

May 21, 2026

Agentic AI in regulated industries

Learnings from Generative AI applications in Fintech

Nuwan Bandara, Head of Technology – Fintech, AWS

Enterprises are doubling down on Agents

340B

Annual GenAI value in
Banking¹

33%

of enterprise software apps
will include agentic AI by
2028, up from less than 1%
in 2024²

15%

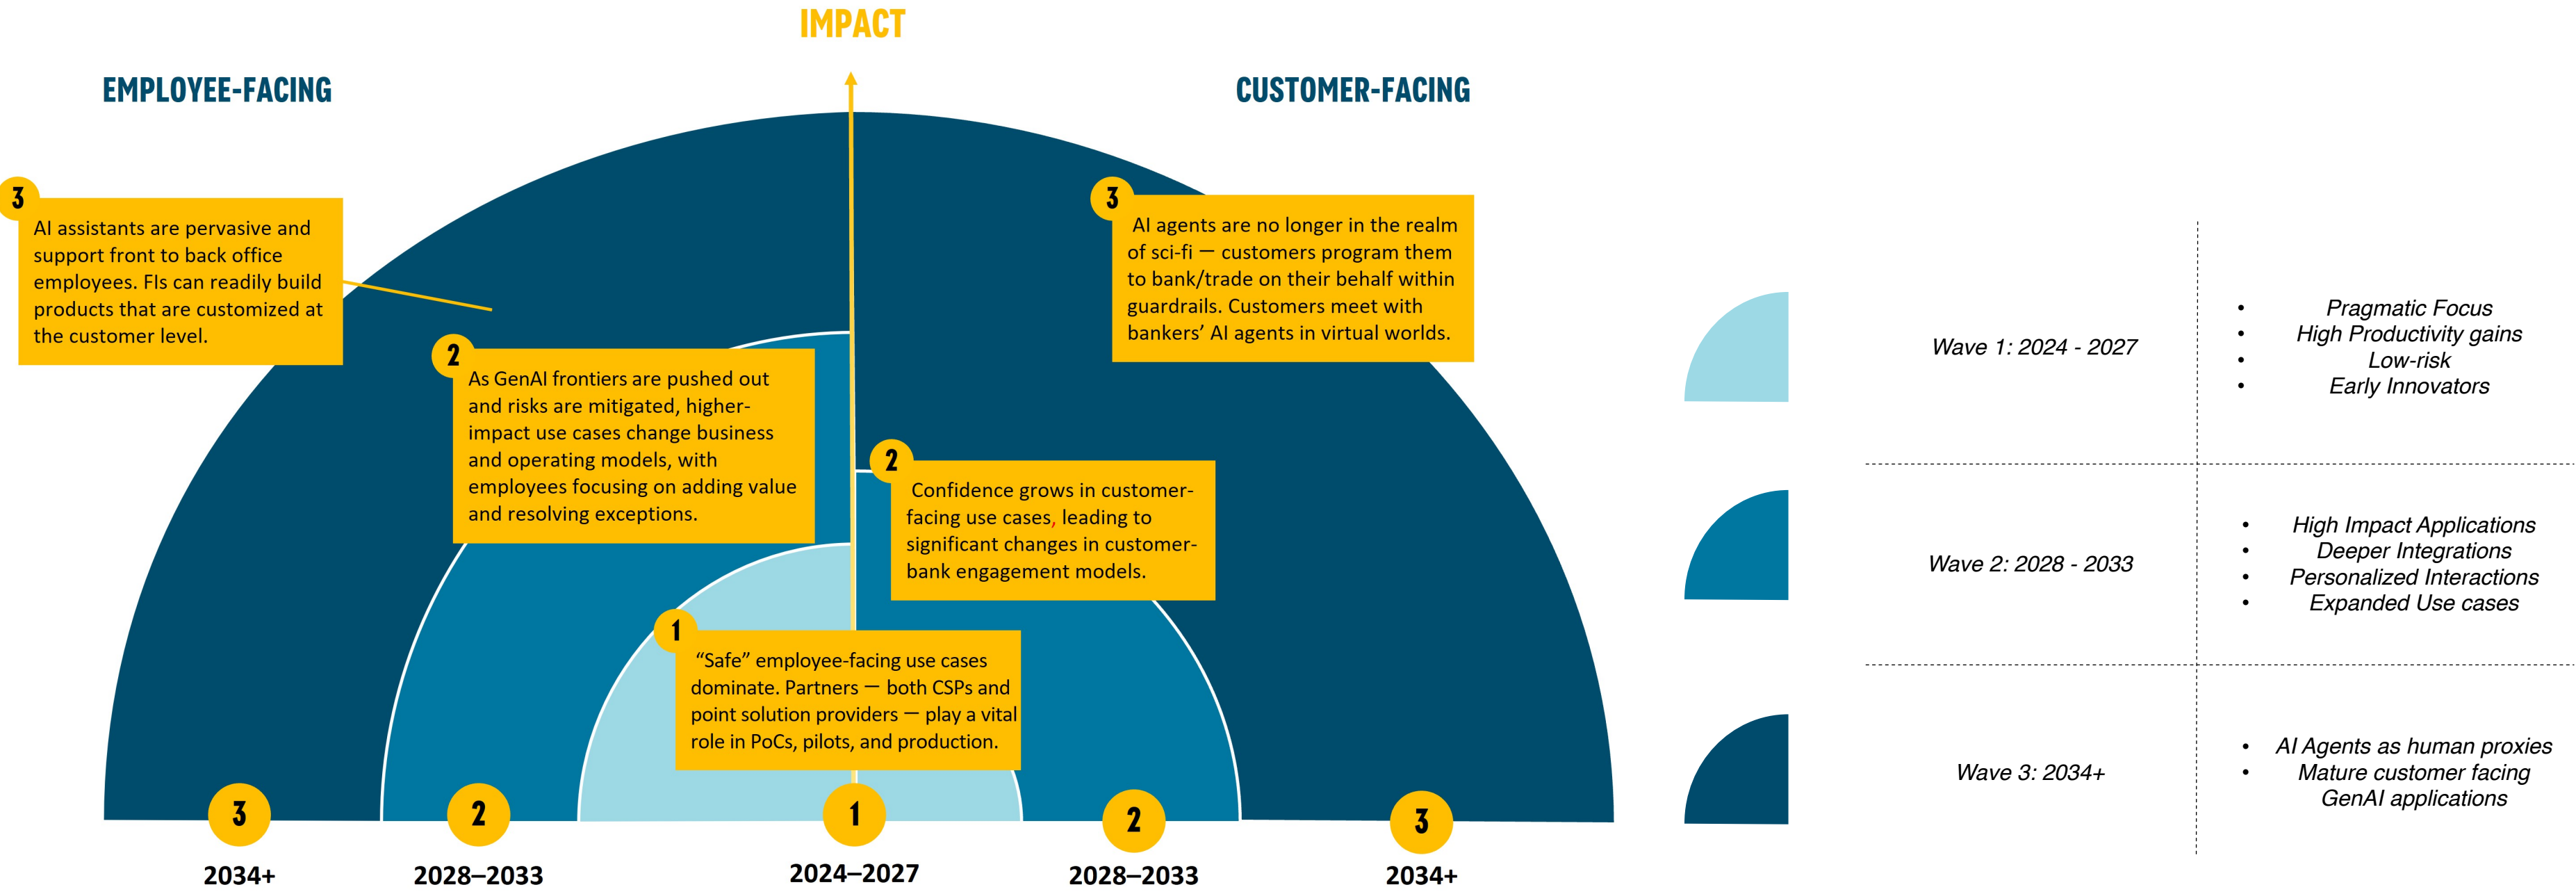
of day-to-day work
decisions will be made
autonomously through
agentic AI by 2028³

¹Mckinsey & Company: [Economic Potential of Generative AI](#)

²Gartner, "Top strategic Technology Trends for 2025," October 2024

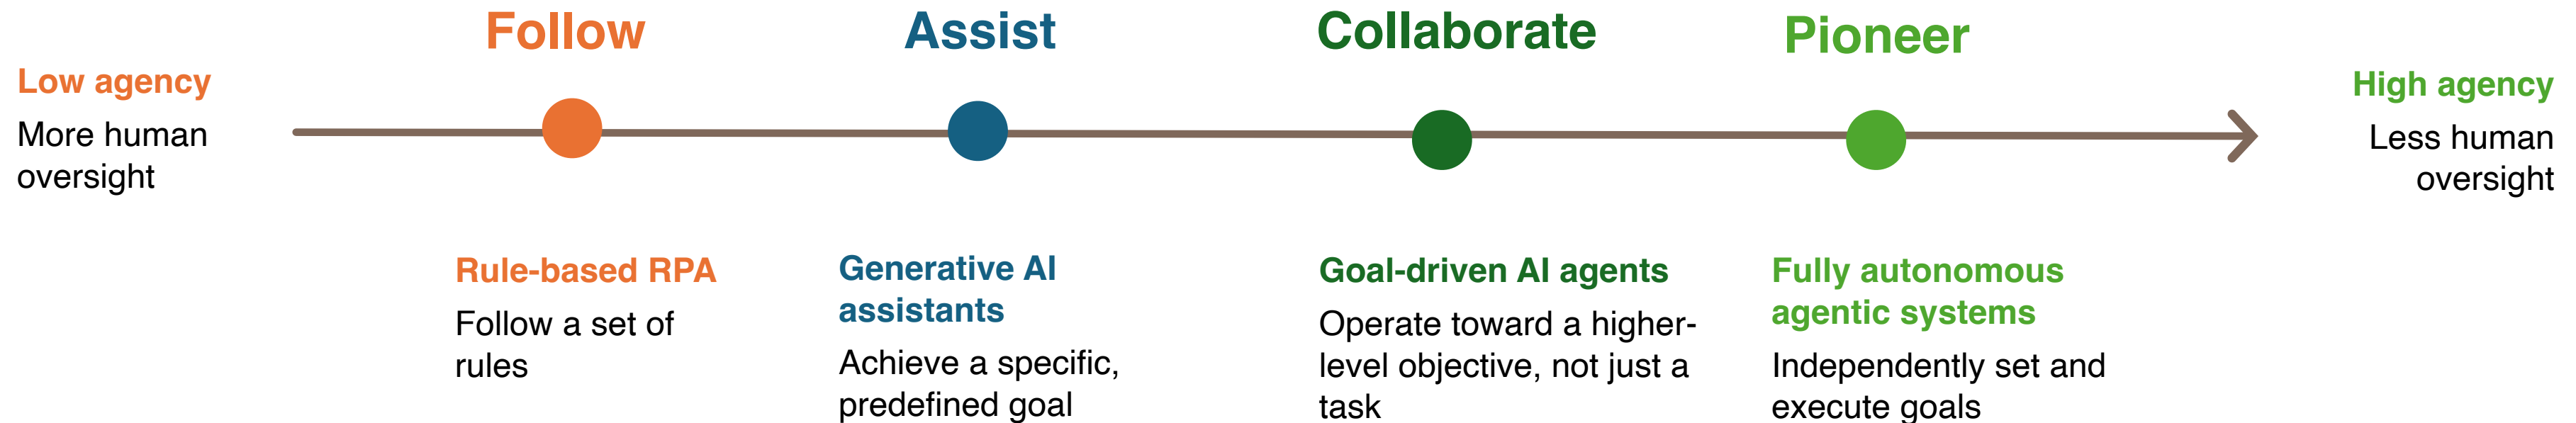
³Gartner, "Top Strategic Technology Trends: Agentic AI—the Evolution of Experience" February 2025

Generative AI Making Waves



Source: [Celent: Generative AI Making Waves](#)

AI maturity scale – from rules to reasoning



Agentic AI will help realize the true promise of LLMs

Gen AI


Primary purpose: Content generation
Dependency: Prompts and input data
Agency level: None (requires more human prompting)
Ability to execute workflows:

Agentic AI

Primary purpose: Goal achievement
Dependency: Access to tools, data and agents
Agency level: Higher (less to no human oversight)
Ability to execute workflows: Medium / high

Key areas of driving business value with agentic AI


Improve
Workplace
Productivity




Accelerate
Software
Development



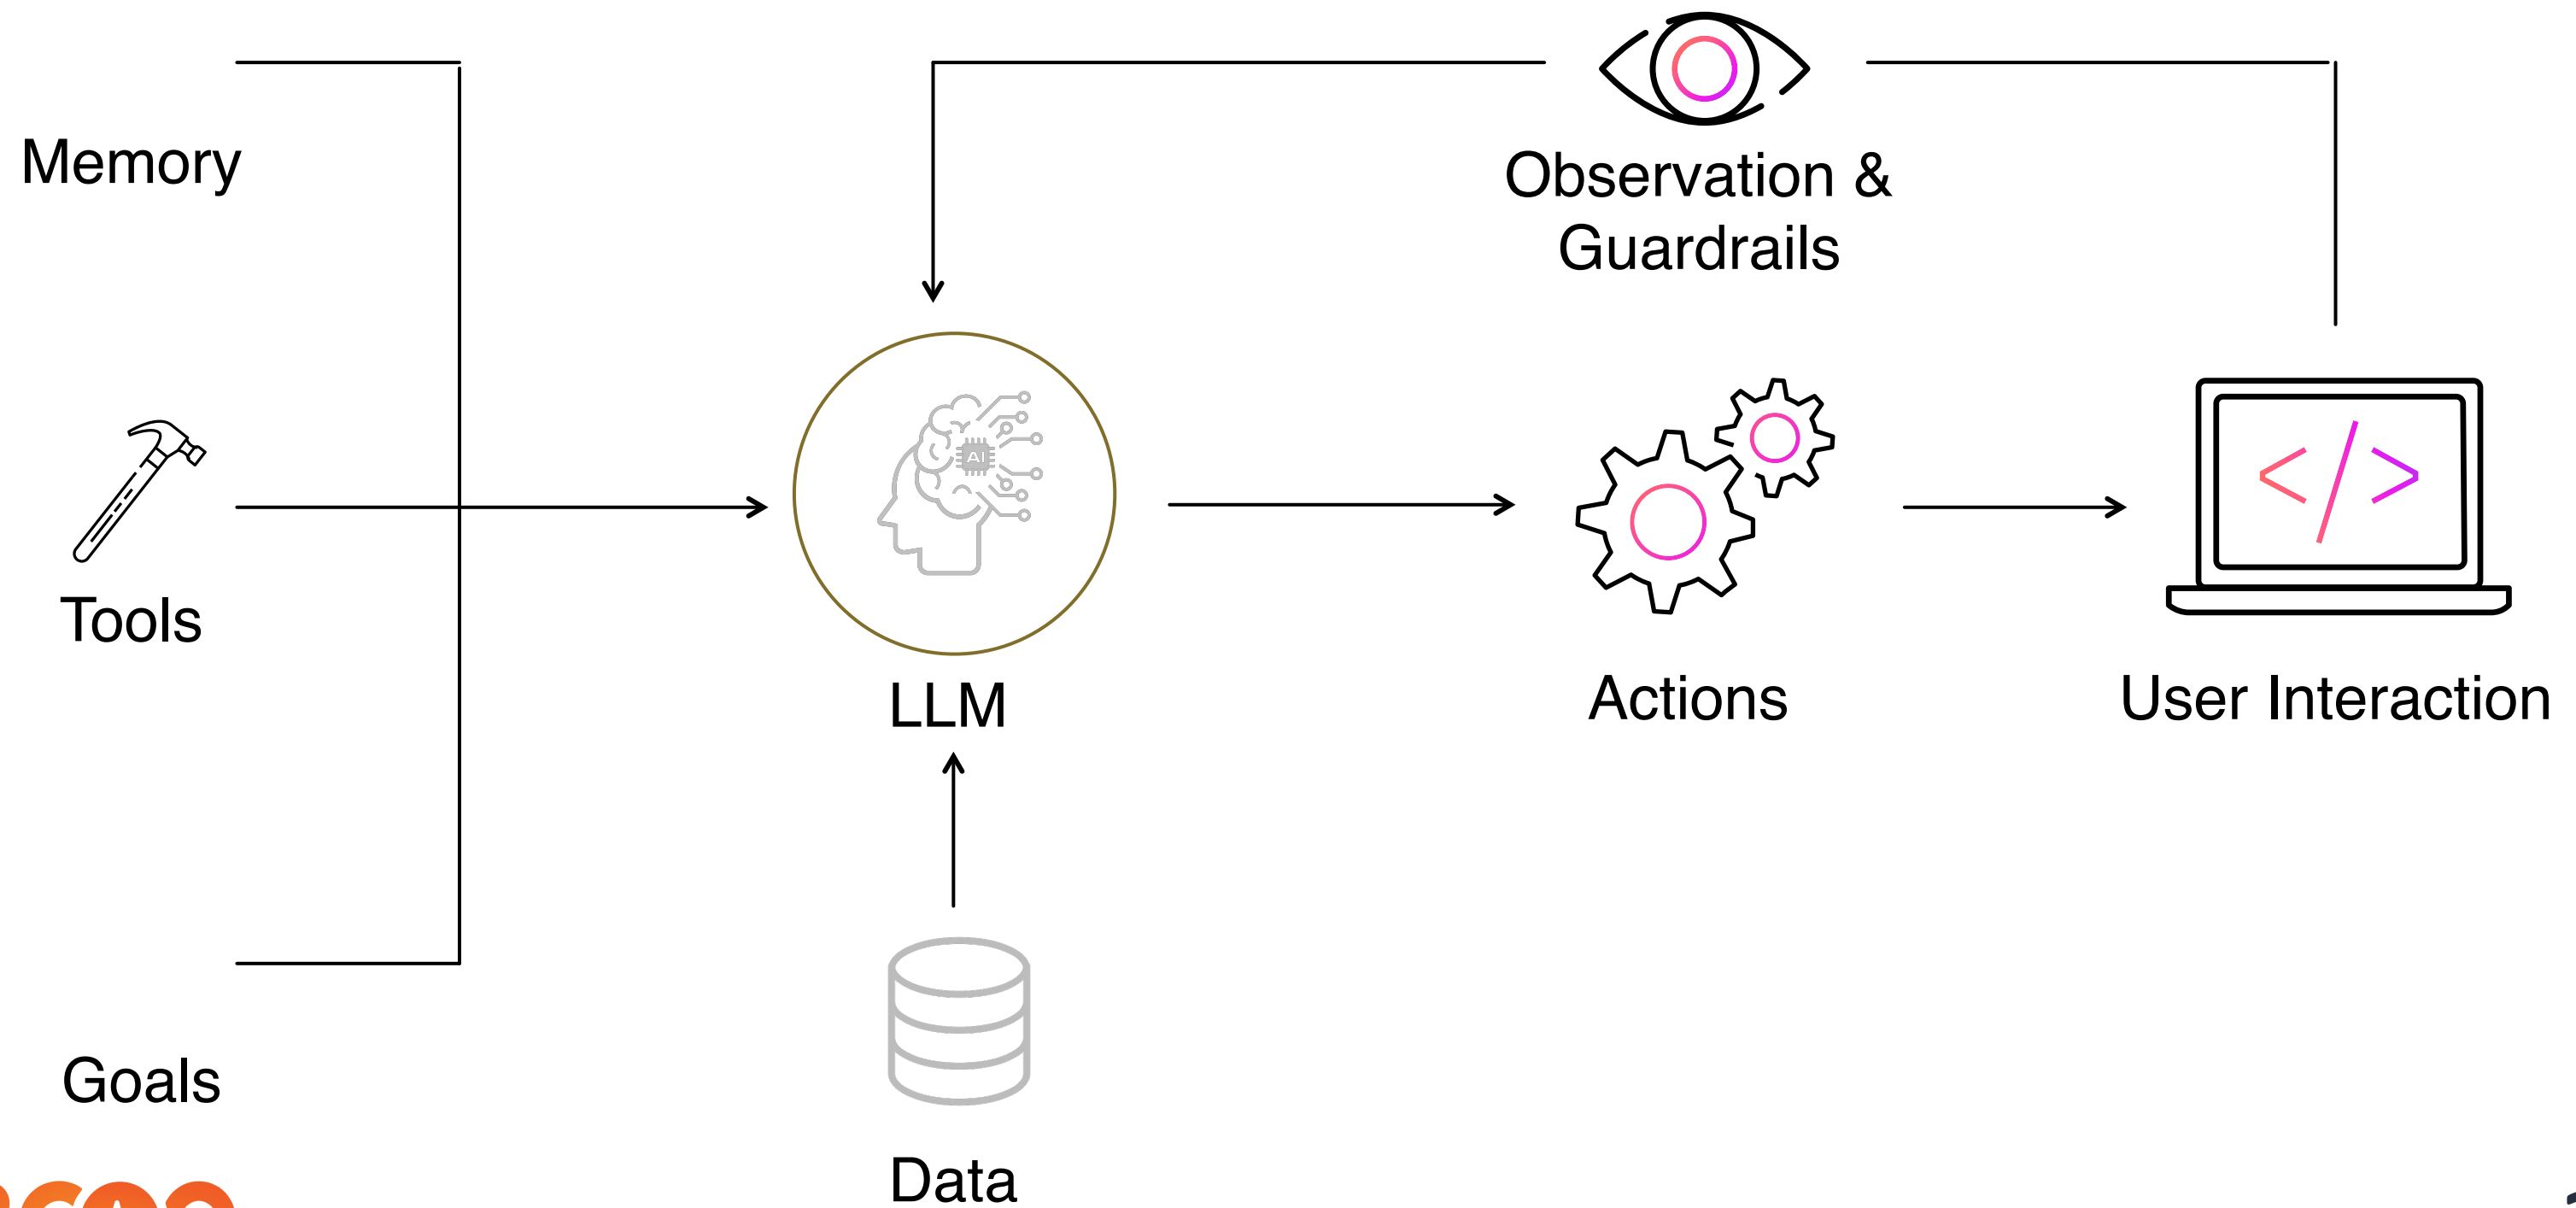
Automate Business
Workflows with
Custom Agents



Transform your
business with
Agentic OS



Building custom agents to reimagine workflows



Finance and risk organizations are facing challenges

Siloed systems

Separate and proprietary systems lead to fragmented data and limited integration

Limited scalability

Regulations stipulate increased on-demand modeling and scenario planning for compliance and risk assessments

Constantly changing regulatory environment

Increased scrutiny on risk, compliance, and transparency with shorter timelines to comply with new regulations

Manual controls

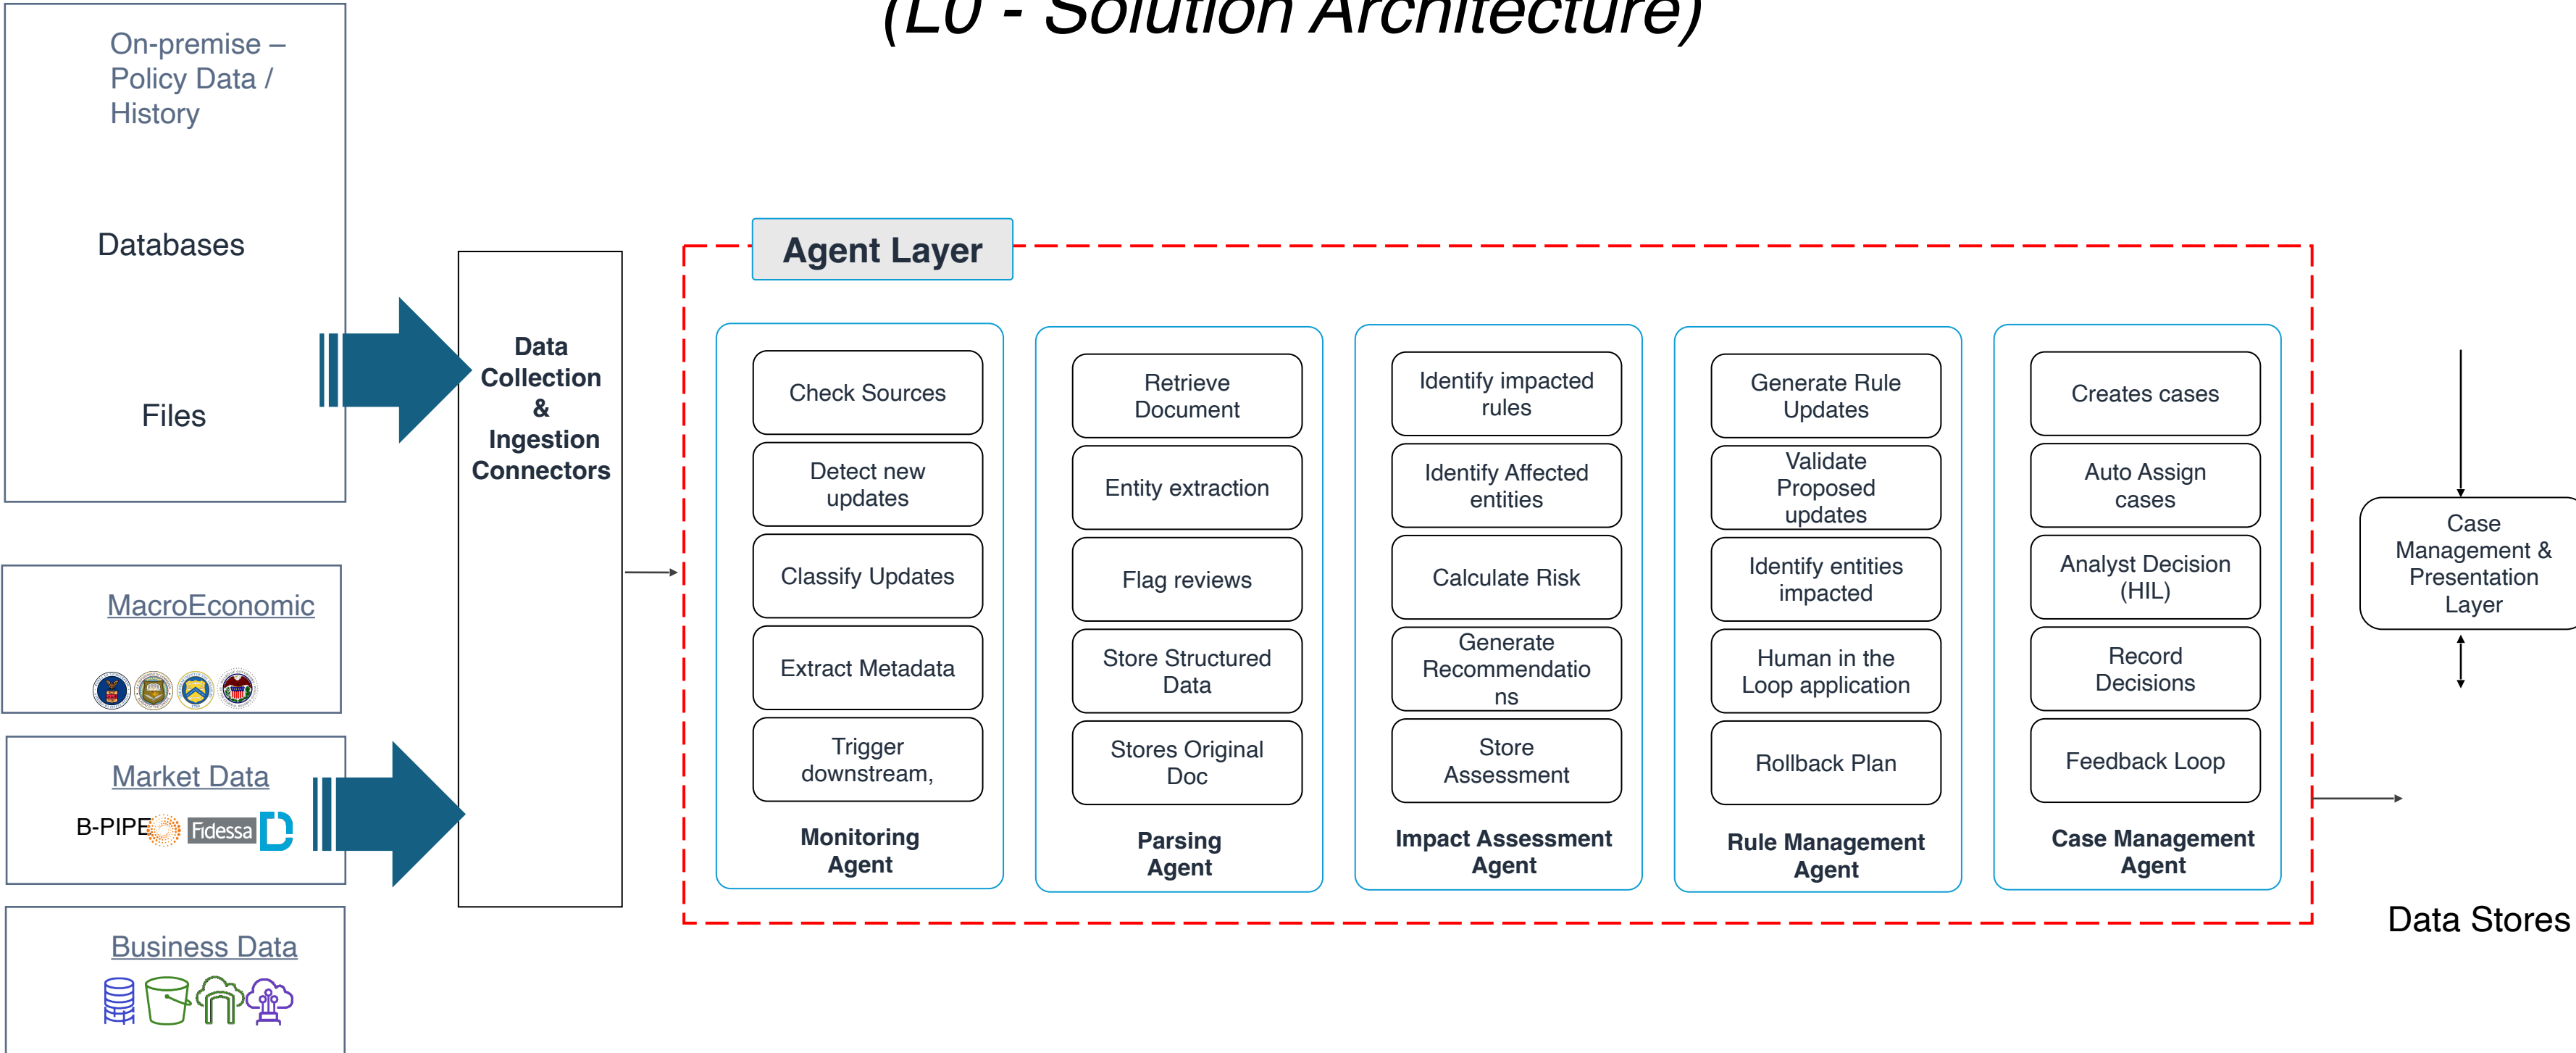
Processes are largely manual, requiring significant human intervention for monitoring, reporting, and auditing.

Slow product innovation

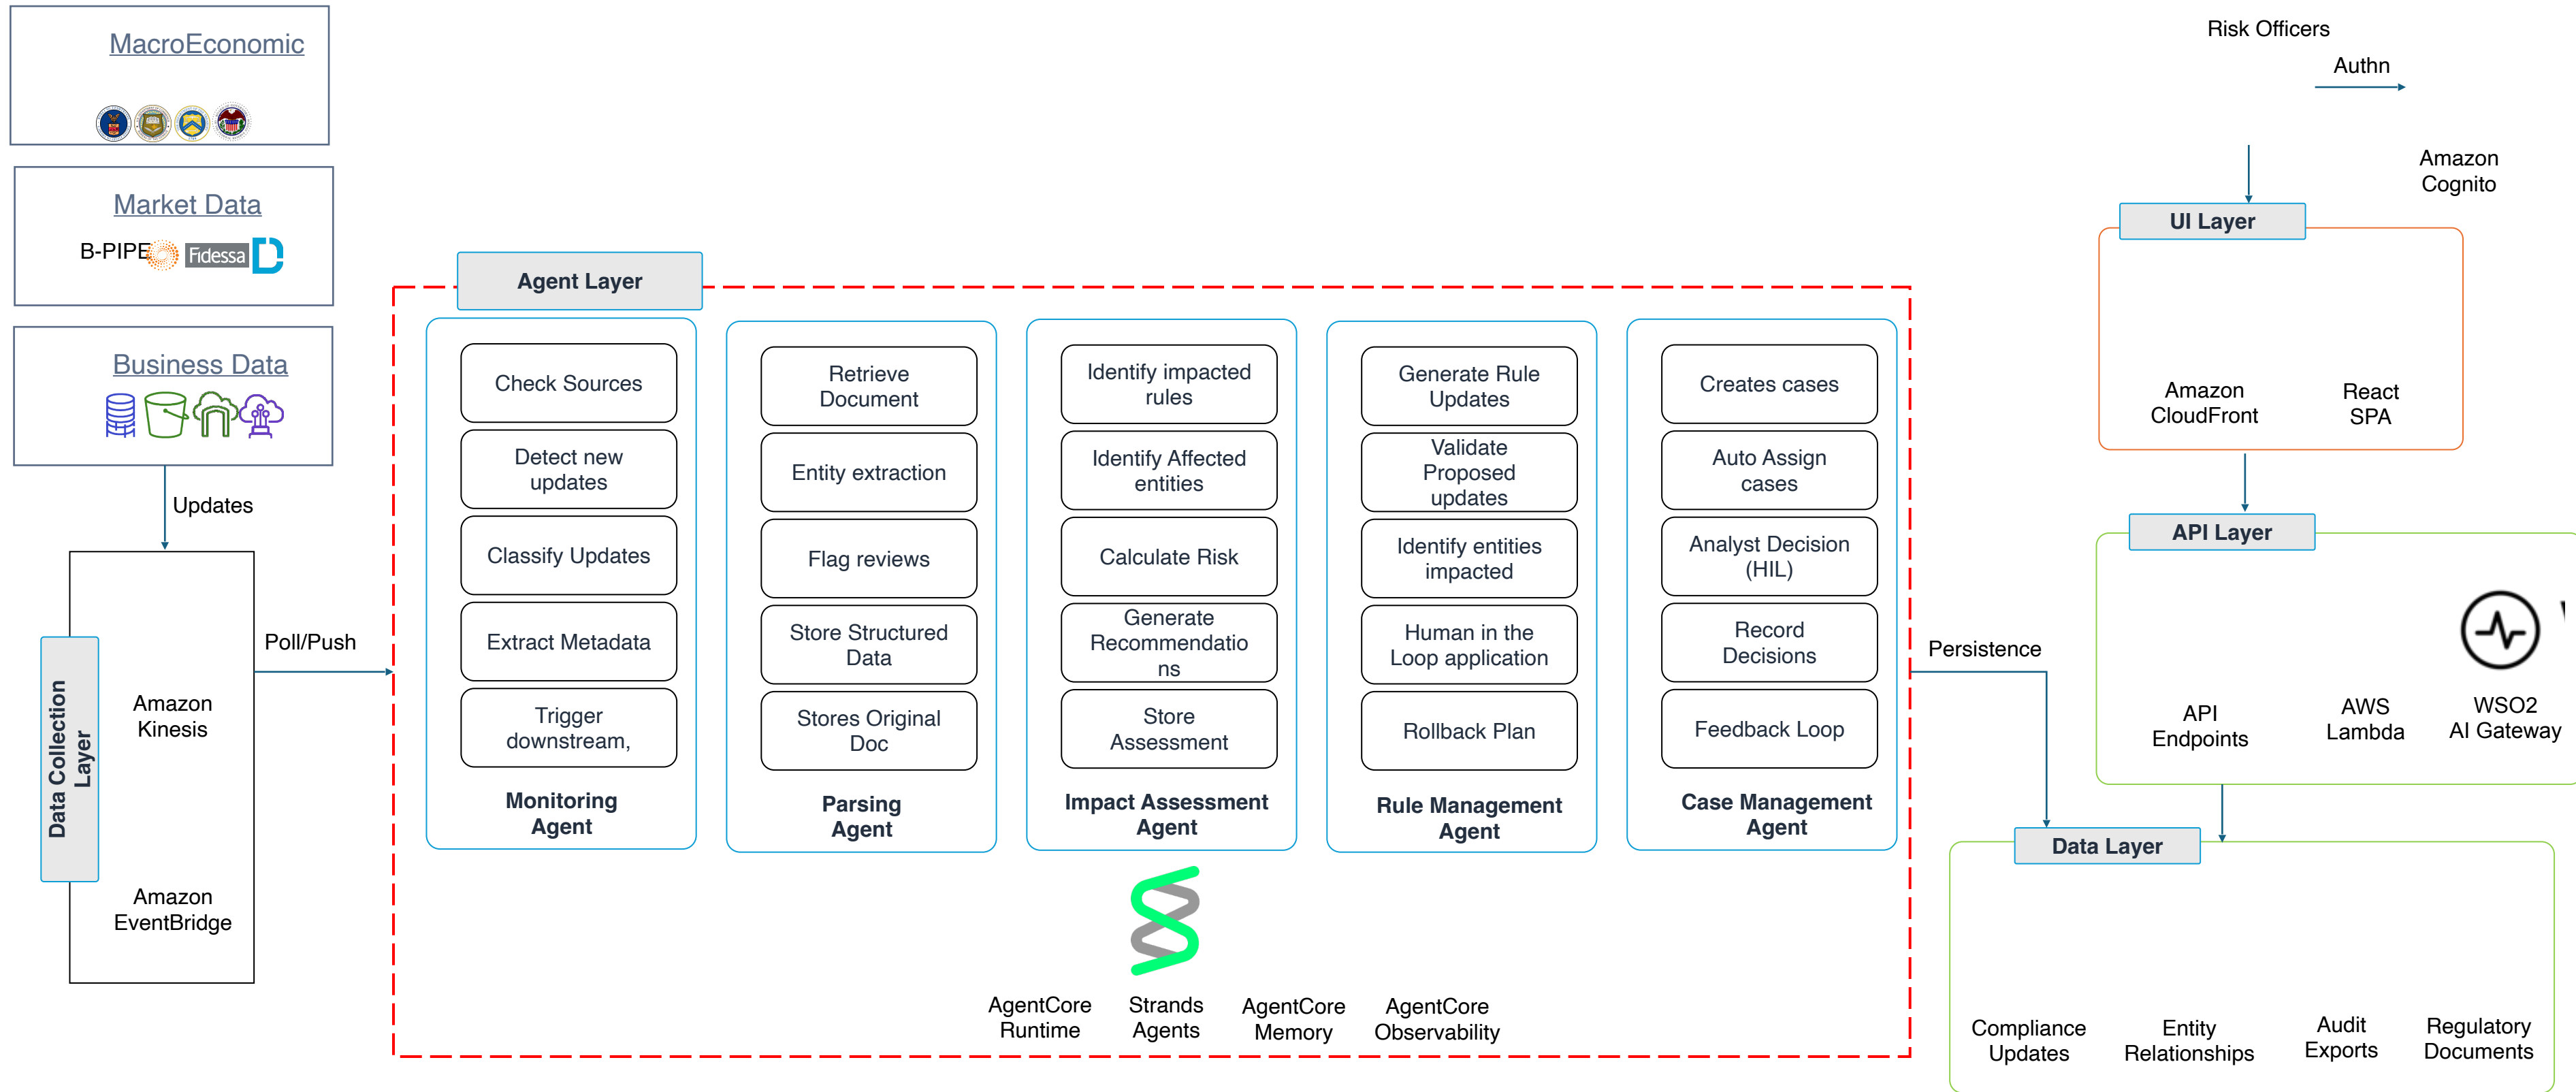
Aggregating and validating financials for new products and the associated risk delays release of new products

AI-Powered Compliance Monitoring & Automation Platform

(L0 - Solution Architecture)



L1 Architecture



When agents act, who is accountable?

Regulators require firms to name a responsible human for AI-driven outcomes.

REGULATORY DIRECTION

FINRA 2026 Oversight Report: GenAI reclassified as "supervised technology." AI governance is explicit enforcement priority. Firms must modify supervisory procedures.

FINRA Rules 3110/3120: AI-driven workflows that perform steps a registered person would supervise = supervisory failure.

SEC Investor Advisory Committee (Dec 2025): Voted to recommend AI disclosure framework for investment advisers.

Singapore IMDA (Jan 2026): Model AI Governance Framework for Agentic AI — addresses delegation chains and multi-agent coordination.

If AI materially influences an outcome, it must be governable to the same standard as a human decision-maker.

EVERY AGENT MAPS TO A HUMAN

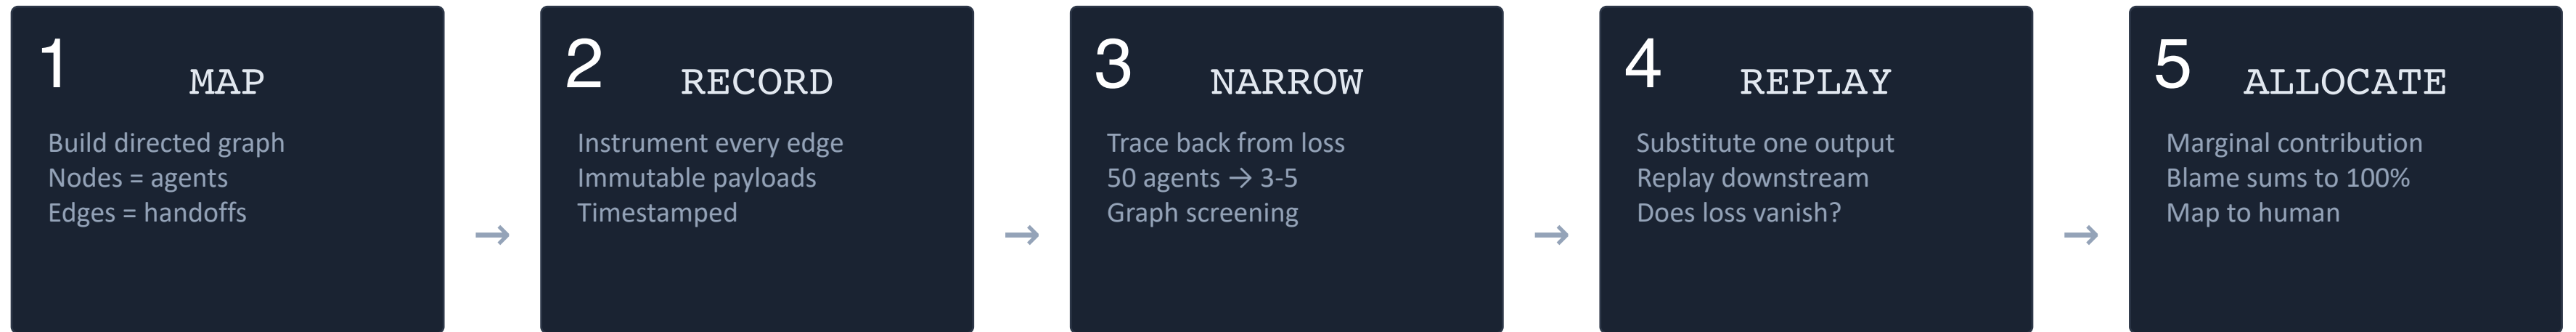
FAILURE MODE	ACCOUNTABLE HUMAN
Agent produces incorrect decision	Agent Builder (quant/engineer)
Policy boundary fails to prevent loss	Policy Author (CIO / risk committee)
Agent drifts or degrades over time	Model Owner (designated supervisor)
Interaction between agents causes harm	System Architect (topology designer)
Agent omits action within its scope	Product Owner (scope definer)
Monitoring fails to catch the above	Head of Model Governance

Loss → Causal Agent (via pipeline) → Responsible Human (via ownership map)

"The system failed" is not an answer. Every allocation terminates at a named person.

Part of AI observability is agent accountability

Five sequential steps to determine responsibility for errors in multi-agent systems



Step	Question Answered	Without It
1 Map	How are agents connected?	You investigate all agents equally — no leads
2 Record	What exactly happened at each handoff?	Reconstructing from memory
3 Narrow	Which agents could have caused this?	You replay 50 agents instead of 5
4 Replay	Which agent's error actually caused the loss?	Blame based on proximity, not causation
5 Allocate	How much blame does each agent deserve?	One agent gets 100% when three contributed

Steps ①–② exist in AgentCore today. Steps ③–⑤ are the accountability gap.

Why Agentic AI Solution?

01

Autonomous Decision-Making with Context

02

Natural Language Understanding

03

Tool Use and External Integration

04

Multi-Step Reasoning Chains

05

Continuous Learning from Human Feedback

The Blueprint

Master the Fundamentals: Optimize your core processes, vendor procurement, and developer workflows before scaling.

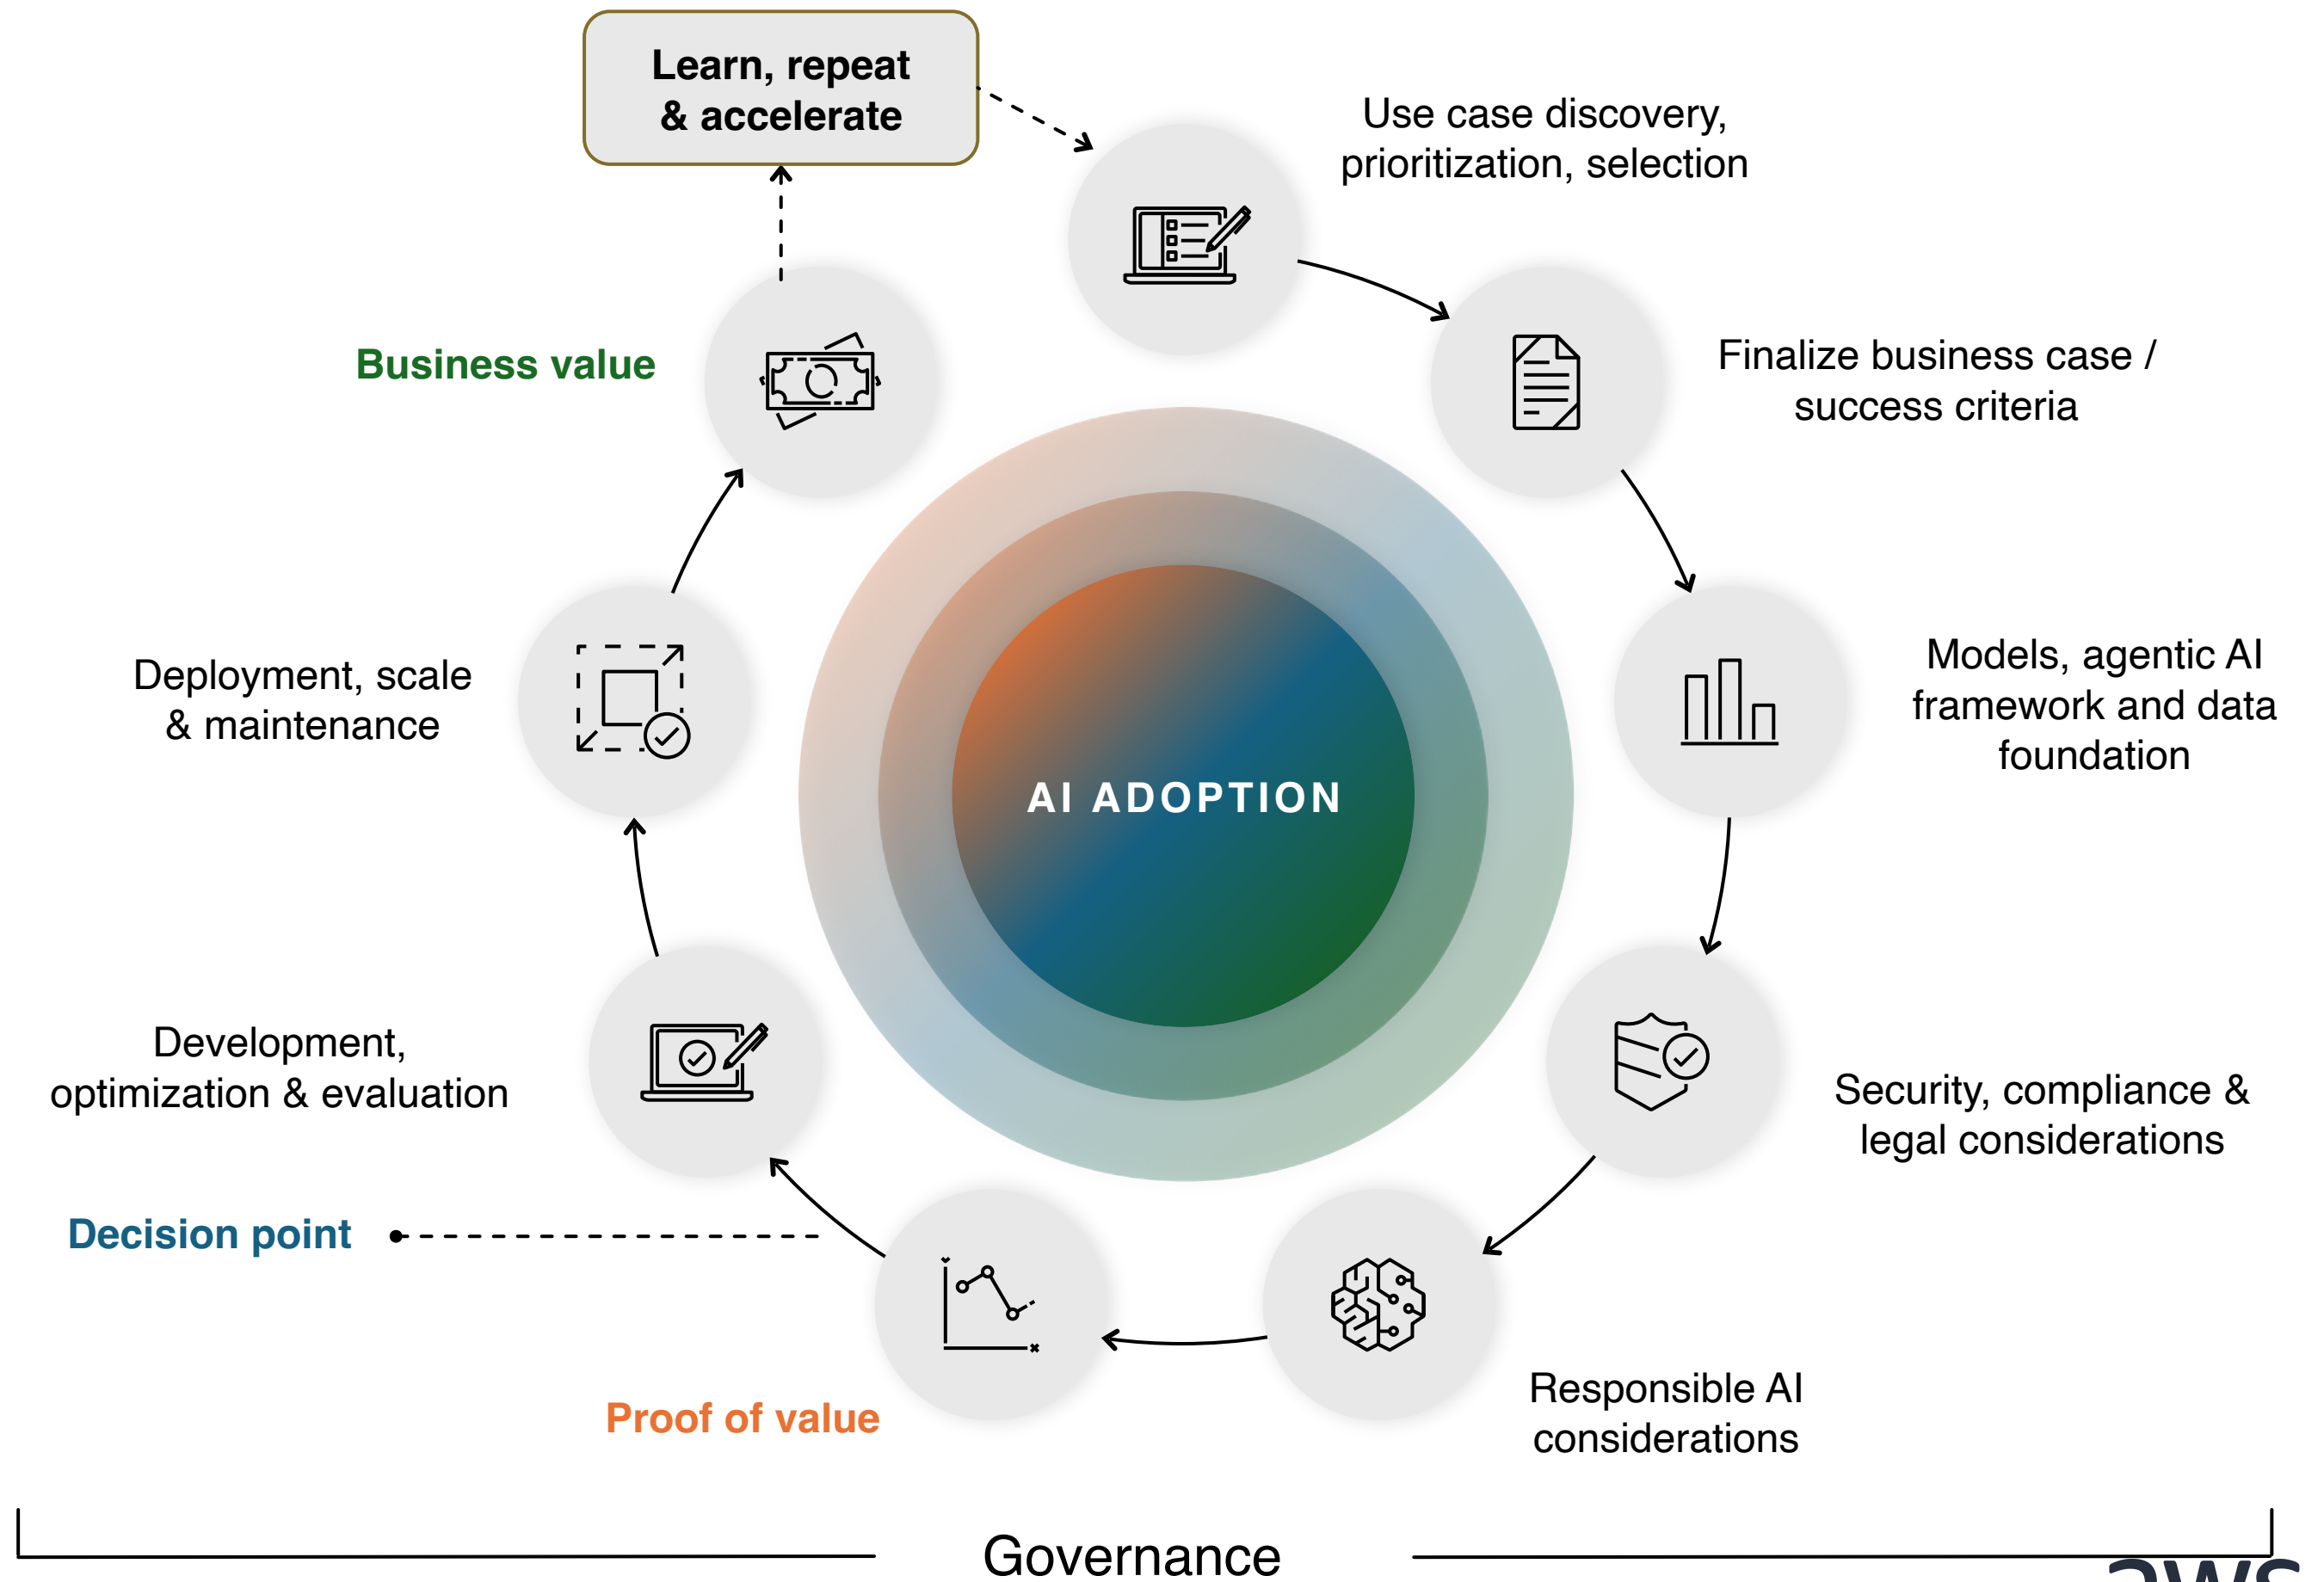
Strategy Over Hype: Invest in building applications transform your businesses over highly-engineered platforms

Iterate Constantly: Technology moves too fast to "time" it —start building and learning now

The Three Anchors: Prioritize top-tier **Talent**, rigorous **Guardrails**, and unshakable **Foundations**.



Start now!



Thank you!

Nuwan Bandara, Head of Technology – Fintech, AWS | nuwan@amazon.com

