The Modernization Engineering Company

SONATA

SONATA SOFTWARE

Whitepaper

From **AI experiments** to **enterprise impact**

Lessons from AWS generative AI assessments across industries

Table of Contents

1. Introduction ---- 3

2. Executive summary ---- 4

3. Customer context lens – who this applies to ---- 5

4. The structured AWS GenAI assessment framework ---- 6

5. Industries impacted by GenAI assessments ---- 7

6. Assessment snapshots ---- 8

7. GenAI Use Case Landscape Map ---- 9

8. Technologies evaluated across assessments ---- 10

9. Enterprise AI maturity model ---- 11

10. Unified agent architecture blueprint ---- 12

11. Key implementation patterns ---- 14

12. Outcomes observed ---- 15

13. The future: Agent-driven enterprises ---- 16

14. Conclusion ---- 17

15. About the author ---- 17

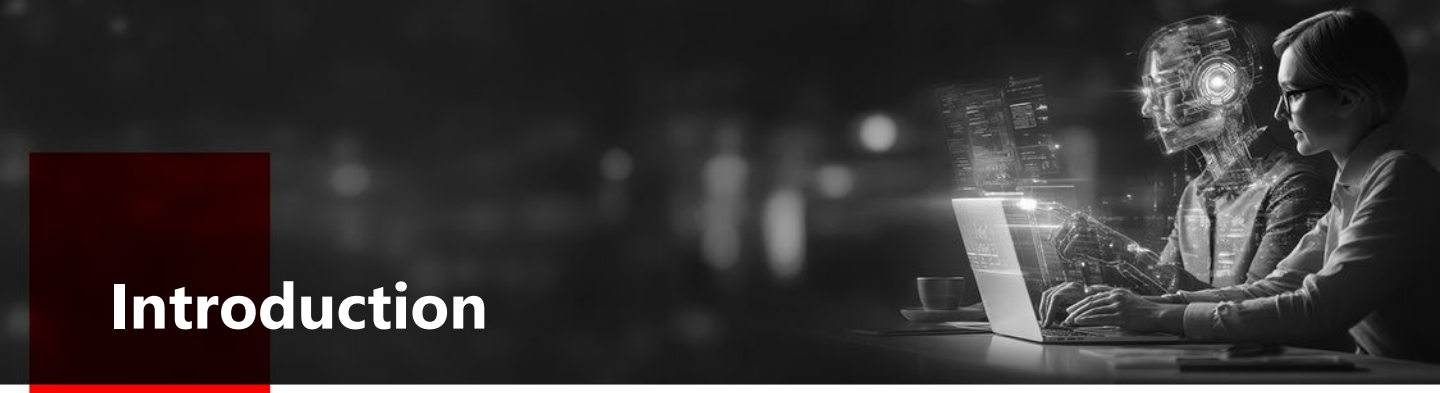

Introduction

Generative AI has crossed a critical threshold. What began as a wave of isolated experiments and proof-of-concepts is now reshaping how enterprises operate, compete, and deliver value. Across industries – from manufacturing and healthcare to retail, logistics, and financial services – organizations are no longer asking whether to adopt generative AI, but how fast they can move from experimentation to measurable impact.

Yet the path from pilot to production remains deceptively difficult. Many organizations find themselves stalled – not because the technology falls short, but because the organizational, architectural, and governance foundations needed to sustain AI at scale have not kept pace with ambition.

Over the past year, Sonata Software conducted AWS Generative AI assessments across organizations spanning six major industries. These engagements revealed a consistent and striking pattern: the enterprises making the most meaningful progress were not necessarily those with the most advanced models or the largest AI budgets. They were the ones that treated AI as a continuous product – with defined ownership, robust data foundations, clear evaluation frameworks, and governance built into the architecture from the start.

This whitepaper distills the key findings from those assessments. It presents a structured framework for evaluating and scaling generative AI, examines industry-specific use cases and technology patterns, and introduces a maturity model to help enterprises understand where they stand – and what it takes to move forward. It also offers a reference architecture for building agent-driven enterprises and strategic guidance on leveraging the combined strengths of AWS and Sonata Software to accelerate adoption while managing risk.

For business and technology leaders navigating the next phase of AI transformation, the insights here provide both a diagnostic lens and a practical roadmap – grounded not in theory, but in what we observed working (and failing) in the field.

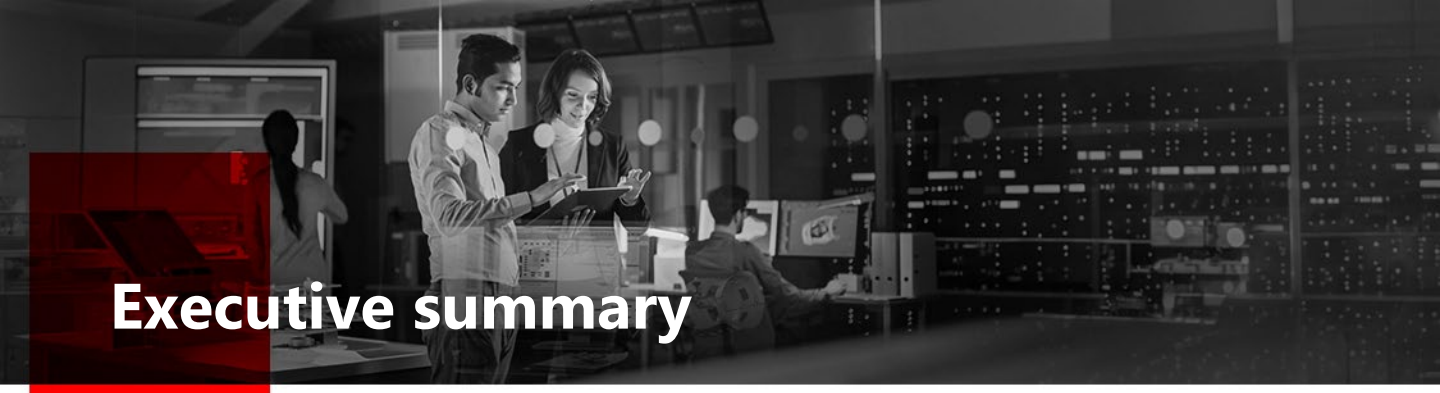

Executive summary

Generative AI has rapidly moved from experimentation to enterprise adoption. Organizations across industries are now looking to move beyond isolated AI pilots toward operational deployments that deliver measurable business outcomes.

Over the past year, Sonata Software conducted AWS Generative AI assessments across multiple organizations, spanning industries such as manufacturing, healthcare, retail, logistics, travel, and financial services.

What enterprises get wrong about GenAI

Based on our AWS GenAI assessments across industries, we observe that most organizations are not failing due to technology limitations - but due to operational and architectural gaps.

We believe the following:

1. **GenAI programs fail at ownership, not model selection:** The biggest gap is not "which model to use," but who owns prompts, knowledge, evaluation, and guardrails over time.
2. **The real bottleneck is AI lifecycle management:** Enterprises treat GenAI as a project, whereas it must be managed as a continuous product lifecycle (design → evaluate → monitor → improve).
3. **Agentic systems introduce a new category of risk:** Risk shifts from "hallucination" to uncontrolled actions across enterprise systems, requiring runtime controls not just policies.
4. **Data readiness is necessary but not sufficient:** Even with good data, organizations struggle without retrieval design, knowledge ownership, and observability.
5. **The winners will industrialize AI, not experiment with it:** Enterprises that standardize agent architectures, evaluation frameworks, and governance models scale significantly faster.

Industry analysts confirm the accelerating momentum.

According to McKinsey, *nearly 65% of organizations now actively use generative AI in at least one business function, a dramatic increase from less than 30% two years ago.*

Similarly, Gartner predicts that by 2027 more than *50% of enterprise applications will embed agentic AI capabilities.*

This whitepaper presents:

- A structured AWS GenAI assessment framework
- Industry-specific use cases and technology patterns
- A maturity model for enterprise AI adoption
- A reference architecture for agent-driven enterprises
- Strategic guidance for scaling AI using AWS + Sonata capabilities

Customer context lens – who this applies to

GenAI adoption patterns are highly sensitive to customer context. Use the lens below to calibrate priorities (use cases, governance depth, and architectural choices) based on whether an organization is an early adopter or a fence-sitter and based on enterprise profile and impact goal.

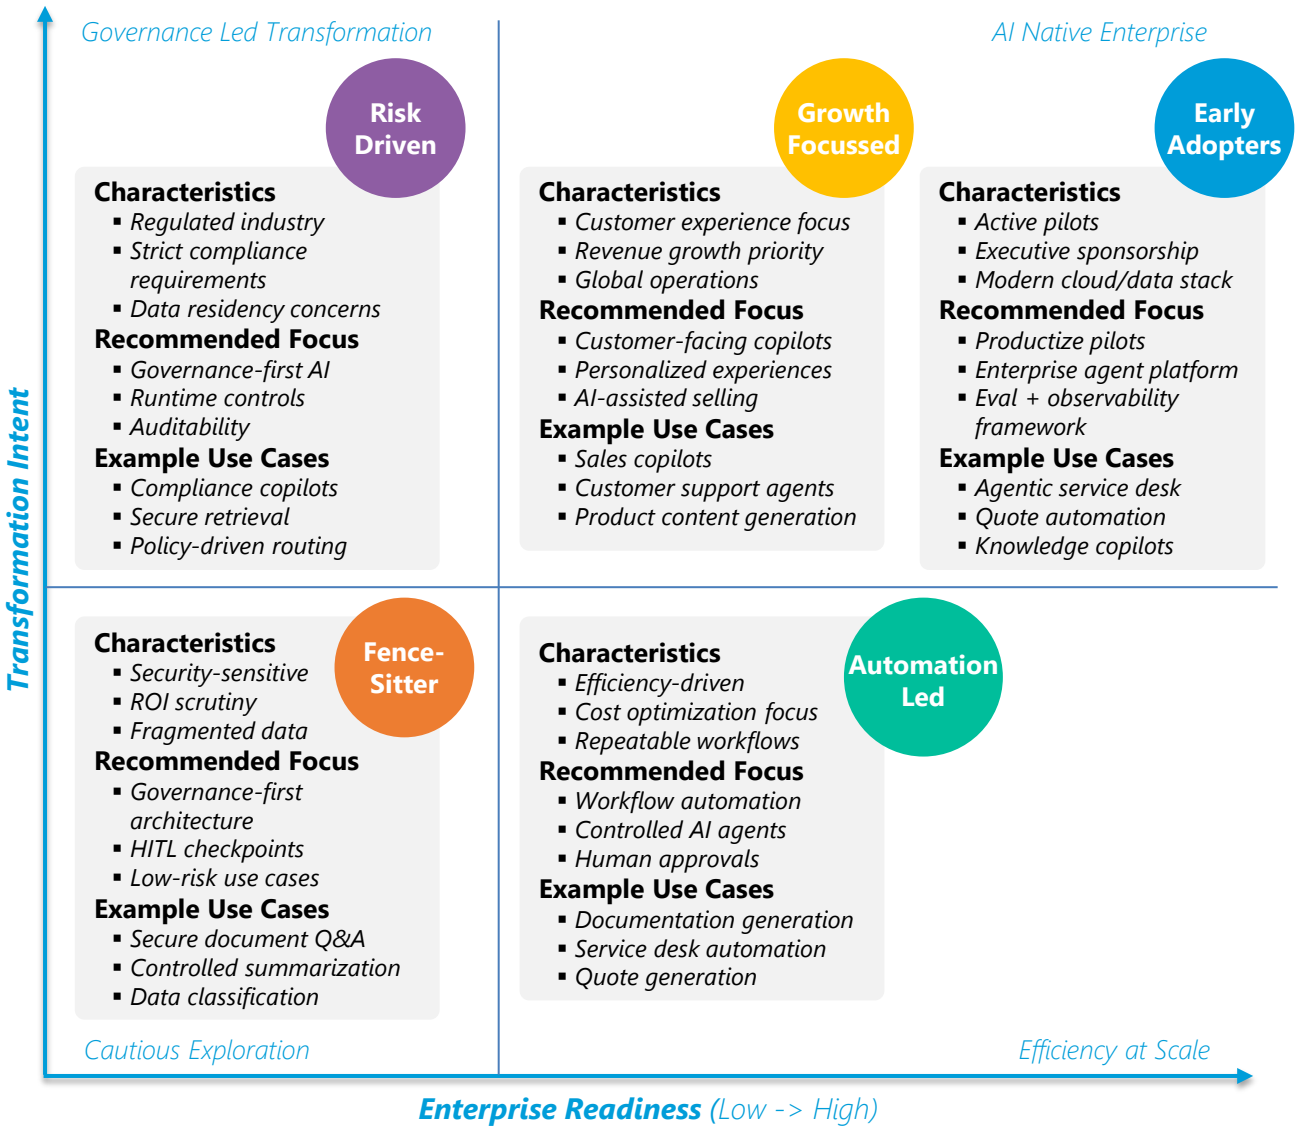

In practice, most organizations start by selecting 2–3 use cases aligned to one dominant impact category, then right-size governance and architecture to their industry and geography. The remaining sections outline a repeatable assessment framework and reference patterns that can be applied across these contexts.

The structured AWS GenAI assessment framework

Across all engagements, Sonata used a structured 4–6-week GenAI assessment framework designed to help organizations transition from AI exploration to implementation.

AWS GenAI Assessment Framework (5-Step Approach):

Stage	Objective	Key Activities	Outcome
Use case discovery	Identify high-impact, feasible GenAI opportunities	• Business and IT workshops • Evaluate use cases (impact, data, feasibility) • Shortlist priority use cases	Prioritized use case backlog aligned to business value
Data readiness assessment	Assess data maturity and availability for GenAI	• Review data architecture and sources • Evaluate document repositories • Check integration readiness • Assess governance and compliance	Data readiness baseline and key gaps identified
Technology and model evaluation	Select right GenAI architecture and tools	• Evaluate foundation models • Assess RAG, orchestration frameworks • Review security and governance capabilities	Recommended architecture leveraging AWS (e.g., Bedrock)
ROI and business case	Quantify value and prioritize investments	• Estimate cost savings and productivity gains • Assess implementation effort • Build ROI-backed prioritization	Ranked use cases with clear business justification
POC and roadmap planning	Enable execution and scale-up strategy	• Define architecture blueprint • Create POC plan and milestones • Establish governance model	Production-ready roadmap with execution plan

Combined strengths

AWS
- Scalable AI infrastructure
- Foundation model platform (Bedrock)
- Enterprise security and compliance

Sonata
- AI engineering expertise
- Industry-specific accelerators
- Enterprise modernization capabilities

Together, AWS and Sonata enable organizations to **accelerate AI adoption while minimizing risk.**

Industries impacted by GenAI assessments

The assessments covered a wide range of industries.

Industry	Example use case	GenAI capability	Business impact
Manufacturing	Sales quote automation	GenAI parses customer emails, validates pricing policies, and generates structured quotes automatically.	• Faster customer response time • Increased sales productivity • Reduced manual workload
Healthcare and life sciences	AI data discovery	GenAI automates dataset documentation and enables natural-language discovery across research datasets.	• Improved research productivity • Faster data discovery • Better governance and compliance
Retail	Process documentation automation	AI agents generate and maintain operational documentation through structured interviews with subject matter experts.	• Reduced documentation effort • Standardized operational processes
Logistics	Customer quote generation	AI agents automate quote requests by interpreting customer requirements and generating structured responses based on pricing policies.	• Faster response time • Reduced operational workload
Travel and hospitality	AI service desk agents	Conversational AI agents automate IT service desk tasks including password resets and knowledge retrieval.	• Reduced support tickets • Faster issue resolution

Assessment snapshots

1. Snapshot A - Early adopter (automation focus)

- Fragmented pilots across business units
- No standard governance or architecture

Insight:

Biggest risk was lack of ownership of prompts, knowledge, and guardrails

Outcome:

- Faster transition from POC → production
- 30–50% reduction in manual effort

2. Snapshot B - Fence-Sitter (risk focus)

- High concern around data leakage and compliance
- Previous AI initiatives stalled

Insight:

Primary blocker was data readiness and governance clarity

Outcome:

- Improved stakeholder confidence
- Defined roadmap for safe adoption

3. Snapshot C - Consumer Enterprise (growth focus)

- Rising customer support costs
- Fragmented knowledge bases

Insight:

Knowledge quality mattered more than model choice

Outcome:

- Faster response times
- Improved customer experience consistency

GenAI Use Case Landscape Map

Across engagements, use cases clustered around several functional areas.

Heat Map

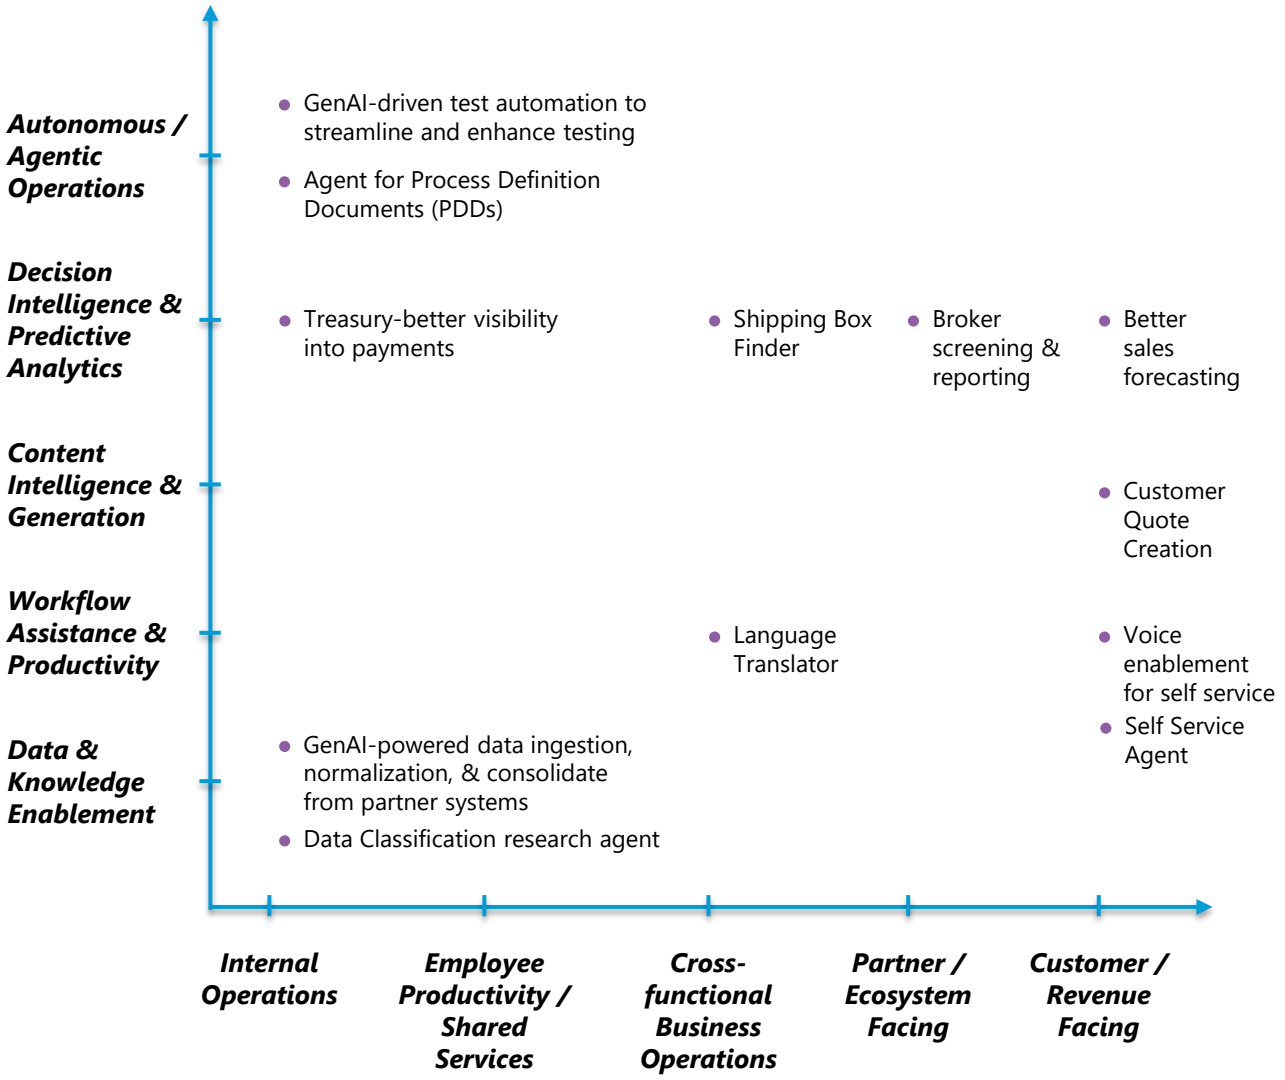

This landscape illustrates that **GenAI adoption spans both customer-facing and internal operational processes.**

However, the strongest early adoption patterns occur in **customer support, document automation, and data engineering.**

Technologies evaluated across assessments

Organizations evaluated technologies across the AWS ecosystem.

Core AI platform

Most enterprises adopted AWS Bedrock, enabling access to foundation models including:

- Anthropic Claude
- Amazon Titan
- Mistral

Document intelligence

Document processing capabilities were enabled using:

- AWS Textract
- Amazon Comprehend

Data platforms

Common data architecture components included:

- Amazon S3
- DynamoDB
- AWS DataZone
- OpenSearch

Workflow orchestration

AI workflows were implemented using:

- AWS Lambda
- AWS Step Functions
- Amazon EventBridge
- LangChain

Enterprise AI maturity model

Across organizations, GenAI adoption follows a five-stage maturity journey.

AI Maturity Model

GEN AI MATURITY LEVELS	Level 1 Exploration	Level 2 Incorporation	Level 3 Proliferation	Level 4 Optimization	Level 5 Sustainable Transformation
The Life Cycle					
AI Governance	Awareness	AI model identification and selection	AI model customization or fine tuning	Process automation and governance	Ethical and sustainability considerations
Data Management	Know the pain points, the Data discovery	Data quality assurance and enablement	Data mobility standardization	Data is a product and sharable	Data products are to be monetized
People	What to What: the skill gap assessment	Skills mapping and enablement	Skill cross-pollination and accessibility to all	AI driven Skill policies at the Organization level	AI mindset and culture
Process	Some standard somewhere	Business outcomes-oriented Business process	Business process operationalization	Business process automation and self-sustenance	Innovation and sustainability-driven Business process
Technology (AI Enablers)	Legacy or no AI tools or infrastructure	Need-based AI tools and infrastructure	CoE team and culture enablement	AI automation and AI/MLops enablement	Sustainable solutions feedback loop

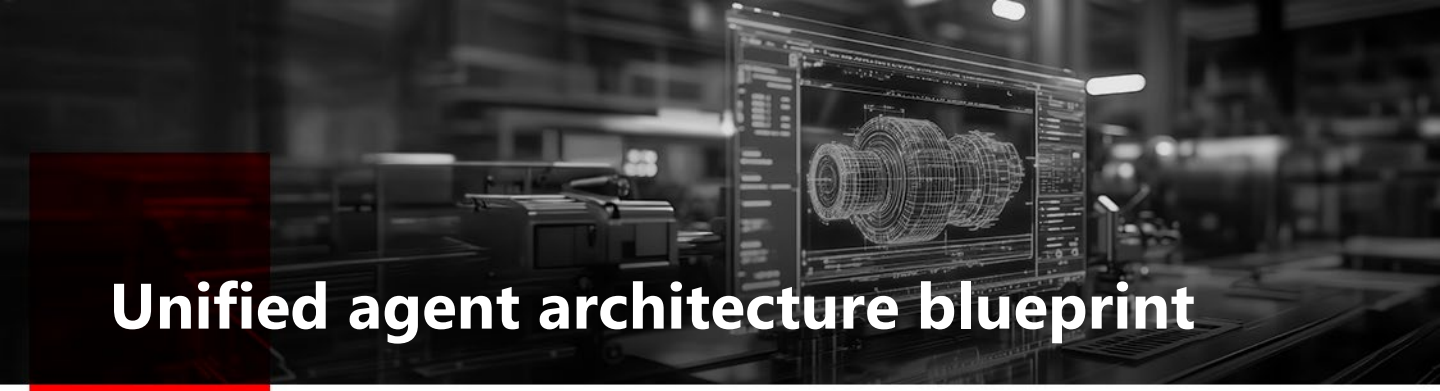

Unified agent architecture blueprint

To scale GenAI across the enterprise, organizations require a standardized agent conceptual architecture.

Architecture

Business Applications & Workflow Automation

- Enterprise Applications Integration
- AI-Driven RPA
- Process Intelligence & Optimization
- Digital Worker Assistants
- AI-Driven Decision Making
- Reducing Human Workload
- Improving Efficiency

Decision Intelligence & Autonomous Agents

- Agentic AI Orchestrator
- Decision Intelligence Engine
- Multi-Agent Collaboration
- Autonomous Execution & Learning
- AI Governance & Compliance Framework
- Adaptive AI with Compliance Framework

AI Core (LLM & Knowledge Processing Layer)

- LLM Model Services
- Memory & Context Manager
- Skill Execution Engine
- Model Registry
- Fine Tuning & Adaptation
- AI Orchestration
- Self-Optimizing AI with Continuous learning

Data Services

- Internal Data Sources
- External Data Sources
- Data Processing & Storage
- Historical + Real-Time data access
- ETL Pipelines, Vector Embeddings, Graph DBs
- ....

Secure Foundation

- Advance Machine Learning/Deep Learning
- Observability
- Natural Language Processing (NLP)
- Reinforcement Learning
- Edge AI
- Blockchain Integration
- Robotics
- Simulation & Modelling

Role	Responsibility
Agent product owner	Defines use cases, KPIs, and lifecycle
Prompt / eval engineer	Designs prompts, evaluates performance
Data steward	Owns knowledge sources and data quality
Platform / engineering team	Builds and scales architecture
Risk and compliance	Defines policies and governance controls

Enterprise-ready agent checklist:

To move from pilots to production, enterprises must implement the following controls:

1. Human-in-the-loop checkpoints for high-risk actions
2. Role-based access and tool permission
3. Audit logs for all agent interactions
4. Tool allowlists and API restrictions
5. Data boundary enforcement (PII, confidential data)
6. Evaluation gates before production rollout
7. Prompt and knowledge versioning
8. Monitoring for drift, latency, and failures
9. Feedback loops for continuous improvement
10. Incident escalation mechanisms

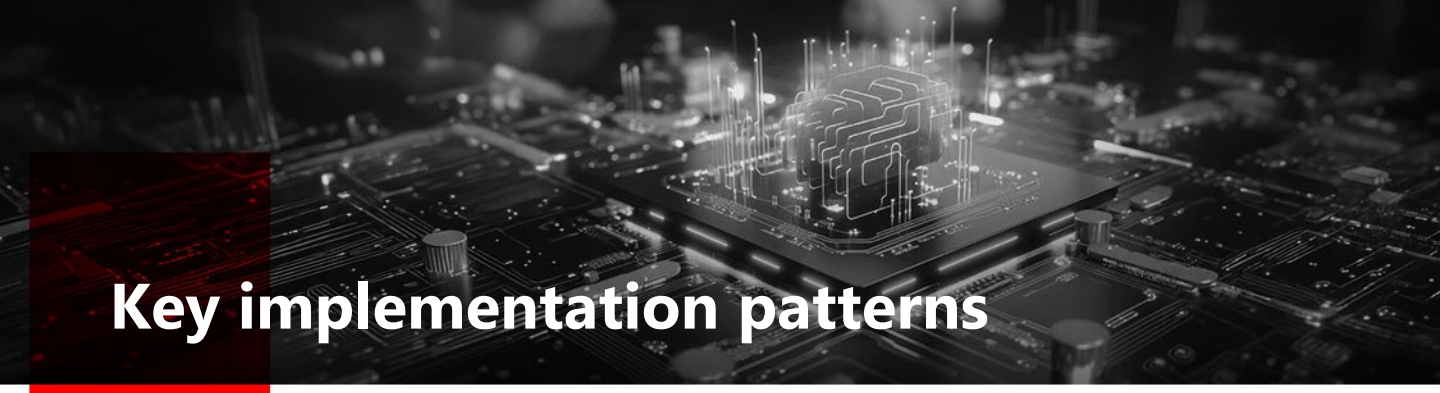

Key implementation patterns

Across engagements, several patterns consistently emerged:

Rise of AI agents

Organizations are moving toward **autonomous AI agents** capable of executing workflows.

Retrieval augmented generation

Most solutions use **RAG architectures** to combine foundation models with enterprise data.

Governance first

Security, privacy, and compliance requirements strongly influence architecture decisions.

Efficiency use cases first

Early adoption focuses on:

- documentation automation
- service desk automation
- process optimization

Across anonymized assessments, organizations that treated agents as long-lived products, not one-time implementations, scaled faster and with fewer risk escalations.

Outcomes observed

Across engagements, organizations observed measurable benefits:

Operational efficiency improvements:

30-60%
reduction in manual effort

Service desk automation:

20-40%
reduction in support tickets

Process acceleration:

50%
faster documentation cycles

Knowledge accessibility:

90%
of enterprise knowledge searchable

Customer experience improvement:

significant reduction
in response times

The future: Agent-driven enterprises

The next phase of enterprise AI will be defined by agent-driven automation—and early signals are already visible across the enterprise. Organizations are progressing from isolated GenAI experiments toward AI agents that can reason, act, and orchestrate workflows across business systems. As this transition accelerates, platforms like Amazon Bedrock AgentCore are emerging as the foundational layer to move agents from proof-of-concept to secure, observable, enterprise-scale execution. Instead of isolated AI tools, organizations will deploy **multiple collaborating AI agents** capable of executing complex workflows.

Examples include:

- documentation agents
- operations agents
- research agents
- customer support agents

These agents will operate within **enterprise agent orchestration platform,** enabling organizations to transform business operations at scale

Conclusion

The findings from these AWS Generative AI assessments point to a clear and urgent conclusion: the window for strategic differentiation in enterprise AI is open now, but it will not stay open indefinitely.

Organizations that have moved beyond pilots – those that have built the data foundations, governance structures, and operating models required for production-grade AI – are already pulling ahead. They are capturing measurable gains: 30–60% reductions in manual effort, 20–40% fewer support tickets, and dramatically faster knowledge access and process cycles. More importantly, they are building institutional muscle that compounds over time.

The central lesson from these engagements is not about technology selection. Foundation models are increasingly commoditized, and the gap between the best and second-best model is shrinking. What separates AI leaders from laggards is execution discipline – the ability to define ownership, manage the AI lifecycle as a continuous product, enforce governance at runtime, and design systems for production from day one rather than retrofitting them after the fact.

Looking ahead, the emergence of agentic AI architecture represents the next major inflection point. Enterprises that standardize now on robust agent frameworks, evaluation pipelines, and cross-functional operating models will be best positioned to deploy the multi-agent systems that will define the next generation of digital operations.

The journey from AI experimentation to enterprise impact is not a single project – it is an ongoing transformation. The organizations that approach it with that mindset, backed by the right partners, architecture, and governance, will be the ones that emerge as AI-native enterprises: faster, more adaptive, and structurally more competitive in an increasingly intelligent business landscape.

About the author

Kasturi Sinha

Sr. Manager - CTO office strategy and business consulting

Kasturi Sinha is a business consulting and AI strategy professional at Sonata Software, specializing in Generative AI transformation, enterprise modernization, and cloud-driven digital strategy. Her work focuses on GenAI assessments, AI operating models, agent-driven architectures, cloud modernization, and enterprise transformation initiatives leveraging AWS and Microsoft ecosystems.

SONATA SOFTWARE

The Modernization Engineering Company

ABOUT SONATA SOFTWARE

Sonata Software is a global IT solutions firm focused on catalysing business transformation initiatives of its clients through deep domain knowledge, technology expertise and customer commitment. The company delivers innovative solutions for Travel, Retail & Distribution and Software Product companies through IP based Platforms, Products and Services, that bring together new digital technologies such as Omni-channel commerce, Mobility, Analytics, Cloud and ERP, to drive enhanced customer engagement, operations efficiency and return on IT investments. A trusted long-term service provider to Fortune 500 companies across both the software product development and enterprise business segments, Sonata seeks to add differentiated value to leadership who want to make an impact on their businesses, with IT.

info@sonata-software.com

www.sonata-software.com