Deloitte.

Unifying operations and analytics

Convergence architecture and the path to closed-loop intelligence

About the Deloitte AI Institute™

The Deloitte AI Institute helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With".

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem comprised of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries, to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, deliver impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in; whether you're a board member or a C-Suite leader driving strategy for your organization, or a hands on data scientist, bringing an AI strategy to life, the Deloitte AI institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet ups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

Introduction

Artificial intelligence (AI) and agentic systems are moving from experiment to enterprise transformation. Yet, even as important investments are being made in AI and its enabling infrastructure, many organizations still struggle with siloed systems. **Often, the insights AI reveals are trapped behind dashboards, rather than driving business outcomes.**

The heart of the challenge is that operational and analytical systems have traditionally been used and managed separately. With the introduction of AI agents, these ecosystems are converging, and a new, dynamic approach to maximizing the business insights from operational and analytical systems is convergence architecture. It is a unifying design that integrates data, AI, and operations into an adaptive intelligence fabric across the technology landscape.

While the path to convergence architecture is iterative and characterized by horizons of maturity and opportunity, in its boldest form, it enables enterprises to evolve to operate as thinking, learning, and acting organizations. The expectations for AI and agentic impact are being tempered by the realities of adoption and the transformation of work. **Convergence architecture offers a new method of turning technology investment into actionable insight.**

Facing a gap between knowing and acting

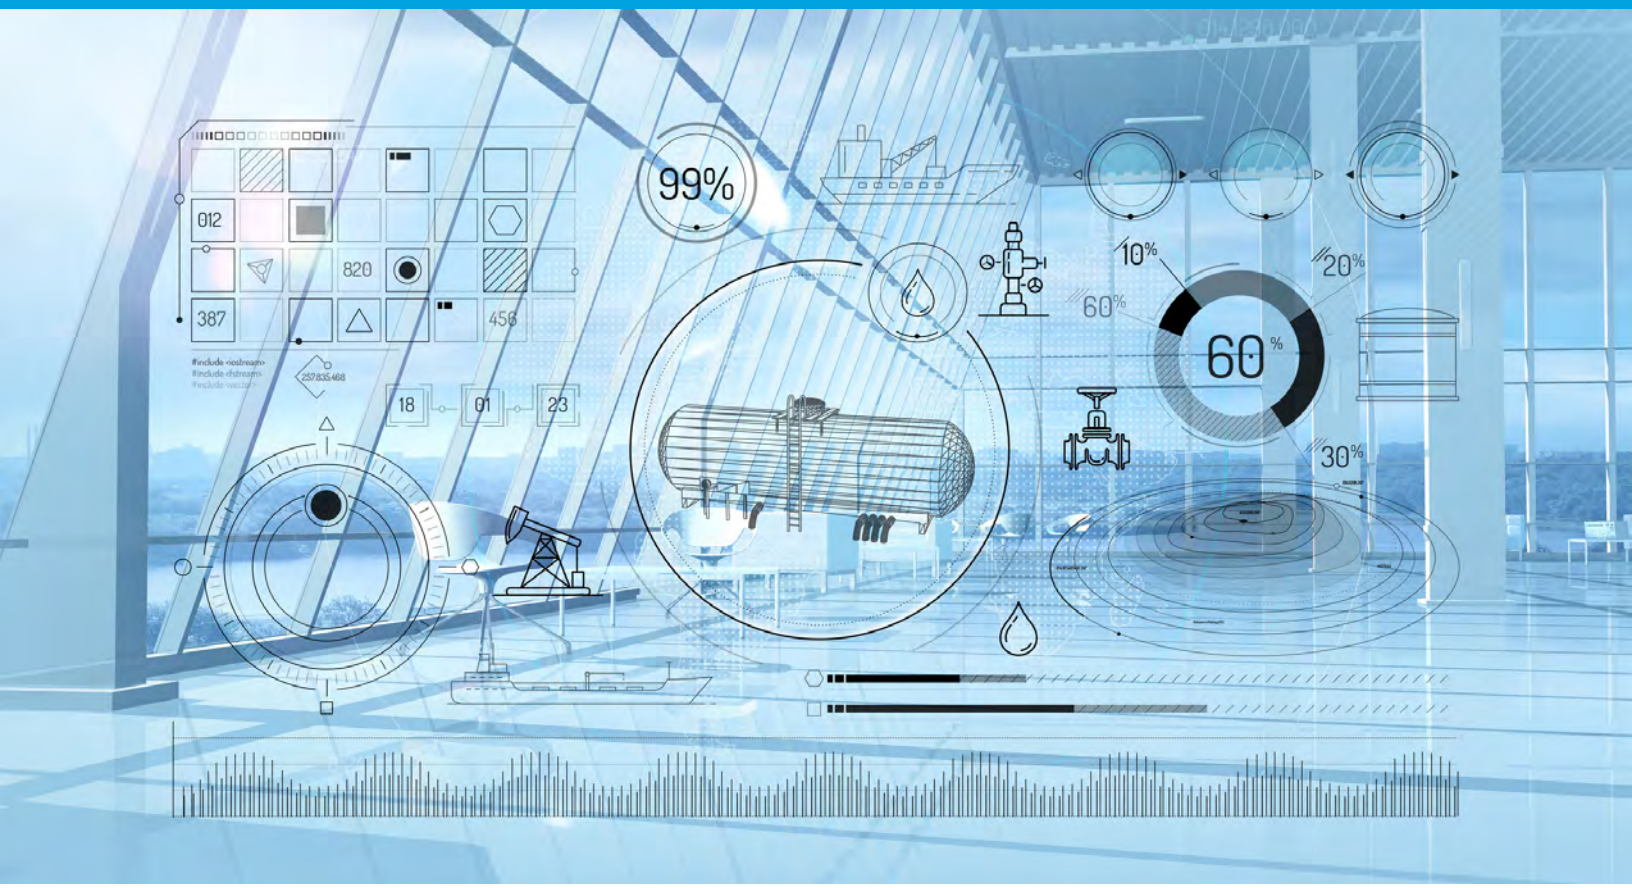

Many enterprises today are facing a widening paradox. Despite significant investments in data platforms, data warehouses, and analytics tooling, many insights remain locked in dashboards, reports, or presentation decks. These insights may inform quarterly reviews but rarely influence operational decisions in real time. The result is a **persistent gap between what enterprises know and what they do.**

The challenge owes to **architectures and barriers that frustrate adoption and create risk** in AI and agentic systems. The causes include:

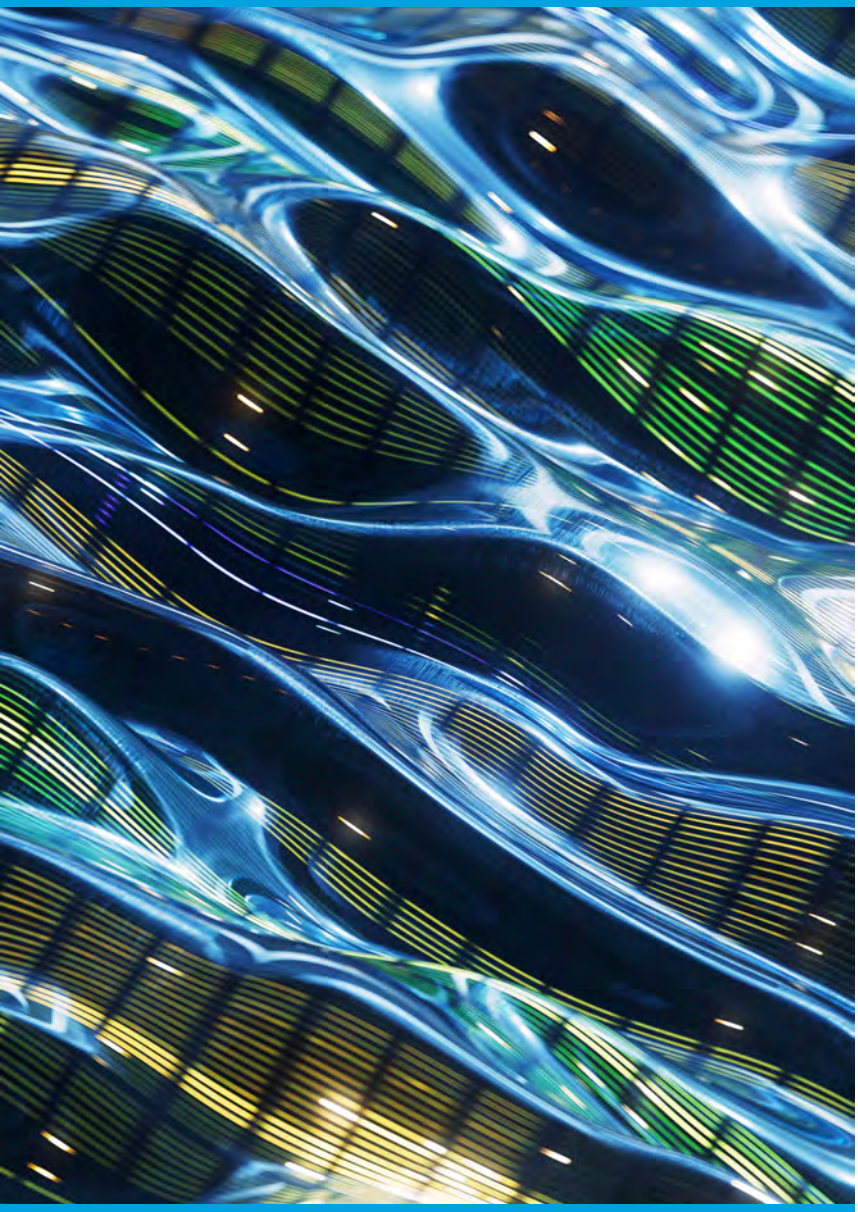

Ultimately, existing architectures are insufficient to close the loop between insight and action. Enterprises need a new method that is **faster, contextual, adaptive, and human-centered.**

An operational–analytical divide
Operational systems (enterprise resource planning [ERP], customer relationship management [CRM], core banking) and analytical systems (warehouses, marts) were not designed to work seamlessly together. They use different schemas, ontologies, and governance rules. Bridging them has historically meant brittle ETL (extract, transform, load) and ELT (extract, load, transform) pipelines, which succeed in batch but fail when low-latency, bi-directional flows are required.

Latency and rigidity
In many enterprises, decision cycles are still measured in weeks or months. By the time data is extracted, transformed, and analyzed, the business context has already shifted. Competitors and customers expect much faster, near-real-time responsiveness.

Data quality and trust gaps
Poor lineage, duplicate records, inconsistent governance, and opaque processes erode trust in data. The lack of confidence undermines adoption, slows innovation, and increases risk.

AI adoption risks
Even as AI models and agentic systems mature, enterprises hesitate to deploy them deeply into operations. Output errors (e.g., hallucination, misinterpretation) can become costly and disruptive, and a lack of guardrails, auditability, and explainability makes enterprises cautious.

Brownfield reality
Most organizations run decades-old systems that are too costly or risky to replace. New AI-native systems must coexist with these legacy platforms, and while this creates complexity, it also presents an opportunity for layered innovation.

Organizational barriers
Beyond technology, questions of ownership, skills, and change management loom large. Who owns the connective "fabric" between data, AI, and operations? How do IT, data, and business teams realign to support agent-augmented workflows? Without clarity, convergence efforts stall.

Experience gap
Many enterprises underestimate the importance of the experience layer. Even the most advanced AI and data fabrics add little value if employees, customers, or partners cannot interact with them intuitively. Convergence must translate into experiences that are contextual, actionable, and trustworthy.

An agile approach to architecture

Convergence architecture is a **unifying design with principles and patterns that connect enterprise operational systems, analytical platforms, and AI** through an agile intelligence fabric.

At its center is an AI core—a system that understands business context, data, and technology, enabling real-time insight-to-action loops, safe agentic orchestration, and experience-first outcomes.

Convergence architecture is not a monolithic platform, nor a rip-and-replace exercise. Instead, it is a **layered, adaptive approach.** Convergence architecture integrates legacy and modern operational and analytical systems and augments them with an AI core that can think, learn, and act. It grounds intelligence in a cognitive knowledge fabric that manages semantics, lineage, and data flows, and orchestrates execution through agents and services. Value is realized through intuitive human and customer experiences.

Key characteristics of convergence architecture include:

Closed-loop intelligence
Insights flowing seamlessly into actions and actions continuously improving insights.

Hybrid integration
Accommodates brownfield realities with overlays, redesigns, and Software-as-a-Service (SaaS).

Dual-mode data flows
Real-time streaming for responsiveness, batch for scale, and compliance.

Experience-first design
Convergence becomes visible through better decisions and journeys.

Adaptive governance
Privacy, lineage, and explainability are woven into the core fabric.

The layered design of convergence architecture

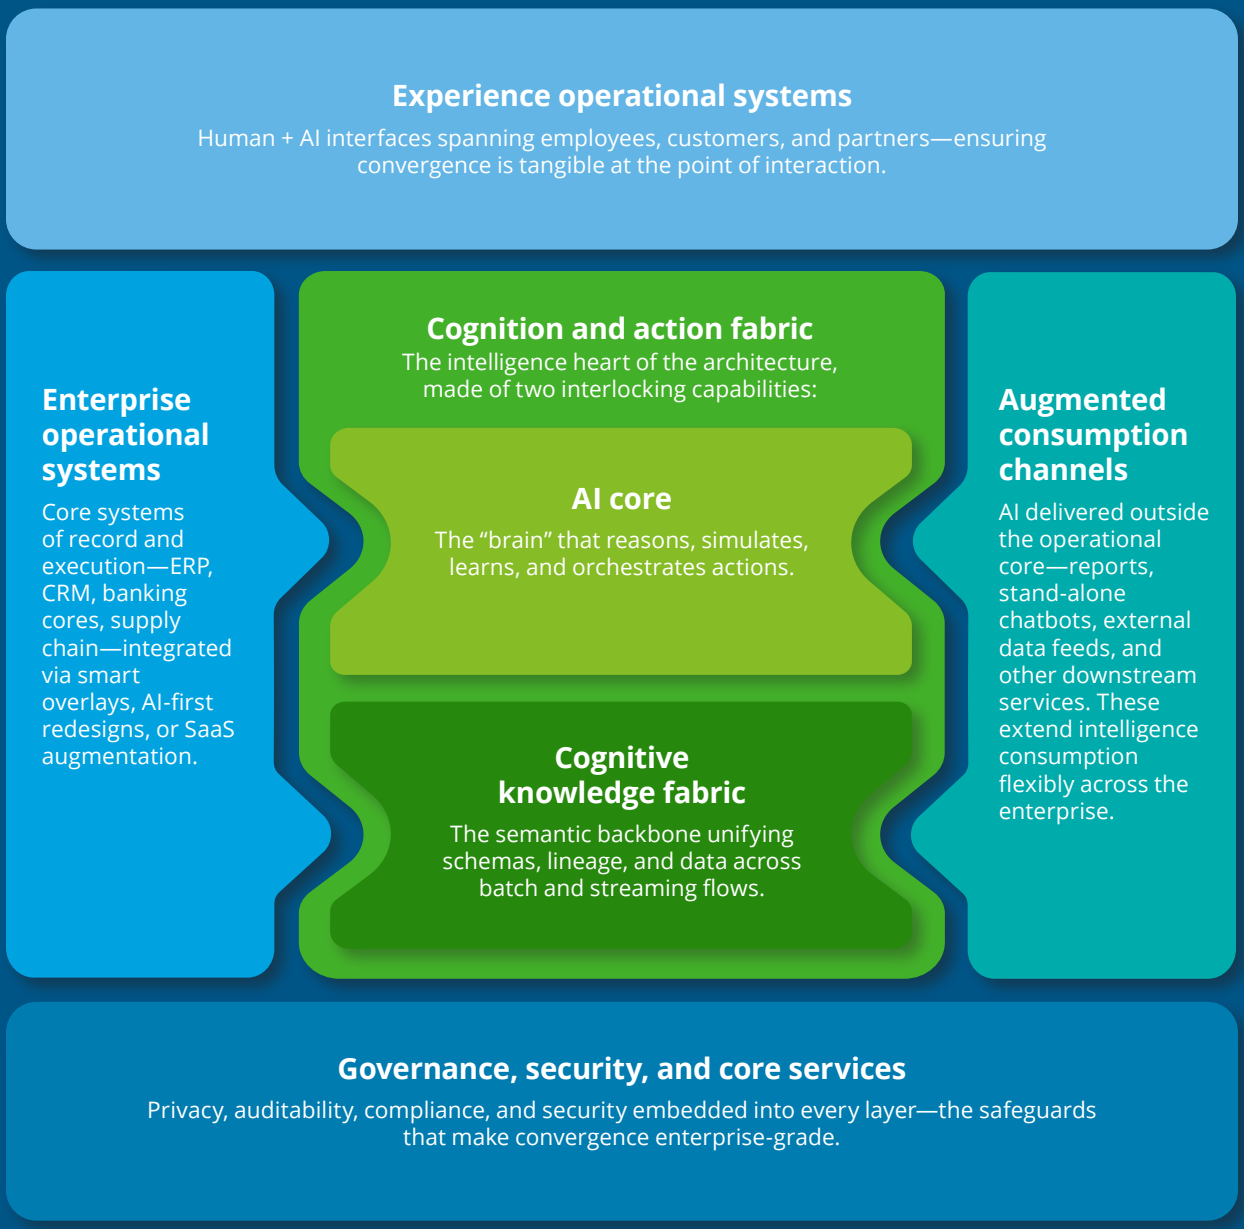

The pillars of convergence architecture

The value and opportunity in convergence architecture depends on **seven interconnected pillars.**

Trustworthy data and knowledge foundation

Traditional metadata, lineage, and governance frameworks designed for centralized warehouses are insufficient, and in their place, enterprises require a cognitive knowledge fabric. This is a living, semantic layer that maps relationships across data domains, processes, and systems. It supports real-time responsiveness (e.g., operational actions) and batch processing (e.g., regulatory reporting large-scale analysis), as well as continuous monitoring of quality, lineage, privacy, and compliance, with automated corrections and alerts. This cognitive knowledge fabric becomes the navigational chart for humans and AI agents, enabling them to act on consistent and reliable information.

AI core for thinking, learning, and acting

At the center of convergence sits an AI core, essentially the brain of the architecture. Unlike traditional engines, this is not a single model but instead a system of capabilities. These include: business and technology awareness, understanding both enterprise processes and the data and technology that enable them; advanced simulations to test decisions before execution, offering scenarios and trade-offs; real-time feedback from monitoring ongoing actions, providing alerts or adjustments to agents and humans; and continuous learning to improve over time. The AI core elevates the enterprise from being system-driven to being intelligence-driven, where every process benefits from reasoning, foresight, and adaptive learning.

Agentic orchestration with contextual grounding

Agents are the operational arms of convergence, turning insights into action. However, their effectiveness depends on deep contextual grounding. Agents must operate within the semantic structure of the cognitive knowledge fabric, and guardrails need to be built in via approval workflows, audit logs, and fallback mechanisms. Human-AI collaboration should be the norm, ensuring that sensitive or high-stakes actions include human oversight. In this way, agents can act as trustworthy collaborators.

Composable AI and model governance

The AI layer is composable, allowing enterprises to combine foundation models for general reasoning, domain-specific models tuned to industry workflows, and third-party SaaS AI services for specialized capabilities. This composability sits within a strong governance framework for monitoring model drift, auditing decisions, measuring fairness, and ensuring explainability.

Event-driven, real-time integration

Convergence architecture requires an event-driven, real-time fabric that enables the continuous flow of intelligence across operational systems, analytical systems, the AI core, and the cognitive knowledge fabric. It permits streaming pipelines for low-latency ingestion and transformation, feedback loops where operational outcomes immediately update analytical systems, and hybrid support for both real-time and batch modes, depending on context.

Human, organizational, and experience alignment

Convergence impacts not only systems but also people and experiences. Capturing insight and value with convergence architecture requires clear ownership of the fabric, often requiring new cross-functional structures for IT, data, and business. Leveraging an experience-first design ensures insights and actions are delivered through intuitive, contextual interfaces. To adapt to new ways of working, enterprises will likely need to invest in upskilling across roles (from data teams to domain experts), as well as drive change management to build human trust in agent-augmented workflows.

Responsible and secure by design

Finally, convergence must be secure, auditable, and compliant. This requires privacy-by-design, with fine-grained access control, end-to-end lineage and traceability for all agentic actions, and security to mitigate adversarial exploitation and prevent agents from being leveraged as attack vectors.

Taken together, these pillars transform enterprise architecture from fragmented and reactive **to unified, intelligent, and proactive.**

Three horizons of adoption

Most enterprises have had core operational systems in place for years, or even decades. Systems for banking, ERP, CRM, supply chain, and claims platforms cannot simply be unplugged in pursuit of architecture suited to the AI era. Convergence architecture accepts this brownfield reality, with **three horizons of adoption that allow enterprises to iteratively evolve and mature.**

1. Smart overlays

AI and agentic capabilities are layered on top of existing systems without modifying the underlying application. **These overlays act as copilots, assistants, or orchestration layers that sit between users and legacy platforms.** Examples include contact center agents that draw from multiple data sources to recommend next-best actions, loan operations copilots that generate data summaries for human officers, and enterprise search and retrieval across siloed datasets.

The potential benefits include fast time-to-value (weeks or months), limited disruption to core operational platforms, and the capacity to experiment and gradually adopt agentic workflows. There are limitations to this adoption pathway. Overlays can become brittle if systems change and constraints (e.g., data latency, a lack of application programming interfaces [APIs], and rigid workflows) in the underlying systems are inherited. It may also be more challenging to achieve deep automation, as agents primarily assist humans rather than act on their behalf. Today, many enterprise AI agents are overlays. They are the way to experiment without significant disruption, but while they are an important first step, they are not the destination.

2. AI-native redesigns

Critical applications or processes are rebuilt with AI-native cores. **These redesigns place the AI core and cognitive knowledge fabric at the center of the workflow,** enabling closed-loop intelligence and real-time feedback. For example, fraud detection engines that adapt and learn from data streams, or supply chain control towers simulating constraints or disruptions and autonomously adjusting planning.

This approach **enables thinking, learning, and acting systems.** It unlocks process redesign (not just process augmentation) by embedding AI in the core of the process. That said, it requires significant investment in technology, governance, and change management. Adoption may be slower, and there may be greater demands for aligning skills, accountability, and the operating model.

3. Third-party SaaS integrations

Instead of building everything in-house, some enterprises **integrate specialized AI services from external providers.** These services bring niche capabilities or external data that could be difficult or costly to replicate internally, such as credit scoring APIs integrated into loan origination, customer sentiment analysis SaaS embedded into CRM, or external data enrichment and fraud feeds powering risk operations.

SaaS integrations offer **quick access to cutting-edge AI, as well as faster scaling in areas outside the enterprise's core differentiation.** They also provide industry benchmarks and network effects (e.g., fraud detection across multiple banks). Yet, capabilities are dependent on vendor roadmaps and offerings; SaaS can create integration overhead; and data governance, privacy, and security may become more complex.

While each adoption pathway brings benefits and limitations, in practice, enterprises will **likely pursue these three approaches in parallel,** with overlays for speed and experimentation, AI-native redesigns for transformative impact in select domains, and SaaS integrations for specialized needs.

Industry opportunities with convergence architecture

Banking

Convergence architecture enables banks to integrate legacy and modern systems, apply AI-driven intelligence at the core, and create adaptive, secure, and customer-centric services.

Real-time risk assessment in lending

Banks need to rapidly evaluate risks for lending, drawing on both internal and external datasets. Convergence architecture can help expedite lending decisions, strengthen regulatory alignment, and surface insights for risk evaluation.

- AI-powered overlays sit atop existing loan origination and risk systems, pulling together datasets, customer profiles, and regulatory updates to give loan officers real-time risk insights within their familiar interface.
- AI-native redesigns fundamentally reimagine risk assessment workflows, with a new AI core ingesting structured/unstructured data, continuously learning from new events, and autonomously surfacing risk scenarios and mitigation strategies.
- SaaS integrations are seamless as the architecture connects to third-party APIs and regulatory reporting SaaS platforms.

Cross-channel customer experience orchestration

Banks seek to deliver seamless, personalized experiences across mobile, web, branch, and call center channels. Convergence architecture helps give customers consistent, engaging experiences, driving retention and higher net promoter scores.

- Unified AI-driven customer engagement overlays provide agents and digital channels with a single, contextual view of each customer, orchestrating next-best actions and personalized responses.
- End-to-end customer journey management is rebuilt as an AI-native service, with real-time journey mapping, intent detection, and adaptive workflow engines that optimize every interaction across all platforms.
- Connects with external customer feedback, sentiment analysis, and personalization SaaS tools to ingest customer behavior signals, benchmark journeys, and auto-adjust engagement strategies.

Instant onboarding and life cycle management

Banks seek to accelerate customer onboarding and seamlessly manage ongoing know-your-customer (KYC) and anti-money laundering (AML) requirements. Convergence architecture can reduce onboarding time and cost, improve compliance, and enable a low-friction user journey that can evolve as regulations or business models change.

- A smart onboarding overlay guides customers and bankers, orchestrating API calls to internal/external verification systems and surfacing resolution workflows for exceptions, while also remaining layered atop existing CRM and KYC systems.
- Onboarding and identity validation processes are rebuilt as event-driven AI pipelines that handle customer input, document verification, decisioning, and auto-triage.
- Flexible plug-ins connect to SaaS KYC/AML providers, fraud detection services, and other authentication tools to enable global scalability and instant policy updates.

Insurance

Insurers operate with complex policy systems and document-heavy processes. Claims handling can be slow and costly, and underwriting relies on static actuarial models that cannot adapt to real-time signals. Convergence architecture shifts insurers from manual, reactive processes to adaptive, AI-augmented workflows, compressing cycle times, lowering loss ratios, and enhancing customer satisfaction.

Real-time claims triage and settlement

Insurers want faster, fairer, and more efficient claims handling to boost customer trust and reduce costs. Convergence architecture can enable a dramatic reduction in claims cycle time, lower fraud exposure, and improved customer satisfaction.

- An AI-enabled claims overlay extracts and summarizes key details from submitted forms, images, and adjuster notes, providing an integrated view and next-step prompts directly within adjuster workbenches.
- The claims workflow is rebuilt as an AI-powered analytics engine that automatically classifies claim complexity, flags anomalies/fraud, and routes simple claims for instant payout, with continuous self-improvement from new data and outcomes.
- APIs connect to external SaaS services such as vehicle/health data providers, fraud analytics engines, and digital payment processors, enabling faster fact validation and automated settlements.

Personalized policy pricing and underwriting

Insurance carriers strive to offer more accurate, behavior-based policy pricing for customers, especially in auto and wellness lines. Convergence architecture helps drive highly competitive and fairer pricing, accelerated policy issuance, and more precise risk management.

- Contextual policy pricing overlays integrate real-time telematics and customer data into underwriter workflows, assisting in personalized risk assessment without overhauling core systems.
- Underwriting engines are reimagined as continuously learning AI models that score risk based on multisource behavioral, environmental, and claims data, rapidly adjusting rates and recommending offers tailored to individual risk profiles.
- Plug-ins bring in external SaaS data (e.g., driving scores, wearable health analytics, and industry risk pools) to enrich internal models and enable dynamic, market-responsive pricing.

Convergence architecture shifts insurers from manual, reactive processes to adaptive, AI-augmented workflows, compressing cycle times, lowering loss ratios, and enhancing customer satisfaction.

Retail and consumer

Retailers juggle volatile demand, global supply chains, and heightened expectations for personalization. Legacy ERP and merchandising systems often cannot handle real-time changes, leading to stockouts, markdowns, or generic customer experiences. Convergence architecture turns retail from batch-driven merchandising into adaptive commerce, in which supply chains flex dynamically and every customer journey feels personal.

Unified omnichannel inventory visibility and fulfillment

Customers expect items to be available and fulfillable from anywhere (i.e., store, warehouse, or third-party partners). Convergence architecture helps reduce stock-outs and overstocks, enable higher fulfillment speed, and improve customer satisfaction with flexible delivery options.

- Inventory management overlays offer real-time stock visibility and fulfillment options (e.g., ship, reserve, pickup) to customers and associates via store apps, kiosks, and web dashboards, all without replacing underlying systems.
- Global inventory and order routing solutions are rebuilt as AI-driven orchestration engines, predicting demand, optimizing facilities and routes, and auto-adjusting safety stock levels in real time.
- The architecture connects to SaaS platforms for distributed order management, last-mile logistics, and third-party marketplace fulfillment to expand sourcing and improve customer service levels.

Streamlined supplier collaboration and dynamic merchandising

Retailers need to react to supply disruptions and changing consumer demand with agility. Convergence architecture enables faster response to market changes, optimized stock, and stronger supplier relationships backed by data.

- Supplier portals and merchandising dashboards overlay AI insights for forecasting, allocation, and ordering, enabling merchants and planners to act quickly from familiar user interfaces.
- The merchandising lifecycle is reinvented as an AI-based process, automatically responding to trends, detecting anomalies, and rebalancing assortment or pricing—all with minimal manual intervention.
- The architecture connects to SaaS partners for real-time market trends, demand sensing, and supplier business-to-business networks, allowing seamless information flow and collaboration across the ecosystem.

Convergence architecture turns retail from batch-driven merchandising into adaptive commerce, in which supply chains flex dynamically and every customer journey feels personal.

Assessing the adoption timeline

Convergence architecture adoption unfolds across horizons, as enterprises build confidence, modernize infrastructure, and rewire organizational processes. **The key is to view convergence architecture as a strategic roadmap, rather than a single implementation.** Importantly, adoption is not linear, wherein one approach is abandoned and replaced with another.

While each enterprise contends with its own strategic goals and AI ambitions, following the roadmap of convergence architecture will likely play out along the following general timeline.

Near-term (1–2 years)

Conditioning the organization for convergence

The near-term focus is likely to be on **quick wins** in customer service, loan operations, claims processing, and merchandising. **Adopting smart overlays for speed and experimentation** helps build confidence in AI deployments by ensuring contextual grounding, human-in-the-loop oversight, and auditability. Even as deep automation remains limited and convergence is visible primarily at the experience layer, smart overlay adoption will help employees begin experiencing the benefits of agent-augmented workflows, and leadership will gain the confidence to scale.

The experience gained in overlay-driven convergence will come alongside efforts around modernizing metadata, quality, and governance frameworks for real-time contexts. This conditions the organization for more ambitious change and establishes the foundation to begin experimenting with more advanced integration in workflow redesign and SaaS adoption.

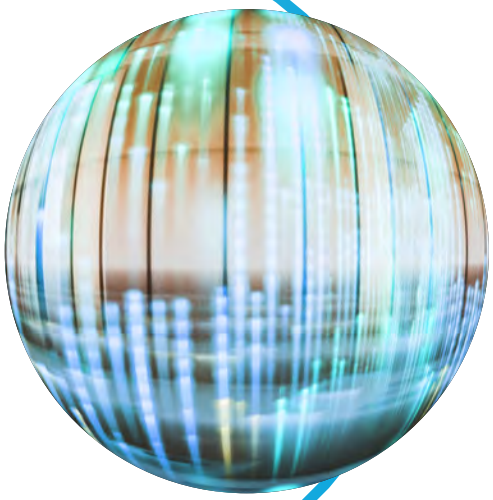

Medium-term (2–5 years)

Streaming and redesign-driven convergence

With early wins and growing confidence, enterprises will begin to tackle deeper integration by **modernizing pipelines and selectively redesigning critical applications.** The focus will likely be on transitioning from batch ETL/ELT to real-time streaming fabrics with closed feedback loops. Building cognitive knowledge fabrics will unify ontologies, semantics, and lineage across domains, and new model governance frameworks will support continuous monitoring, drift detection, and explainability.

Convergence architecture shifts from AI assisting humans to humans and AI co-orchestrating workflows. Here, select domains achieve closed-loop intelligence where insights automatically shape operations. This is also the phase where organizational redesign begins, focusing on ownership of the fabric, new operating models, and new skill profiles.

Long-term (5+ years)

AI-native enterprise convergence

With maturity and experience, enterprise architecture itself becomes intelligence-driven by design. **The AI core and cognitive knowledge fabric become central to operations, and agents act safely and securely across systems with minimal human intervention.** Transformation will focus on embedding AI cores as the strategic decision layer across the enterprise, with governed, enterprise-wide adoption of agentic orchestration grounded in contextual knowledge. Dual-mode operation will permit real-time responsiveness and large-scale analytics, and third-party ecosystems and offerings (e.g., SaaS, data feeds, partner systems) will be woven into the fabric.

This is the ambitious vision with convergence architecture, where enterprises operate as thinking, learning, and acting systems. Insights seamlessly flow into action, and outcomes continuously feed back into intelligence. With this, customers, employees, and partners experience organizations as adaptive, context-aware, and proactive. **Even as this is a goal, however, smart overlays, redesign-driven convergence, and SaaS will coexist, depending on enterprise goals and the need to maximize value, manage costs, and mitigate risk.**

Charting a path forward

To move from vision to execution, leaders should consider **six guiding principles** for convergence architecture adoption.

1

Think fabric, not islands

Avoid fragmented, siloed AI deployments, and design for a connective intelligence fabric that links operational systems, analytical platforms, and AI.

2

Start with trust

Data quality, governance, and lineage are the foundation. Without trustworthy data and clear guardrails, the enterprise will be challenged to scale AI responsibly.

3

Adopt portfolio thinking

Balance overlays (quick wins), AI-native redesigns (transformational bets), and SaaS integrations (specialized accelerators).

Prioritize experience

The value of convergence architecture is realized at the experience layer, where employees, customers, and partners interact. Design for human-AI collaboration from the start.

4

Governance by design

AEmbed privacy, security, auditability, and explainability into each layer, while maintaining that governance should enable innovation, rather than impede it.

5

Organize for change

Clarify ownership of the fabric, invest in skills, and adapt operating models to agent-augmented workflows.

6

Guided by these principles, enterprises can take a phased approach to adoption.

In phase one (first 12 months), establish the foundations for convergence architecture. Stand up a cognitive knowledge fabric by unifying metadata, lineage, and key ontologies, and deploy smart overlays in high-value use cases (customer service, claims, loan ops) to build trust. This is also the period where governance, security, and monitoring frameworks are defined.

In phase two (one to three years), focus on scaling real-time convergence. Shift from batch ETL/ELT to streaming-first data pipelines with feedback loops, and selectively re-architect AI-native applications in critical domains (e.g., fraud, claims, supply chain). During this time, expand governance to cover model drift monitoring, explainability, and fairness checks.

In phase three (three to five years or longer), embed AI core and enterprise-wide convergence. The enterprise is prepared to operationalize the AI core as the central intelligence layer across business units. Agentic orchestration is expanded across front, middle, and back offices, and external SaaS and ecosystems are integrated into the cognition and action fabric. At this point, the enterprise should be redesigning organizational roles and experiences for human-AI collaboration at scale.

The journey begins now

The opportunities with AI and agents are found through convergence architecture. In the future, enterprises will not simply run systems. They will operate as thinking, learning, and acting organizations. Recognizing this future state and the incremental steps required to reach it, enterprise leaders should begin adoption now. This means building the fabric with overlays and governance foundations, investing selectively in redesigns where the payoff is transformational, and embracing SaaS where it accelerates capability. Above all, **design for closed-loop intelligence, with insights flowing seamlessly into actions and actions continuously improving insights.**

Authors

Prakul Sharma
AI & Insights Practice Leader
Deloitte Consulting LLP
praksharma@deloitte.com

Brijraj Limbad
Senior Manager – AI & Data
Deloitte Consulting LLP
blimbad@deloitte.com

Naveen Balawat
Specialist Leader, AI & Data – Financial Service
Deloitte Consulting LLP
nbalawat@deloitte.com

Deloitte.

About Deloitte

Deloitte refers to one or more Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms. Please see www.deloitte.com/us/about for a detailed description of the legal structure of Deloitte LLP and its subsidiaries. Certain services may not be available to attract clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.