McKinsey & Company. (2025). *Reimagining tech infrastructure for and with agentic AI*. McKinsey Technology.

MIT Sloan Management Review. (2025). *The emerging agentic enterprise: How leaders must navigate a new age of AI*. MIT Sloan Management Review.

OpenAI. (2025). *BNY: Building the future of financial services with AI*. OpenAI Case Studies.

PwC. (2025). *PwC and Microsoft Copilot: Enterprise AI deployment at scale*. PwC Case Studies.

Salesforce. (2025). *Agentforce ROI case studies: Verified outcomes from enterprise deployments*. Salesforce Research.

NIST. (2023). *AI risk management framework (AI RMF 1.0)*. National Institute of Standards and Technology.

European Parliament. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council — Artificial Intelligence Act*. Official Journal of the European Union.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). *ReAct: Synergizing reasoning and acting in language models*. arXiv:2210.03629.

Shinn, N., Cassano, F., Labash, B., Gopinath, A., Narasimhan, K., & Yao, S. (2023). *Reflexion: Language agents with verbal reinforcement learning*. arXiv:2303.11366.

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., & Chen, Z. (2024). *A survey on large language model based autonomous agents*. Frontiers of Computer Science, 18(6).

Goldrich, J. (2025). *Four strategic tensions in the agentic enterprise*. Medium / TechTonic Shifts.

HatchWorks. (2025). *Orchestrating AI agents: Production patterns for enterprise deployments*. HatchWorks Engineering Blog.

Andersen Institute. (2025). *Agentic AI success patterns: Empirical analysis of 50 enterprise deployments*. Andersen Institute Research.

STRIDE Research Consortium. (2024). *Task suitability scoring for agentic AI: The STRIDE framework*. arXiv:2512.02228.