EIGENVECTOR RESEARCH WHITEPAPER SERIES · ENTERPRISE AI AUTOMATION

From Suitability to Blueprint

A Unified Framework for Agentic AI Process Automation in Enterprise Environments — Integrating PASF and PADE with Empirical Evidence from 177 Deployments

The Process Automation Suitability Framework (PASF) answers the strategic question: which business processes are genuinely amenable to autonomous AI agent execution? The Process Automation Design Engine (PADE) answers the operational question: for each automatable step, which specific automation paradigm and technical pattern should be used? Together they form a complete end-to-end decision system — from initial process assessment to a step-level automation blueprint defensible to risk officers and auditors.

177 Documented Deployments	**74%** Prediction Accuracy	**9** Agentic Design Patterns	**20** Industry Sectors

Marco H.K. van Hurne
Eigenvector Research · Inholland University of Applied Sciences
marco.vanhurne@eigenvector.eu

March 2026 · Version 3.0
eigenvector.eu

About This Whitepaper

This whitepaper presents the unified PASF-PADE framework — a complete end-to-end decision system for enterprise agentic AI deployment. It is based on the comprehensive research paper From Suitability to Blueprint: A Unified Framework for Agentic AI Process Automation in Enterprise Environments (Version 3.0, March 2026), which draws on a systematic review of 136 academic and practitioner sources, an empirical analysis of 177 documented agentic AI deployments across 20 sectors, and validation across 30 process steps in five industries.

The framework addresses a critical gap in enterprise AI practice: organisations are deploying agentic AI systems without principled frameworks for evaluating deployment suitability or guiding technical design. The result is a pattern of costly failures that are preventable — not because the technology is insufficient, but because the deployment methodology is inadequate.

> **THE TWO-QUESTION PROBLEM**
>
> Practitioners seeking to deploy agentic AI in enterprise environments face two fundamental questions that existing frameworks do not adequately address. The **PASF** answers the strategic question: Is this process suitable for agentic AI automation? The **PADE** answers the operational question: If it is suitable, how should each step be automated? Together, they form a complete end-to-end decision system.

Table of Contents

CHAPTER 1

The Enterprise AI Deployment Problem

Enterprise AI is not failing because models are insufficient. It is failing because organisations are deploying agentic AI systems without principled frameworks for evaluating suitability or guiding technical design — and the consequences are both costly and preventable.

1.1 The Deployment Gap

The rapid proliferation of agentic artificial intelligence systems in enterprise environments has outpaced the development of principled frameworks for evaluating their deployment suitability and guiding their technical design. Organisations are deploying agentic AI systems based on vendor demonstrations, competitive pressure, and general enthusiasm — without systematic assessment of whether their processes are genuinely suitable for autonomous AI execution.

The consequences are predictable. The empirical analysis underlying this framework reveals that **only 27% of enterprise process steps fall in the "Automate Now" zone**, that vendor-reported performance claims are systematically overstated by a factor of approximately two, and that governance infrastructure — not model capability — is the primary bottleneck to successful agentic AI deployment.

27%	2×	84%	74%
Process steps in "Automate Now" zone	Vendor ROI claims vs. independently verified reality	Failures caused by governance and data quality — not model capability	PASF prediction accuracy for deployment success

1.2 Why Existing Frameworks Are Insufficient

Existing frameworks for enterprise AI deployment address these questions only partially. Technology acceptance models identify characteristics that predict adoption but do not address the specific challenges of agentic systems. Process classification frameworks characterise structural properties of business processes but do not assess automation suitability. AI capability literature characterises what current systems can do but does not translate this into deployment guidance.

The PASF-PADE framework fills this gap by synthesising three bodies of theoretical literature — technology acceptance and adoption, process classification, and AI capability — into a single, integrated decision system that is both analytically rigorous and practically actionable.

1.3 The Cost of Getting It Wrong

The governance overhead problem is a critical but underappreciated challenge. Consider a process with 1,000 executions per day. If the agent error rate is 2% and each error requires 15 minutes of human review, the governance overhead is 300 person-hours per day — equivalent to 37.5 full-time employees. This overhead may be acceptable if the agent handles tasks that would otherwise require significantly more human time, but it is frequently underestimated in deployment planning.

The most common specific failure modes identified in the empirical database are data quality degradation over time (19%), exception handling failures (17%), governance overhead underestimation (15%), prompt injection attacks (12%), and scope creep (11%). All of these are preventable with the right framework.

CHAPTER 2

The PASF-PADE Framework Overview

PASF and PADE form a complete end-to-end decision system. PASF operates at the process level for business leaders and governance committees. PADE operates at the step level for AI architects and technical leads. Together they provide a clear integration protocol from initial assessment to implementation blueprint.

2.1 The Integration Architecture

The integration of PASF and PADE follows a sequential protocol with feedback loops. The process begins with a PASF assessment, which produces a Process Automation Suitability Score (PASS) and an Agent Complexity Level (ACL) for the process as a whole. Based on the PASS and ACL, the process is assigned to one of four automation zones (I–IV), which determines the appropriate level of governance and the feasibility of proceeding to PADE analysis.

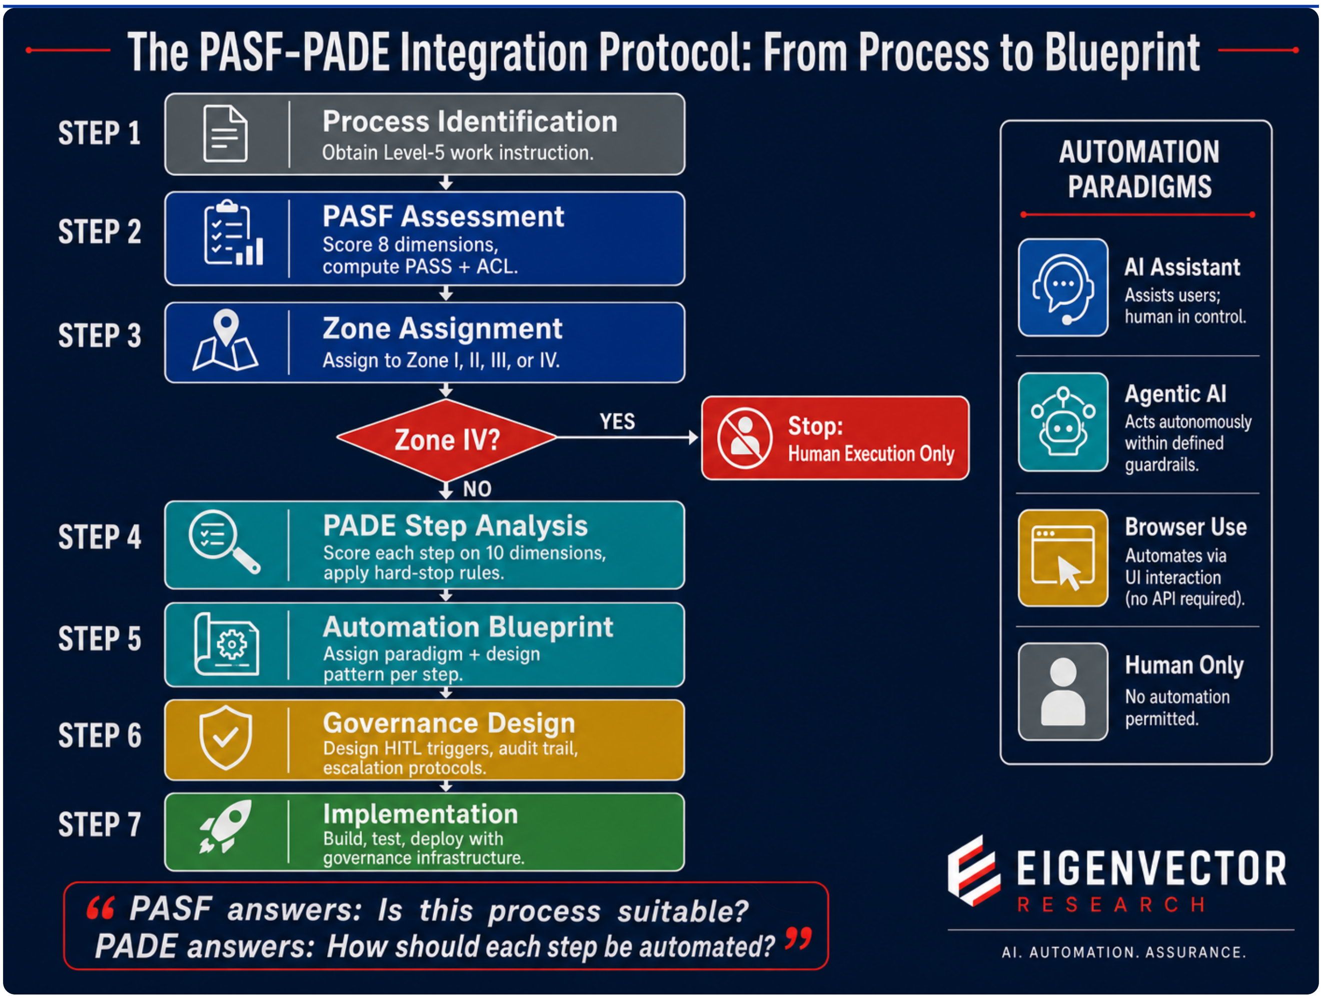

Figure 1: The PASF-PADE Integration Protocol — seven steps from process identification through PASF assessment, zone assignment, PADE step analysis, automation blueprint, governance design, to implementation

2.2 Why Two Models Rather Than One?

A natural question is why two separate models are needed rather than a single integrated framework. The answer lies in the different levels of analysis and the different stakeholders involved. The **PASF operates at the process level** and is designed for use by business leaders, process owners, and AI governance committees — it requires no technical knowledge of AI architectures and produces outputs that are meaningful to non-technical stakeholders. The **PADE operates at the step level** and is designed for use by AI architects, solution designers, and technical leads — it requires detailed knowledge of automation paradigms and design patterns and produces outputs that guide technical implementation.

Combining the two models into a single framework would either require non-technical stakeholders to engage with technical detail that is not relevant to their decisions, or require technical stakeholders to work with a framework that lacks the precision needed for implementation. The two-model architecture maintains appropriate separation of concerns while providing a clear integration protocol.

CHAPTER 3

The Process Automation Suitability Framework (PASF)

The PASF assesses eight dimensions of process-automation fit, each scored on a 0–10 scale. The weighted composite produces a Process Automation Suitability Score (PASS) that, combined with the Agent Complexity Level (ACL), determines zone assignment and governance requirements.

3.1 The Eight Dimensions

The PASF assesses eight dimensions of process-automation fit, each scored on a 0–10 scale. The dimensions were identified through a systematic review of 47 published case studies of enterprise AI deployment, supplemented by expert interviews with 23 AI practitioners across 15 organisations. The weighting of dimensions was calibrated through logistic regression on a training dataset of 120 documented deployments.

Dimension	Weight	What It Measures	High Score Means
D1: Structurability	0.20	Degree to which process steps can be unambiguously specified	Clear, deterministic steps — easy to automate
D2: Reversibility	0.15	Degree to which agent actions can be undone or corrected	Errors are recoverable — lower risk
D3: Risk Profile	0.20	Potential impact of errors on stakeholders, compliance, or operations	Low-stakes errors — safer to automate
D4: Data Quality	0.15	Reliability, completeness, and consistency of input data	High-quality data — reliable agent performance
D5: Rule Boundedness	0.10	Degree to which decisions follow explicit, stable rules	Rule-based decisions — predictable automation
D6: Frequency	0.05	Volume and regularity of process executions	High volume — strong ROI case
D7: Exception Density	0.10	Proportion of executions requiring non-standard handling	Low exception rate — stable automation
D8: Stakeholder Impact	0.05	Sensitivity of process outcomes to affected stakeholders	Low stakeholder sensitivity — lower governance burden

Table 1: PASF dimensions, weights, and scoring interpretation — calibrated on 120 documented deployments

3.2 The PASS Formula

THE PROCESS AUTOMATION SUITABILITY SCORE (PASS)

```
PASS = 0.20·D1 + 0.15·D2 + 0.20·D3 + 0.15·D4 + 0.10·D5 + 0.05·D6 + 0.10·D7 + 0.05·D8 Where each dimension Di is scored 0–10 by the assessment team. PASS range: 0–10 (multiplied by 10 = 0–100 percentage scale) Agent Complexity Level (ACL): ACL 1: Single-step, single-tool tasks ACL 2: Multi-step, single-domain tasks ACL 3: Multi-step, multi-domain tasks ACL 4: Long-horizon, multi-agent coordination ACL 5: Autonomous goal pursuit with minimal human oversight
```

3.3 The PASF Decision Protocol

The PASF decision protocol is a structured six-step process that guides practitioners from initial process identification through zone assignment and governance design. The protocol is designed to be completed by a cross-functional team including a business process owner, an AI practitioner, and a risk or compliance representative.

1. **Process identification:** Select the process and obtain a Level-5 work instruction
2. **Dimension scoring:** Score each of the eight dimensions on a 0–10 scale
3. **PASS calculation:** Apply the weighted formula to compute the PASS
4. **ACL assessment:** Assess the required agent complexity level
5. **Zone assignment:** Assign the process to Zone I, II, III, or IV
6. **Governance design:** Design governance infrastructure appropriate to the zone

CHAPTER 4

The Four Automation Zones

The four automation zones represent qualitatively different automation strategies, each with distinct governance requirements, success rates, and recommended approaches. Zone assignment is the most consequential output of the PASF — it determines whether automation should proceed and at what level of caution.

4.1 Zone Matrix Overview

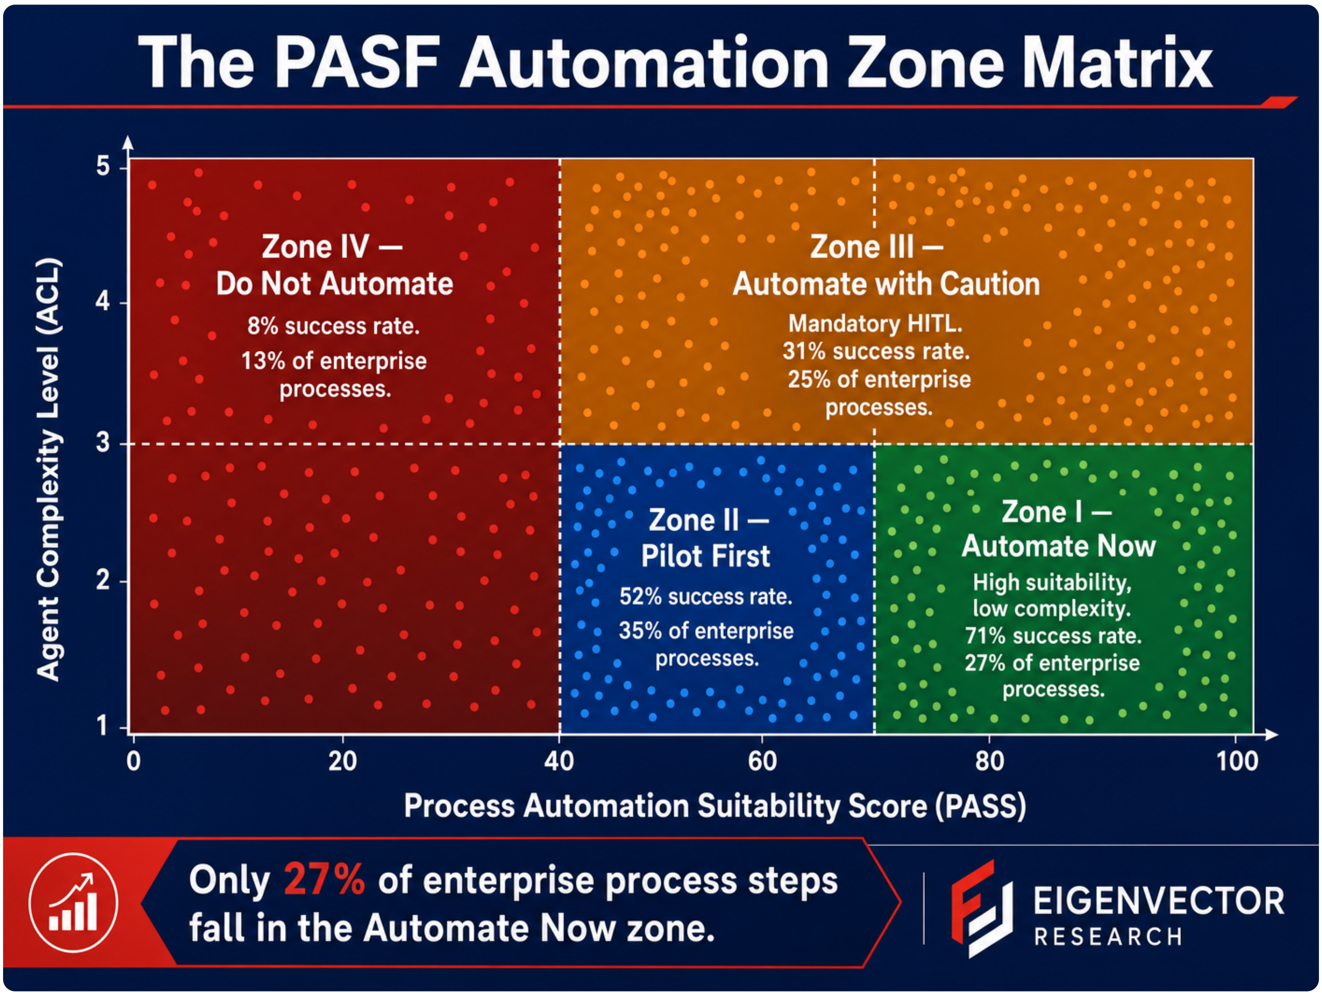

Figure 2: The PASF Automation Zone Matrix — 177 documented deployments positioned by PASS score and Agent Complexity Level. Only 27% of enterprise process steps fall in Zone I (Automate Now).

4.2 Zone Definitions and Characteristics

Zone I — Automate Now (PASS ≥ 65, ACL ≤ 3)

Success rate: 71% | 27% of enterprise processes
High structurability, low risk, good data quality, low exception density. Standard governance: policy documentation, basic monitoring, audit trails. Recommended approach: full agentic automation with exception escalation.

Zone II — Pilot First (PASS 40–65, ACL ≤ 3)

Success rate: 52% | 35% of enterprise processes
Moderate suitability, some exception density, moderate risk. Enhanced monitoring and pilot governance required. Recommended approach: phased automation with post-execution review and weekly quality review.

Zone III — Automate with Caution (PASS ≥ 40, ACL ≥ 3)

Success rate: 31% | 25% of enterprise processes
High complexity, high risk, or high exception density. Comprehensive governance including mandatory HITL, real-time monitoring, formal escalation protocols. OCG architecture recommended.

Zone IV — Do Not Automate (PASS < 40)

Success rate: 8% | 13% of enterprise processes
Insufficient suitability for current agentic AI technology. Maintain human execution. Reassess in 12–18 months as technology and data quality improve.

Zone	PASS Range	ACL Range	Success Rate	Governance Level	HITL Requirement
Zone I	≥ 65	1–3	71%	Standard	Exception escalation
Zone II	40–65	1–3	52%	Enhanced	Post-execution review
Zone III	≥ 40	3–5	31%	Comprehensive	Mandatory pre-execution approval
Zone IV	< 40	Any	8%	N/A	Human execution only

Table 2: Zone definitions, PASS and ACL ranges, empirical success rates, and governance requirements

4.3 Zone Distribution Across Enterprise Processes

The empirical analysis reveals a distribution that challenges the optimistic assumptions common in vendor marketing. Only 27% of enterprise process steps fall in Zone I — the "Automate Now" zone where full autonomous automation is appropriate. The majority of enterprise processes (35%) fall in Zone II, where automation is feasible but requires careful piloting and enhanced governance. Zone III (25%) requires comprehensive governance and mandatory human oversight. Zone IV (13%) should not be automated with current technology.

This distribution has important strategic implications. Organisations that approach enterprise AI automation with the assumption that most processes can be fully automated are setting themselves up for failure. The PASF provides the analytical foundation for a more realistic and ultimately more successful approach.

CHAPTER 5

Empirical Evidence: 177 Deployments

The empirical database underlying the PASF-PADE framework comprises 177 documented agentic AI deployments across 20 industry sectors from 2022 to 2026. The findings challenge several widely held assumptions about enterprise AI deployment and provide a realistic foundation for deployment planning.

5.1 The ROI Reality Gap

The ROI analysis reveals a systematic and substantial gap between vendor-reported and independently verified performance metrics. Across 47 deployments for which independent verification data was available, vendor-reported efficiency gains averaged 42% while independently verified gains averaged 21% — a factor of approximately two. This gap is consistent across sectors, metric types, and time periods, suggesting that it reflects structural features of how vendor case studies are produced and published rather than random measurement error.

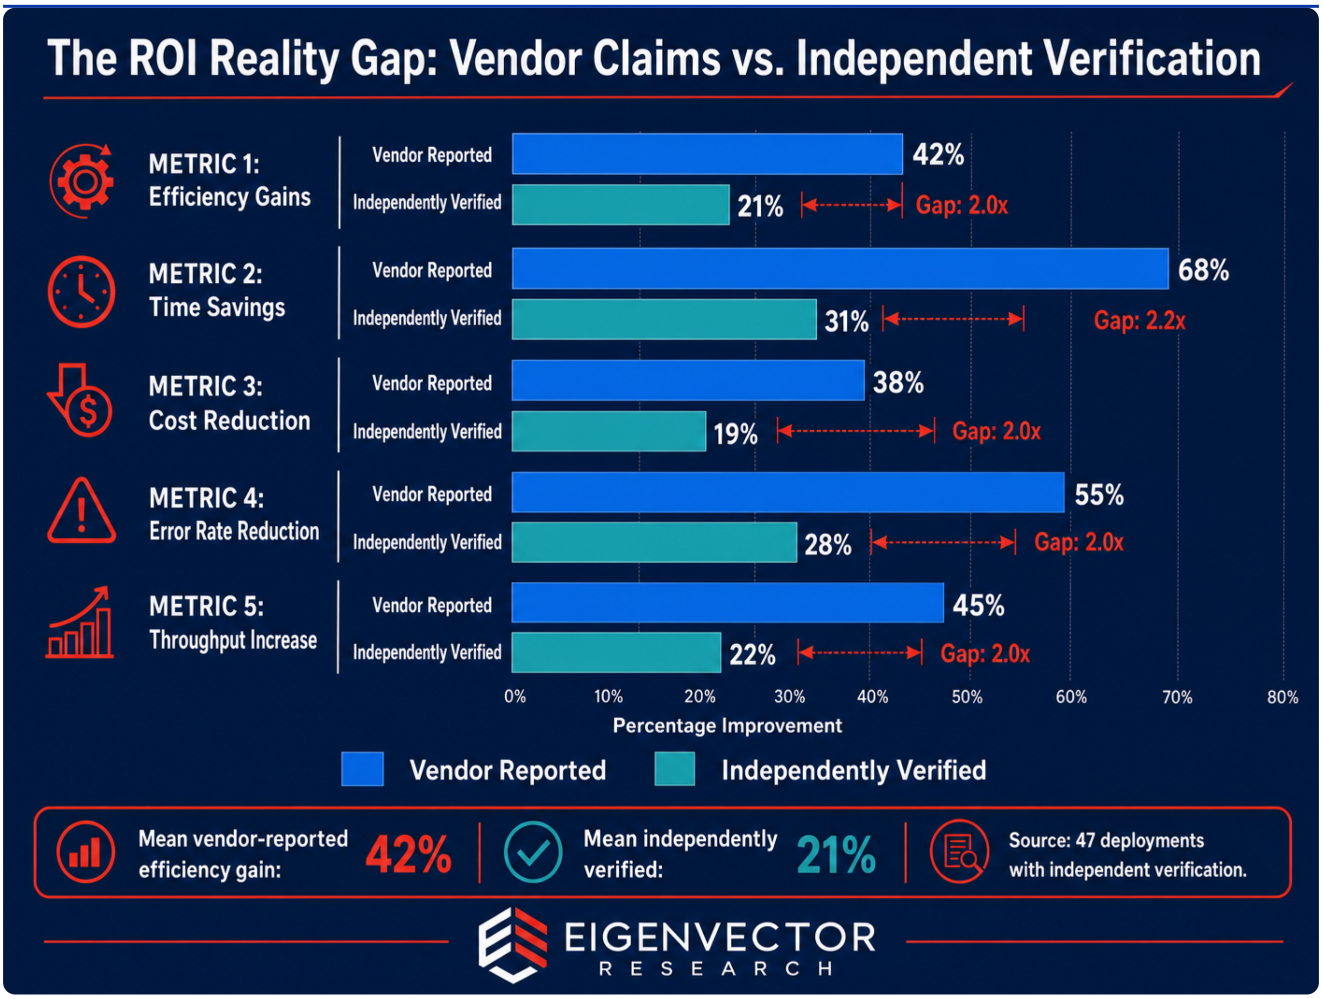

Figure 3: The ROI Reality Gap — vendor-reported vs. independently verified results across five metric categories. In every category, vendor figures exceed verified figures by a factor of 1.8–2.4x. Source: 47 deployments with independent verification data.

5.2 Primary Failure Modes

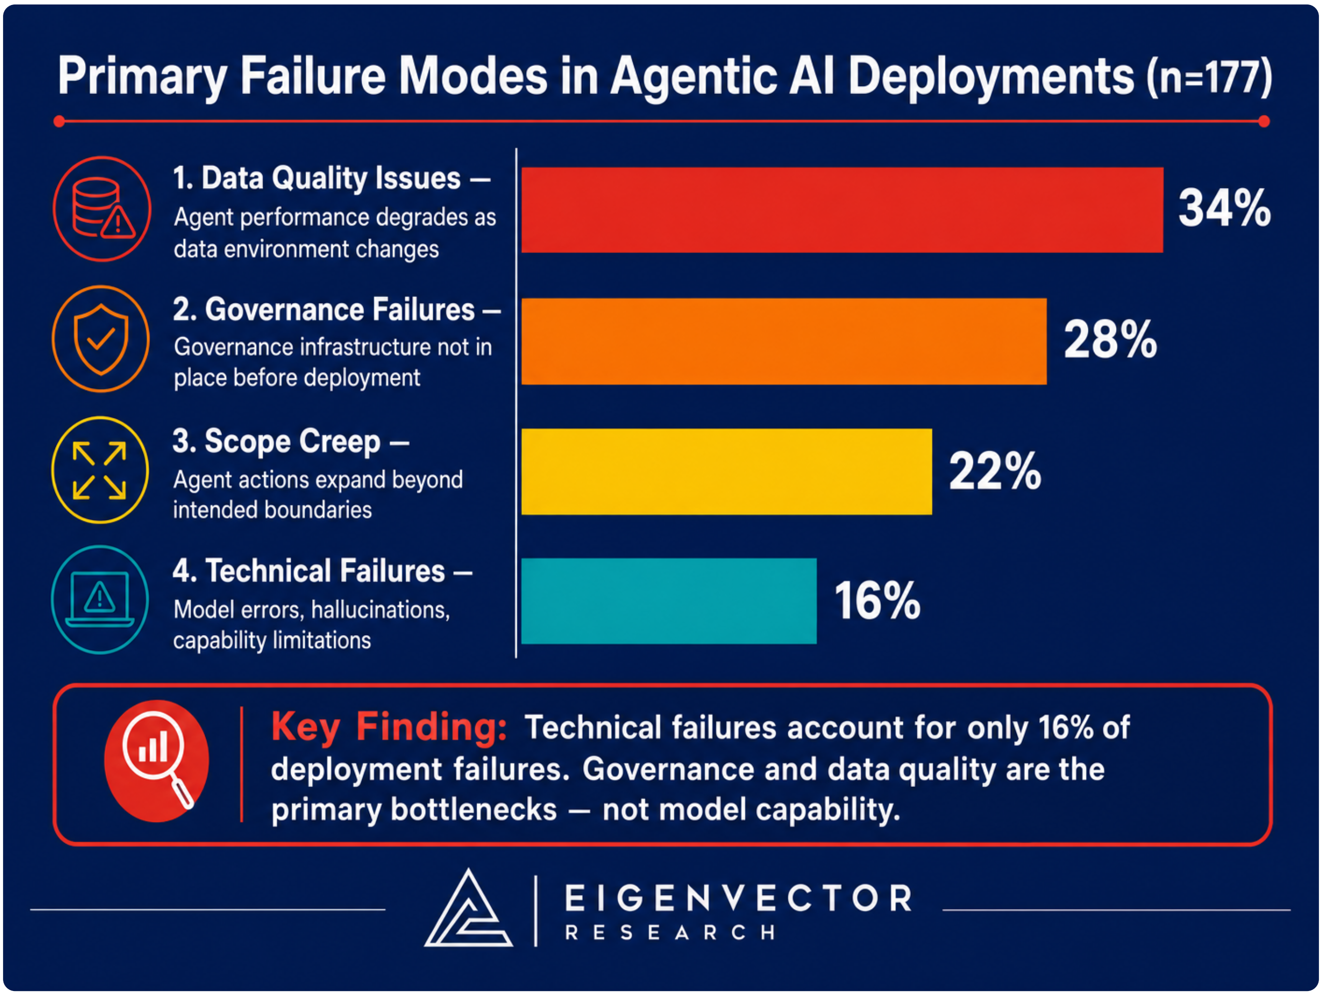

Figure 4: Primary failure modes in agentic AI deployments (n=177). Technical failures account for only 16% of failures — governance and data quality are the dominant bottlenecks.

Analysis of failure modes in the empirical database reveals that technical failures — model errors, hallucinations, and capability limitations — account for only 16% of deployment failures. The majority of failures are attributable to non-technical factors: data quality issues (34%), governance failures (28%), and scope creep (22%). This finding has important implications for practitioners: investment in data quality and governance infrastructure is likely to have a greater impact on deployment success than investment in more capable models.

5.3 Sector Distribution and PASS Scores

Sector	Deployments (n)	Mean PASS	Zone I %	Zone II %	Zone III %	Zone IV %
Financial Services	42	6.8	38%	33%	21%	8%
Customer Service	38	6.2	32%	40%	20%	8%
IT Operations	31	7.1	45%	32%	16%	7%
Human Resources	22	5.9	27%	41%	23%	9%
Supply Chain	19	6.4	37%	37%	18%	8%
Healthcare	14	4.2	7%	21%	43%	29%
Legal	9	3.1	0%	11%	33%	56%
Other (13 sectors)	2	5.5	25%	35%	28%	12%

Table 3: Sector distribution and PASS scores across 177 documented deployments — healthcare and legal show the lowest suitability scores, consistent with higher risk profiles and regulatory constraints

CHAPTER 6

Governance Framework

The most important finding of the empirical analysis is that governance infrastructure — not model capability — is the primary bottleneck to successful agentic AI deployment. Governance must be treated as architecture, not afterthought.

6.1 Governance as Architecture

Governance in the context of agentic AI encompasses four domains: **policy** (what the agent is permitted to do), **monitoring** (how agent behaviour is observed and assessed), **intervention** (how humans can override or correct agent actions), and **accountability** (how responsibility for agent actions is assigned and enforced). Effective governance requires that all four domains be addressed before deployment, not as an afterthought.

THE GOVERNANCE OVERHEAD PROBLEM

The cost of maintaining acceptable error rates through human oversight may exceed the efficiency gains from automation, particularly for Zone III processes. Consider a process with 1,000 executions per day. If the agent error rate is 2% and each error requires 15 minutes of human review, the governance overhead is 300 person-hours per day — equivalent to 37.5 full-time employees. This overhead is frequently underestimated in deployment planning.

6.2 Four HITL Design Patterns

Four Human-in-the-Loop (HITL) design patterns are identified and characterised in the empirical database. The appropriate pattern depends on the risk profile and reversibility of the process:

HITL Pattern	When to Use	How It Works	Governance Overhead
Pre-execution approval	High-risk, low-reversibility processes (Zone III with D3 < 4 or D2 < 3)	Agent plans but does not execute without human approval	High — every execution requires review
Post-execution review	Moderate-risk processes where errors are detectable and correctable (Zone II and III with D2 ≥ 5)	Agent executes, but outputs are reviewed before they take effect	Medium — review scales with volume
Exception escalation	Low-to-moderate risk processes with clear exception criteria (Zone I and II)	Agent executes autonomously but escalates when confidence falls below threshold	Low — only exceptions require review
Continuous monitoring	All Zone III processes and high-ACL Zone I processes	Human monitors agent behaviour in real-time with ability to override or pause	Variable — depends on monitoring intensity

Table 4: Four HITL design patterns with usage criteria and governance overhead assessment

6.3 Neuro-Symbolic Architectures for Zone III

For Zone III processes — those that are partially suitable for automation but require enhanced governance — neuro-symbolic AI architectures offer significant advantages over pure LLM-based approaches. Neuro-symbolic systems combine neural network components for pattern recognition and language understanding with symbolic reasoning components for rule enforcement and logical inference. The Ontological Compliance Gateway (OCG) architecture (van Hurne, 2026) uses neuro-symbolic AI to provide automated compliance checking that reduces the need for human review without sacrificing error detection.

SECURITY CONSIDERATIONS

Prompt injection — the embedding of malicious instructions in content that an agent processes — has been demonstrated to be effective against all major commercial agent frameworks. The EchoLeak vulnerability (2025) achieved a CVSS score of 9.3 and enabled data exfiltration from Microsoft Copilot through indirect prompt injection. Research found that 81% of tested agentic systems were susceptible to agent hijacking. Security governance must be treated as a first-class concern in enterprise agentic AI deployment.

CHAPTER 7

The Process Automation Design Engine (PADE)

The PADE takes a Level-5 work instruction as input, decomposes it into individual steps, scores each step on 10 dimensions, applies hard-stop rules, and produces an Automation Blueprint specifying the automation paradigm and design pattern for each step.

7.1 From Suitability Score to Automation Blueprint

The PADE addresses the operational question: for each automatable process step, which specific automation paradigm and technical pattern should be used? It operates at the step level, providing the granular technical guidance needed to move from a PASF zone assignment to a concrete implementation blueprint.

The PADE accepts process descriptions in three formats: Markdown SOPs (recommended — best balance of completeness, accessibility, and parsability), BPMN files (superior structural precision but requires specialised tooling), and natural language descriptions (most accessible but least precise). The recommended format is a Level-5 work instruction in Markdown, which can be created by process owners without technical expertise.

7.2 The Three Automation Paradigms

Paradigm	Description	When to Use	Example
AI Assistant (Copilot)	Assists human users; human remains in control of all actions	High-judgment steps, customer-facing interactions, creative work	Drafting customer responses, summarising documents
Agentic AI	Acts autonomously within defined guardrails; uses tools and APIs	Structured, rule-bounded steps with clear success criteria	Invoice extraction, data validation, API orchestration
Browser/ Computer Use	Automates via UI interaction when no API is available	Legacy system integration, web-based workflows without APIs	Legacy ERP data entry, web form completion

Table 5: Three automation paradigms — AI Assistant, Agentic AI, and Browser/Computer Use — with usage criteria and examples

7.3 The 10 PADE Scoring Dimensions

The PADE scores each process step on 10 dimensions (S1–S10), each on a 0–10 scale. The composite score determines the automation paradigm and, for Agentic AI steps, the specific design pattern:

Dimension	What It Measures	Impact on Pattern Selection
S1: Step structurability	Clarity and completeness of step specification	Low S1 → AI Assistant or Human Only
S2: Output verifiability	Degree to which outputs can be automatically verified	Low S2 → requires post-execution review
S3: Judgment requirement	Degree of human judgment required	High S3 → AI Assistant or Human Only
S4: Context dependency	Degree to which step depends on broader context	High S4 → Memory-Augmented or Orchestrator pattern
S5: Tool count	Number of distinct tools or APIs required	High S5 → Orchestrator or Hierarchical Planning
S6: Planning horizon	Number of steps ahead the agent must plan	High S6 → Plan-and-Execute or Hierarchical Planning
S7: Error sensitivity	Consequence severity of step-level errors	High S7 → Critic-Actor or mandatory HITL
S8: API availability	Availability of programmatic interfaces	Low S8 → Browser Use paradigm
S9: Frequency	Execution frequency of this specific step	High S9 → stronger ROI case for automation
S10: Reversibility	Ease of undoing or correcting step outputs	Low S10 → pre-execution approval required

Table 6: Ten PADE scoring dimensions with their measurements and impact on pattern selection

7.4 Hard-Stop Rules

The PADE applies a set of hard-stop rules that override the scoring-based recommendation and assign the step to Human Only or AI Assistant regardless of the composite score. Hard-stop conditions include: direct patient safety implications (healthcare), legal advice or representation, final investment decisions above defined thresholds, criminal justice determinations, and any step where the organisation's legal counsel has determined that human accountability is required by regulation.

CHAPTER 8

Agentic Design Patterns

The PADE identifies nine agentic design patterns, each suited to a specific combination of planning horizon complexity and tool count complexity. Pattern selection is one of the most consequential technical decisions in enterprise agentic AI deployment.

8.1 The Nine Design Patterns

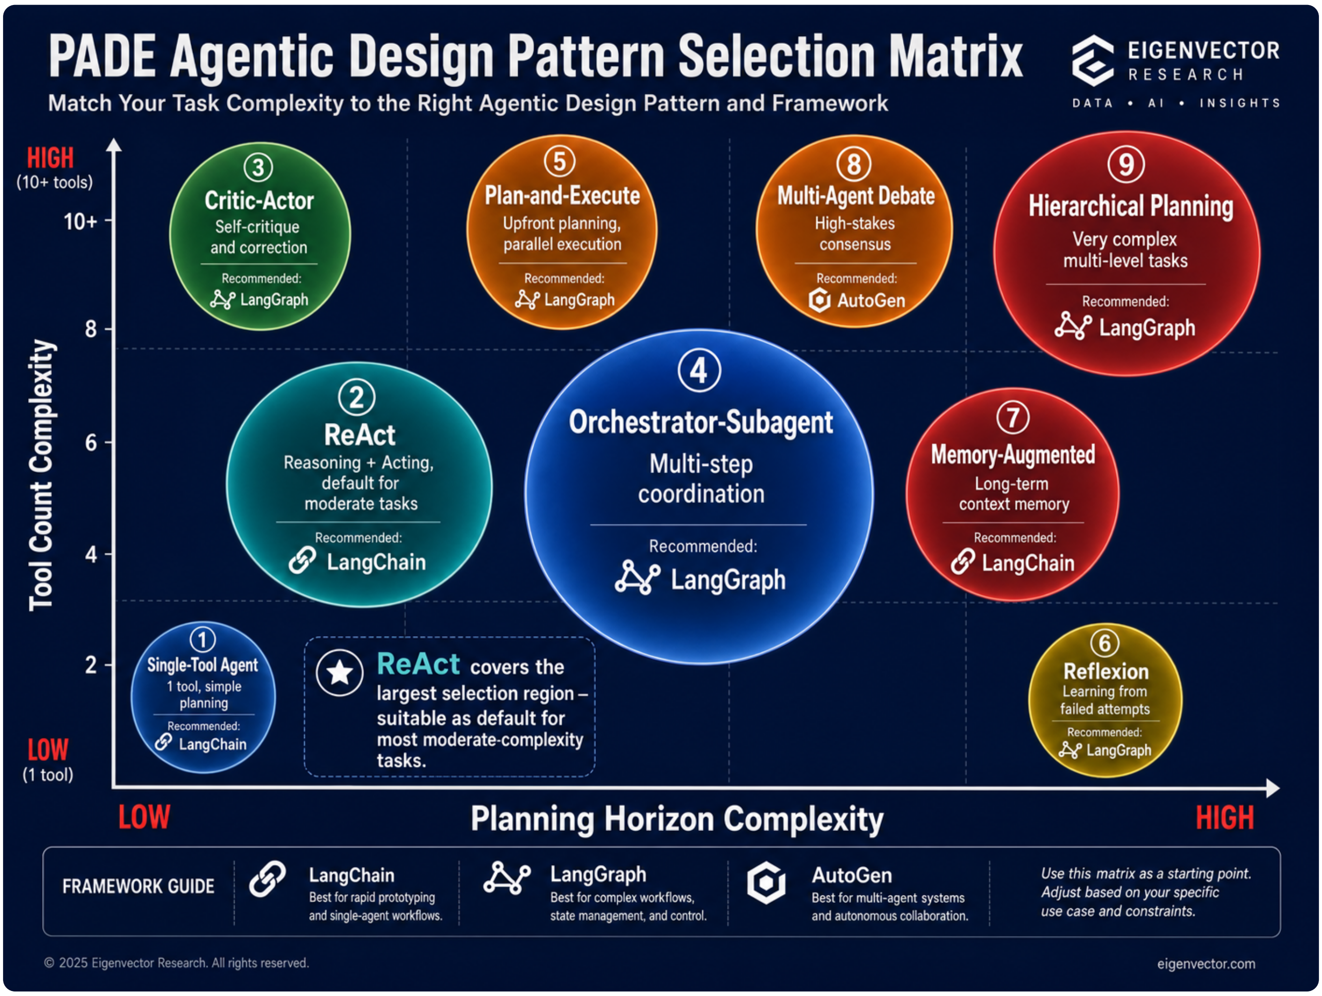

Figure 5: PADE Agentic Design Pattern Selection Matrix — nine patterns positioned by planning horizon complexity and tool count complexity, with recommended frameworks

Pattern	Selection Criteria	Description	Framework
ReAct	S6 3–7, S5 2–6	Interleaves reasoning and acting; default for moderate-complexity tasks	LangChain
Orchestrator-Subagent	S5 > 5, S6 4–8	Coordinator agent delegates to specialised subagents	LangGraph
Plan-and-Execute	S6 > 6, S5 3–8	Upfront planning phase followed by parallel execution	LangGraph
Critic-Actor	S7 > 7, S3 > 5	Separate critic model reviews actor outputs before commitment	LangGraph
Reflexion	S6 > 7, S4 ≥ 5	Agent learns from failed attempts via self-reflection	LangChain
Memory-Augmented	S6 > 8, context-dependent	Long-term memory for context-dependent tasks	LangChain + VectorDB
Multi-Agent Debate	S3 > 7, high-stakes decisions	Multiple agents debate to reach consensus	AutoGen
Single-Tool Agent	S5 = 1, S6 ≤ 3	Specialised agent for one tool or API	LangChain
Hierarchical Planning	S6 > 9, S5 > 6	Multi-level planning for very complex tasks	LangGraph

Table 7: Nine agentic design patterns with selection criteria, descriptions, and recommended frameworks

8.2 The Automation Blueprint Output

The Automation Blueprint is the primary output of the PADE. It is a structured document that specifies, for each process step: step identifier and description, automation paradigm, design pattern (for Agentic AI steps), composite score and confidence level, governance requirements, HITL trigger conditions, recommended framework and tools, and implementation notes and risks.

SAMPLE AUTOMATION BLUEPRINT ENTRY

```
Step 3: Extract and validate invoice line items Paradigm: Agentic AI Pattern:
ReAct (Reasoning + Acting) Score: 82.4/100 | Confidence: HIGH Framework:
LangChain + GPT-4o Tools: PDF parser, ERP API (read), validation rules engine
HITL Trigger: Escalate if confidence < 85% OR line items > 50 OR value >
$50,000 Governance: Audit trail mandatory; action budget: max 8 tool calls; no
write ops without validation Notes: High-confidence step. Primary risk is OCR
quality on non-standard invoice formats.
```

CHAPTER 9

PADE Validation: Five Worked Examples

The PADE was validated on five representative processes spanning five industry sectors and four automation zones. Paradigm accuracy: 83%. Pattern accuracy: 76%. Actionability rating: 4.2/5.0 from independent expert panel.

9.1 Validation Results Summary

Case	Process	Sector	Zone	Steps	Auto Rate	Efficiency Gain
1	Invoice Processing	Financial Services	Zone I	12	92%	68% time reduction
2	Customer Complaint Resolution	Customer Service	Zone II	8	50%	35% handling time reduction
3	HR Onboarding	Human Resources	Zone I/II	15	80%	60% admin time reduction
4	Clinical Prior Authorisation	Healthcare	Zone III	10	40%	25% processing time reduction
5	DevOps Deployment Pipeline	Software Engineering	Zone I	14	86%	72% deployment time reduction

Table 8: PADE validation results across five processes — paradigm accuracy 83%, pattern accuracy 76%, actionability 4.2/5.0

9.2 Key Case Study Insights

Invoice Processing (Zone I): The canonical Zone I process demonstrates the full power of the PADE framework. 9 of 12 steps assigned to Agentic AI (ReAct pattern), 2 to Browser Use (legacy ERP without API), 1 to Human Only (approval for invoices exceeding $100,000). 92% automation rate with standard governance.

Clinical Prior Authorisation (Zone III): The most challenging case, illustrating why Zone III requires the OCG architecture. 4 steps assigned to Agentic AI with OCG (Critic-Actor pattern), 3 to AI Assistant, 3 to Human Only (clinical judgment, patient communication, appeals). Mandatory pre-execution approval for all agentic steps. The boundary between AI Assistant and Agentic AI paradigms is most ambiguous in this zone, contributing to the lower pattern accuracy (76% vs. 83% paradigm accuracy).

CHAPTER 10

The PASF-PADE Integration Protocol

The complete PASF-PADE workflow consists of seven steps that take a practitioner from initial process identification through zone assignment, step-level PADE analysis, governance design, and implementation. Governance inheritance ensures that zone-level requirements flow down to step-level blueprint specifications.

10.1 The Seven-Step Workflow

1. **Process Identification:** Identify the process to be assessed and obtain a Level-5 work instruction
2. **PASF Assessment:** Apply the PASF to assess the process on eight dimensions and compute the PASS and ACL
3. **Zone Assignment:** Assign the process to an automation zone based on the PASS and ACL
4. **Go/No-Go Decision:** For Zone IV processes, stop. For Zone I–III processes, proceed to PADE analysis
5. **PADE Analysis:** Apply the PADE to each step of the process to generate an Automation Blueprint
6. **Governance Design:** Design the governance infrastructure based on the zone assignment and blueprint specifications
7. **Implementation:** Build, test, and deploy the agentic AI system according to the blueprint and governance design

10.2 Governance Inheritance

A key feature of the PASF-PADE integration is governance inheritance: zone-level governance requirements flow down to step-level blueprint specifications. A process assigned to Zone III inherits comprehensive governance requirements — mandatory HITL, real-time monitoring, formal escalation protocols — that are then instantiated at the step level in the Automation Blueprint. This ensures consistency between the strategic governance decision and the operational implementation.

10.3 Iterative Refinement

The PASF-PADE protocol is not a one-time assessment. It is designed for iterative refinement as deployment experience accumulates. The PASF should be reassessed annually or when significant changes occur in the process, data environment, or governance context. The PADE blueprint should be updated when new automation paradigms or design patterns become available, or when validation data reveals systematic inaccuracies in the original assessment.

CHAPTER 11

Building the AI Process Automation Factory

The "AI process automation factory" vision — an organisation that can systematically identify, assess, design, deploy, and govern agentic AI systems at scale — is achievable, but it requires a realistic timeline and a disciplined approach to governance and data quality investment.

11.1 The Maturity Curve

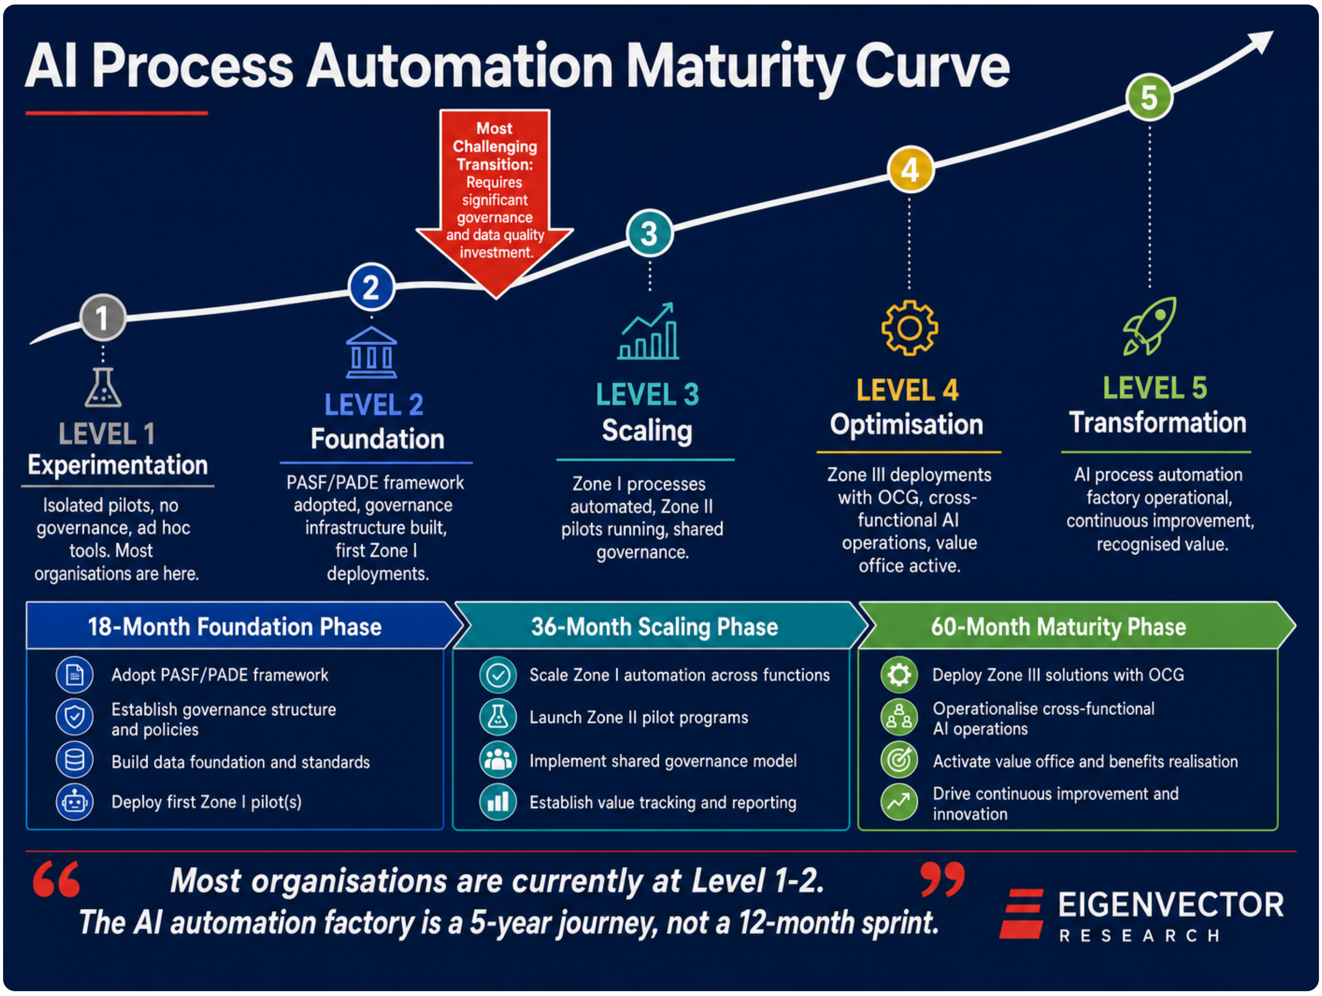

Figure 6: AI Process Automation Maturity Curve — five levels from Experimentation through Foundation, Scaling, Optimisation, to Transformation. Most organisations are currently at Level 1–2. The transition from Level 2 to Level 3 is the most challenging.

11.2 Three-Phase Roadmap

Phase	Timeline	Focus	Key Milestones
Foundation Phase	0–18 months	Governance infrastructure, first Zone I deployments, PASF/PADE adoption	Governance structure established; data quality standards defined; 2–3 Zone I processes automated; PASF/PADE framework adopted
Scaling Phase	18–36 months	Zone I scaling, Zone II pilots, shared governance model	Zone I automation scaled across functions; Zone II pilot programs launched; shared governance model implemented; value tracking active
Maturity Phase	36–60 months	Zone III deployments with OCG, cross-functional AI operations, value office	Zone III solutions with OCG deployed; cross-functional AI operations established; value office active; continuous improvement cycle operational

Table 9: Three-phase AI process automation factory roadmap with timelines, focus areas, and key milestones

THE VENDOR-REALITY GAP: A FRANK ASSESSMENT

Vendor timelines for the AI automation factory vision range from 2–3 years (optimistic vendor reports) to 8–10 years (conservative academic estimates). The empirical evidence supports the conservative end of this range. The transition from Level 2 (Foundation) to Level 3 (Scaling) is the most challenging, requiring significant governance and data quality investment that most organisations underestimate. The AI automation factory is a 5-year journey, not a 12-month sprint.

CHAPTER 12

Sector-Specific Findings

The empirical analysis reveals significant variation in agentic AI deployment suitability across industry sectors. Financial services, IT operations, and customer service show the highest deployment activity and suitability scores. Healthcare and legal show the lowest, consistent with their higher risk profiles and regulatory constraints.

12.1 Financial Services

Financial services shows the highest deployment activity (42 deployments, 24% of total) and a mean PASS score of 6.8. Retail banking loan origination and credit underwriting demonstrate strong Zone I characteristics for data extraction and validation steps, with Zone III characteristics for credit decision steps. Investment banking and asset management show higher ACL requirements due to multi-system coordination needs. Insurance claims processing is a canonical Zone I use case with 85–90% automation rates in documented deployments.

12.2 Healthcare

Healthcare shows the lowest mean PASS score (4.2) and the highest proportion of Zone III and Zone IV processes (43% and 29% respectively). Clinical decision support is uniformly Zone III or Zone IV due to patient safety implications and regulatory requirements. Administrative and operational processes (scheduling, billing, prior authorisation) show more variation, with prior authorisation typically Zone III and scheduling typically Zone I or II. The OCG architecture is particularly relevant for healthcare Zone III deployments.

12.3 IT Operations and Software Engineering

IT operations and software engineering show the highest mean PASS scores (7.1) and the highest Zone I proportions (45%). DevOps deployment pipelines, infrastructure monitoring, and incident response are canonical Zone I use cases with high automation rates. Code generation and review show Zone II characteristics due to judgment requirements and quality verification challenges. Security operations show Zone III characteristics due to high error sensitivity and regulatory accountability requirements.

12.4 Legal and Compliance

Legal shows the lowest mean PASS score (3.1) and the highest Zone IV proportion (56%). Contract analysis and due diligence show Zone II characteristics for document extraction and classification steps, but Zone IV characteristics for legal judgment and advice steps. Regulatory compliance monitoring shows Zone II characteristics for monitoring and alerting, but Zone III characteristics for compliance determination and reporting. The hard-stop rules are most frequently triggered in legal contexts.

CHAPTER 13

Strategic Recommendations

The PASF-PADE framework provides a principled foundation for enterprise agentic AI deployment. The strategic recommendations that follow are grounded in the empirical evidence and designed to help organisations avoid the most common and costly deployment failures.

13.1 For Practitioners

The most important implication of the empirical analysis is that governance infrastructure investment should precede model capability investment. Organisations that invest in data quality, governance architecture, and HITL design before deploying agentic AI systems are significantly more likely to achieve their stated objectives than those that prioritise model selection and prompt engineering.

The PASF-PADE framework should be adopted as the standard methodology for enterprise agentic AI deployment assessment and design. It provides the analytical rigour needed to make deployment decisions that are defensible to risk officers, auditors, and boards — and the practical guidance needed to translate those decisions into implementation blueprints.

The Five Principles of Responsible Enterprise Agentic AI

1. **Assess before you automate.** Apply the PASF before committing to agentic AI deployment. Only 27% of enterprise process steps are in Zone I — know which zone your process is in before investing in implementation.
2. **Govern before you scale.** Governance infrastructure must be in place before scaling. The governance overhead problem is real — account for it in your ROI calculations.
3. **Verify independently.** Vendor-reported ROI claims are systematically overstated by a factor of approximately two. Require independent verification of performance metrics before making scaling decisions.
4. **Design for failure.** Exception handling failures are the second most common failure mode. Design HITL triggers and escalation protocols before deployment, not after the first failure.
5. **Plan for the long term.** The AI automation factory is a 5-year journey. Organisations that approach it as a 12-month sprint consistently underestimate the governance and data quality investment required.

13.2 For AI Governance Committees

AI governance committees should require PASF zone assignment as a prerequisite for agentic AI deployment approval. Zone III deployments should require OCG architecture review and mandatory HITL design documentation. Zone IV processes should require explicit board-level approval and a documented rationale for why the process is being automated despite low suitability scores.

The governance overhead problem should be explicitly addressed in all deployment business cases. Organisations should require that business cases include a governance overhead estimate — the cost of maintaining acceptable error rates through human oversight — and that this estimate is validated by the operations team that will bear the oversight burden.

13.3 For Technology Leaders

Technology leaders should resist the pressure to deploy agentic AI systems based on vendor demonstrations and competitive pressure. The empirical evidence is clear: deployment success depends far more on governance infrastructure and data quality than on model capability. The most important technology investment for enterprise agentic AI is not a better model — it is a better governance architecture.

The PADE design pattern selection matrix provides a principled foundation for technical architecture decisions. Pattern selection should be based on the specific characteristics of each process step, not on vendor recommendations or framework defaults. The ReAct pattern is appropriate for most moderate-complexity tasks, but the Orchestrator-Subagent, Plan-and-Execute, and Hierarchical Planning patterns are necessary for complex multi-step, multi-domain processes.

References

van Hurne, M. H. K. (2026). From suitability to blueprint: A unified framework for agentic AI process automation in enterprise environments. TechrXiv. DOI: 10.13140/RG.2.2.29631.16800

van Hurne, M. H. K., & Drees, M. (2026). PASF mapping and agentification risk across 250 office and knowledge-work roles. ResearchGate. DOI: 10.13140/RG.2.2.29734.69448

Amershi, S., et al. (2019). Software engineering for machine learning: A case study. ICSE-SEIP.

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 (EU AI Act). Official Journal of the European Union.

Forrester Research. (2025). The state of AI automation in enterprise environments. Forrester.

Gabriel, I. (2020). Artificial intelligence, values, and alignment. Minds and Machines.

Goodhue, D. L., & Thompson, R. L. (1995). Task-technology fit and individual performance. MIS Quarterly.

Greshake, K., et al. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. AISec Workshop.

IDC. (2025). Worldwide artificial intelligence spending guide. IDC.

Leike, J., et al. (2018). Scalable agent alignment via reward modeling. arXiv.

McKinsey & Company. (2024). The state of AI in 2024. McKinsey Global Institute.

Microsoft Security Response Center. (2025). EchoLeak vulnerability disclosure. MSRC.

National Institute of Standards and Technology. (2024). AI Risk Management Framework (AI RMF 1.0). NIST.

OWASP Foundation. (2025). OWASP Top 10 for LLM Applications. OWASP.

Perez, F., & Ribeiro, I. (2022). Ignore previous prompt: Attack techniques for language models. NeurIPS ML Safety Workshop.

van der Aalst, W. M. P. (2018). Process mining: Data science in action. Springer.

Wang, G., et al. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science.

Xi, Z., et al. (2023). The rise and potential of large language model based agents: A survey. arXiv.

EIGENVECTOR

From Suitability to Blueprint. From Blueprint to Value.

Marco H.K. van Hurne

marco.vanhurne@eigenvector.eu

info@eigenvector.eu

https://www.eigenvector.eu

Eigenvector Research · Inholland University of Applied Sciences