cognizant

Confronting the AI velocity gap: A new architecture for enterprise operations

Contents

Executive summary

Agentic AI has moved from experimentation to boardroom priority, promising a future where autonomous digital workers execute complex business processes with unprecedented speed and scale. This vision is driving massive investment and reshaping technology roadmaps across the industry. However, for business and technology leaders accountable for outcomes, the challenge is no longer understanding what agentic AI can do—but determining how to deploy it reliably, at scale and with measurable business impact.

Despite rapid advances, many organizations face a widening gap between AI investment and realized value. This disconnect, described by Cognizant CEO Ravi Kumar S as the AI velocity gap, reflects the absence of a practical bridge between compelling demonstrations and enterprise-grade execution.

This white paper, “Confronting the AI velocity gap: A new architecture for enterprise operations”, the first in a three-part series, offers a pragmatic blueprint for navigating the initial, and most critical phase of this transformation. It deconstructs the core components of agentic AI and confronts the challenges of implementation head-on. The paper’s central argument is that success requires a fundamental architectural shift away from traditional human-led augmentation.

We examine two competing implementation models and advocate for a differentiated approach anchored in straight-through processing (STP) from day one. The STP-first design philosophy redefines the role of human involvement, moving it from an active human-in-the-loop (HITL) participation to a strategic human-over-the-loop (HOTL) governance. Finally, the paper introduces a four-stage maturity model and the modern operations architecture required to guide organizations from their current state of automation toward truly autonomous and resilient enterprise processes.

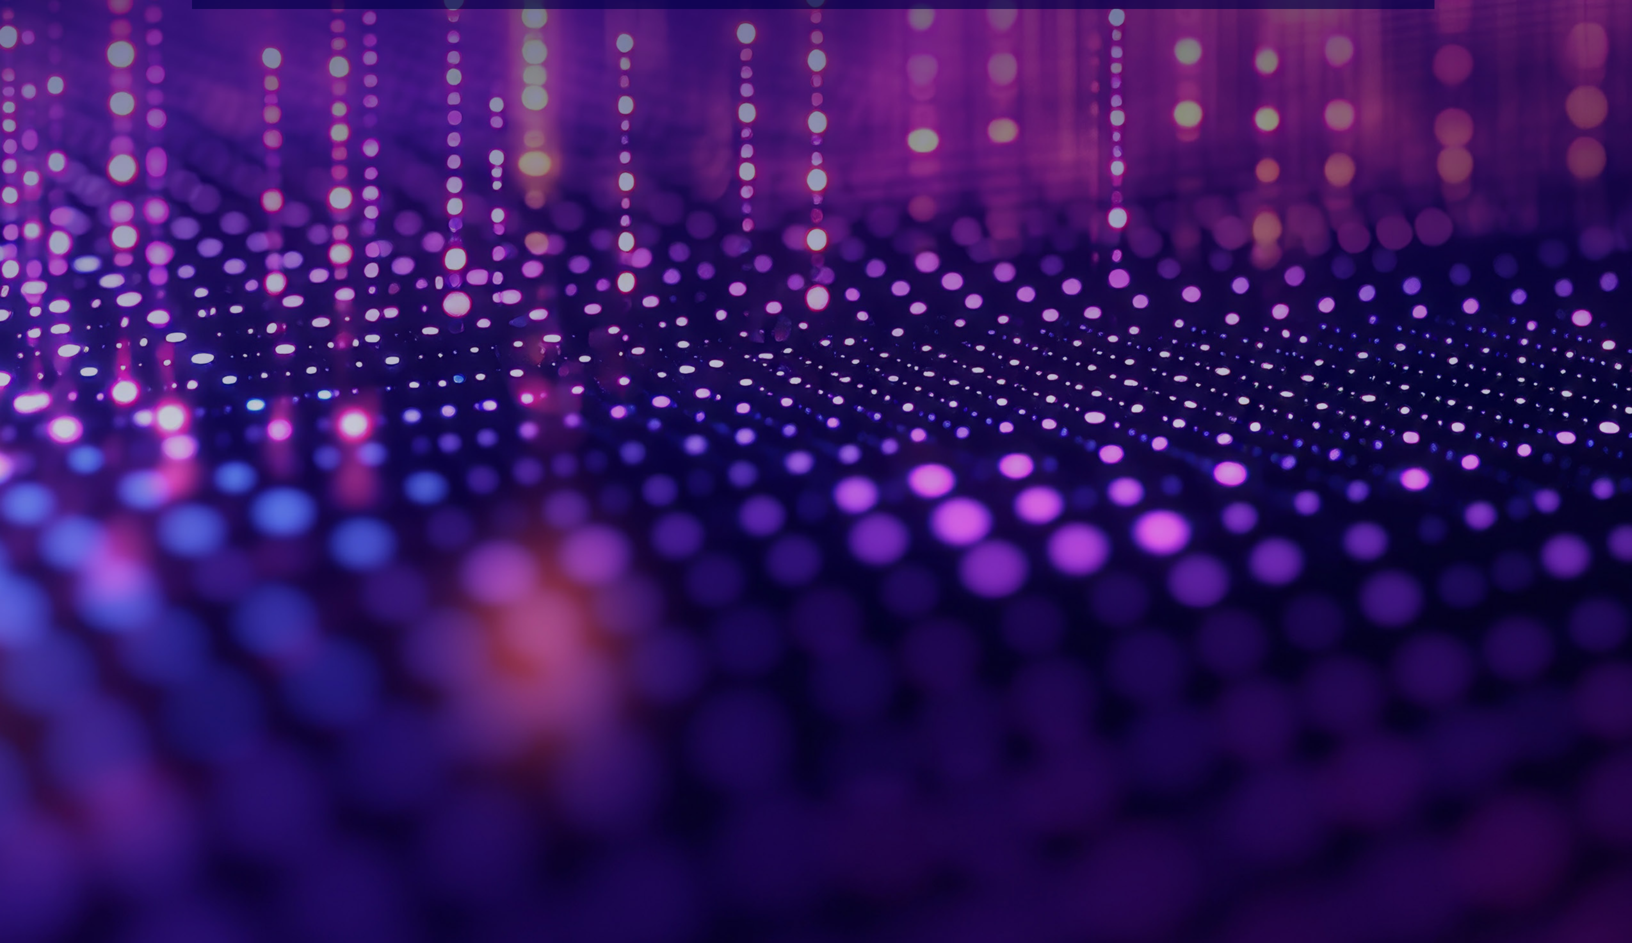

Beyond the hype: Confronting the AI velocity gap

Every executive is grappling with the promise of agentic AI. We are told that autonomous agents will not just augment human capabilities, but will take over entire business functions, driving unprecedented efficiency. However, amid this optimism, a critical conversation is avoided. We're excited about the potential benefits, but that excitement can pull us away from asking the questions that matter most. So, we have to ask: What are the true, measurable benefits we can expect? What is the real cost of building, deploying and maintaining these agentic solutions? And how will the maturity of these systems actually progress over time?

This white paper is the first in a three-part series designed to answer these questions. The overarching goal of the series is to provide a comprehensive playbook for building an agentic enterprise—an organization architected around autonomous execution, with humans providing oversight, judgment and strategic direction. This journey begins here, with clarity on the fundamentals, the architectural trade-offs and what it truly takes to cross the AI velocity gap.

While the principles discussed apply broadly, we will draw heavily on examples from the business process outsourcing (BPO) industry. No other sector provides a better real-world laboratory for agentic AI. The scale of its operations, the diversity of its processes and the relentless pressure for efficiency make it the ideal setting for testing these new models. The lessons learned in this domain offer a powerful blueprint for any enterprise seeking to transform its core operations.

Agentic AI: The catalyst for transformation

What makes this new era of automation possible is the emergence of agentic AI. An AI agent is not just a model; it is a system that can perceive its environment, reason about a goal, create a plan and execute that plan by orchestrating a set of tools. This ability to act autonomously is what separates agentic AI from the predictive and generative AI of the past.

The core components of an AI agent

While the specific implementations vary, most AI agents share a common set of core components that work together to enable autonomous operation.

- Perception: The ability to ingest information from the outside world, including text, images, audio and structured data.
- Reasoning engine: Typically powered by a large language model (LLM) and/or large reasoning model (LRM), this is the component that understands the user's intent, formulates a plan to achieve it and adapts that plan as new information becomes available. Planning capabilities transform AI from a reactive system into a proactive agent.
- Tool use: The ability to use a set of external tools to execute the plan. These tools can range from simple API calls to complex software applications. This might include searching databases, calling web services, executing code, updating systems or sending communications. Tool use dramatically expands what AI can accomplish beyond just generating text.
- Memory: The ability to store and retrieve information from past interactions, allowing the agent to learn and improve over time. This includes both short-term memory (the context of the current conversation) and long-term memory (a knowledge base of facts and procedures). Memory enables agents to improve over time and handle complex, multistep processes.
- Self-monitoring and metacognition: Advanced agentic systems can evaluate their own performance and confidence, recognize when they lack needed information or capabilities, seek help or additional information when needed, and detect and recover from errors. This metacognitive capability is crucial for reliable autonomous operations.
- Context understanding: Agents need to interpret and apply domain-specific knowledge, understand business rules and policies, recognize relevant contextual factors and align decisions with organizational values and objectives. This is where context engineering becomes critical.
- Learning mechanisms: Agents should be able to learn from feedback and outcomes, identify patterns in successful and unsuccessful actions, refine decision-making models based on experience and expand capabilities over time. Learning enables continuous improvement

Transformative applications of agentic AI in BPO

Agentic AI enables several transformative capabilities in BPO contexts that go far beyond simple task automation as below:

- Autonomous process execution: AI agents can independently execute end-to-end business processes with minimal human intervention. Consider a claims processing agent that receives a claim submission, extracts relevant information from documents, verifies coverage based on policy terms, assesses the claim amount based on documentation and historical data, detects potential fraud indicators, makes an approval decision for routine cases, generates and sends decision communications and updates all the relevant systems. The agent handles the entire process autonomously, escalating to humans only when it encounters situations beyond its capabilities.
- Dynamic process optimization: AI agents can continuously analyze process performance and make real-time adjustments to improve outcomes. They identify bottlenecks and inefficiencies, test alternative approaches, implement improvements that prove effective, and adapt to changing conditions without requiring human intervention. This enables processes to improve continuously rather than waiting for periodic optimization initiatives.
- Intelligent exception handling: Rather than simply escalating all exceptions to humans, AI agents can identify and resolve many exceptions autonomously. They recognize when a situation deviates from the norm, determine whether they can handle the exception based on their knowledge and capabilities, attempt resolution using available tools and information and escalate to humans only when necessary, providing comprehensive context to facilitate human decision-making.
- Proactive issue resolution: AI agents can anticipate potential problems and take preventive actions. They monitor conditions that typically lead to problems, predict likely issues based on patterns, take proactive steps to prevent problems from occurring and alert humans about risks that require strategic attention. This shift from reactive to proactive operation creates significant value.

Challenges and limitations

Despite the immense potential, agentic AI is not a magic fix. Deploying reliable, enterprise-grade agents requires overcoming significant technical and organizational challenges:

- Reliability and hallucination: LLMs can hallucinate, which means invent facts or make logical errors. Ensuring that agents act reliably and factually requires grounding them in authoritative knowledge sources and implementing robust validation mechanisms.
- Explainability and trust: The black box nature of some AI models can make it difficult to understand why an agent made a particular decision. Building trust with users, auditors and regulators requires systems that can explain their reasoning in a clear and transparent way.
- Bias and fairness: AI systems can perpetuate or amplify biases present in training data. Ensuring that agentic AI systems make fair decisions across different demographic groups and situations requires ongoing attention and mitigation efforts.
- Security and adversarial robustness: As AI agents gain more autonomy and access to systems, they become potential targets for manipulation or attack. Ensuring security and robustness against adversarial inputs is critical.
- Alignment with human values: Ensuring that autonomous AI agents make decisions aligned with human values and organizational objectives is a fundamental challenge. As agents become more capable, the stakes of misalignment increase.

Successfully navigating these challenges requires a pragmatic and engineering-led approach. It requires a deep understanding of the existing technology landscape and a willingness to use a combination of modern APIs with established and more pragmatic automation techniques to get the job done.

Organizations should start with clear and bounded use cases (where the value is clear and risks are manageable), invest in robust context engineering, implement strong governance and oversight mechanisms, maintain appropriate human involvement for high-stakes decisions, continuously monitor and improve agent performance, and prepare for the organizational changes that agentic AI will require.

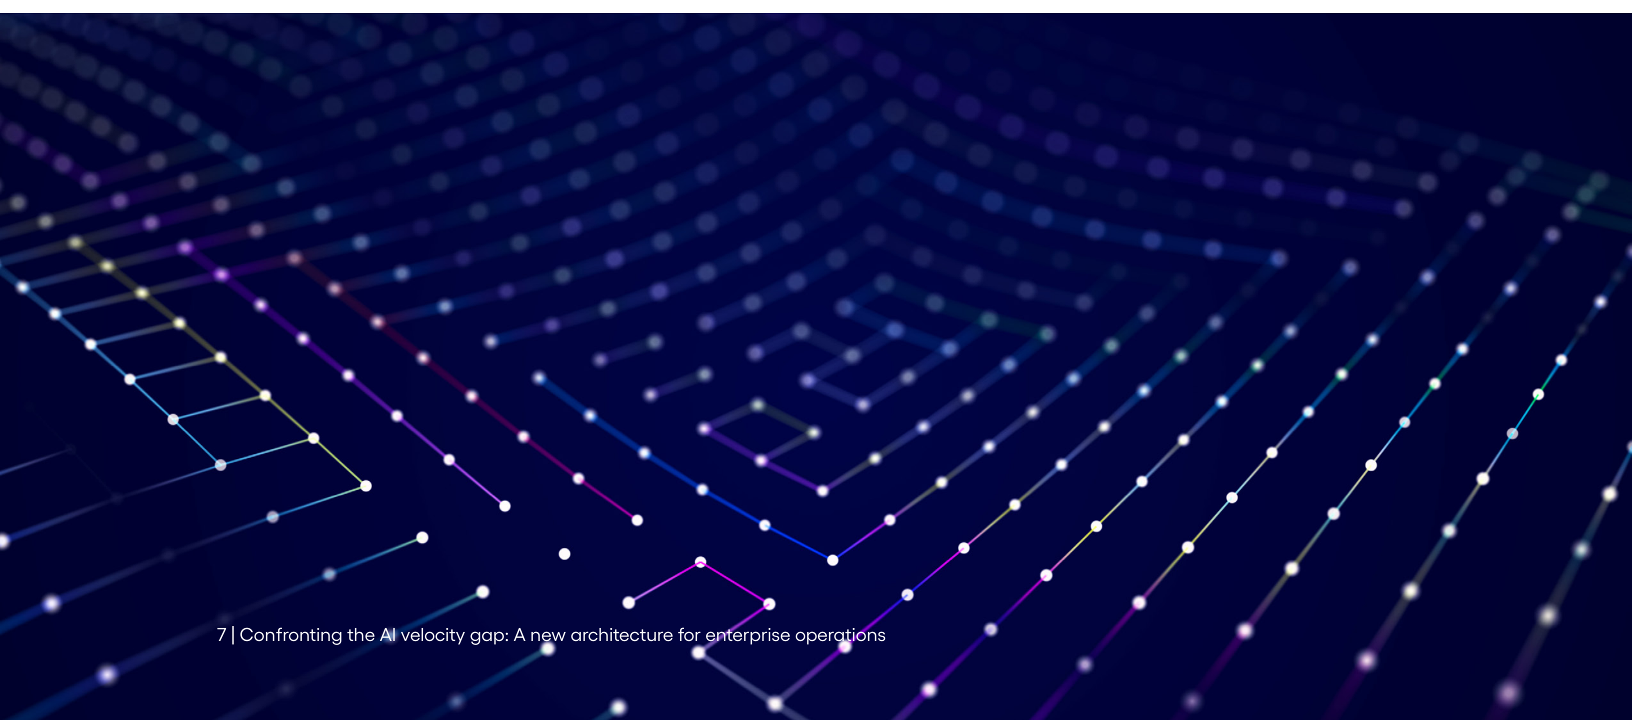

At the implementation crossroads: Traditional versus STP-first agentic AI

Enterprises today stand at a critical crossroads. The traditional approach to AI design and implementation is characterized by siloed projects, incremental automation and a focus on augmenting existing human-led processes. It is proving insufficient to unlock the full potential of agentic AI. So, a fundamentally different, more holistic approach is required; one that is architected from the ground up for autonomous operation, with straight-through processing (STP) as its design principle. This is the differentiated STP-first agentic AI design.

The traditional agentic AI design approach: Human-led, AI-enabled

The traditional model of agentic AI design is best described as human-led, AI-enabled. It focuses on augmenting the capabilities of human workers, who remain the primary actors at every critical step of the process. The AI acts as an assistant, providing recommendations and data analysis, but the final decision-making authority and quality control rest firmly in human hands.

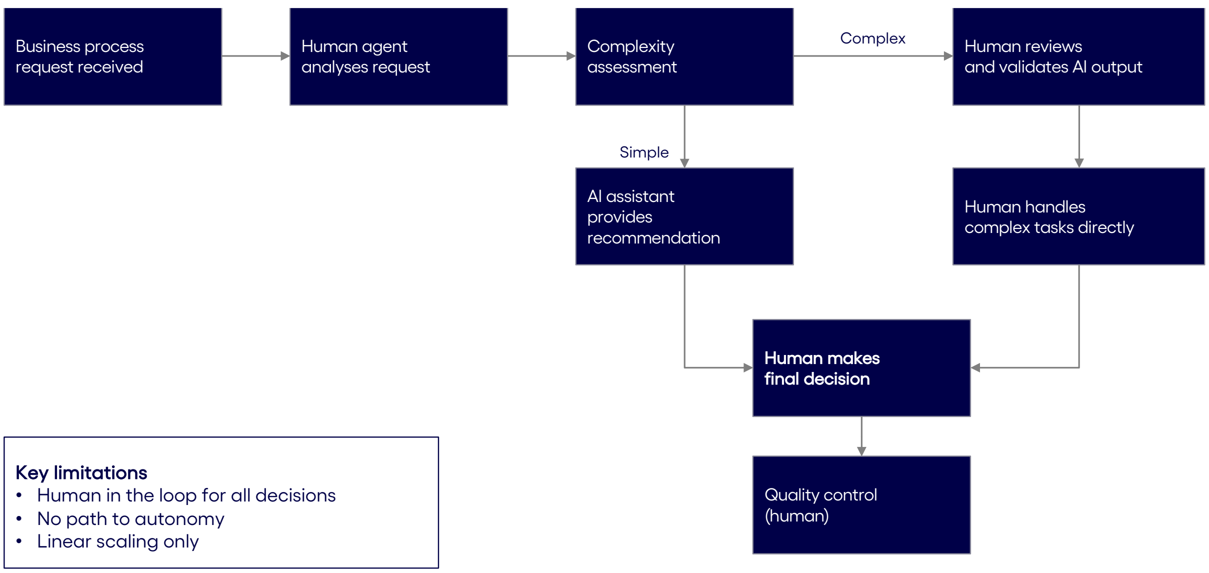

This approach, illustrated in the diagram above, typically follows this workflow:

- A business process request is received and analyzed by a human agent (sometimes by an AI agent).
- A complexity assessment is carried out by a human. Simple, routine tasks are routed to an AI assistant for a recommendation, but this recommendation must be reviewed and validated by a human.
- Complex or strategic tasks are handled directly by a human, who may be supported by AI-powered data analysis tools.
- A human makes the final decision in most cases.
- Following the process execution, a human performs quality control. If the quality check fails, it triggers human intervention and correction loop.

While this model can deliver incremental efficiency gains, it suffers from a fundamental flaw that it is not designed for autonomy. The constant need for human review, validation and decision-making creates a hard ceiling on scalability and speed. It treats AI as a bolt-on to existing processes rather than as a transformative core, leading to a fragmented technology landscape and a failure to capture the exponential value of true automation.

The proposed differentiated approach: AI-led, human-enabled

The proposed STP-first architecture represents a paradigm shift to an AI-led, human-enabled model. It is not about augmenting existing processes. It is about reimagining them for a world where AI is the primary actor. This approach is built on a foundation of process reimagination, modular architecture and progressive autonomy.

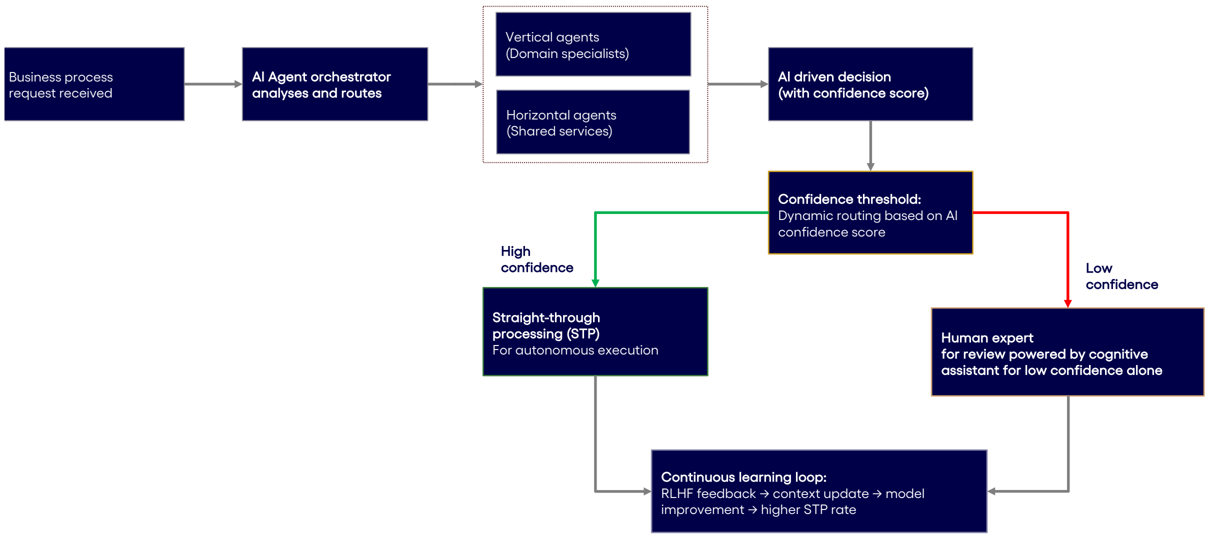

This architecture, designed with STP as its north star, operates very differently:

- A business process request is received by an AI agent orchestrator.
- The orchestrator performs context analysis and selects the appropriate vertical (specialist) and horizontal (generalist) agents for the task.
- The selected agents collaborate to perform specialized and cross-functional tasks, leading to an AI-driven decision.
- The core of the model is the confidence threshold. If the AI agent's confidence in its decision is high, the process is executed autonomously, achieving STP.
- The task is escalated to a human expert for review only when the agent's confidence is low. The human's decision is then fed back into a continuous learning loop, which updates the context and improves the agent over time.

This model is designed to overcome the limitations of the traditional approach. By designing STP-first architecture, it creates a clear path to autonomy. The modular agent network allows for scalability and flexibility. And the continuous learning loop ensures the system becomes more intelligent and efficient over time, progressively reducing the need for human intervention.

The following table summarizes the fundamental differences between the two approaches:

Dimension	Traditional AI implementation	Differentiated STP-First Agentic AI Architecture
Design philosophy	Augment existing human-led processes	Reimagine processes for AI-led orchestration
Architecture	Siloed point solutions, human-centric	Modular, composable agent network, STP-first
Human role	Primary actor, reviewer and decision-maker	Governor, exception handler and AI trainer
Scalability	Linear (more people for more human-in-the-loop volume)	Nonlinear (10 times volume without 10 times headcount)
Value capture	Incremental, task-level efficiency	Holistic, end-to-end process automation
Path to STP	No clear path; designed for human review	STP is the default path for high-confidence tasks
Learning	Ad hoc or nonexistent	Systematic and continuous via RLHF feedback loop
Explainability	Dependent on point solution	Full traceability with reasoning and confidence scores

The exception processing myth and STP-first design

For decades, the goal of process automation has been to handle common cases and route the exceptions to a human. This has led to a mindset where automation is seen as a tool for reducing the volume of work, but not for fundamentally changing its nature. Agentic AI requires a new north star and that is what STP is. The goal is not to reduce exceptions, but to design a system where the default path is fully autonomous, and human intervention is a carefully managed, value-adding activity, not a routine part of the flow.

Defining the human role: Human-in-the-loop versus human-over-the-loop

As we shift from a human-led to an AI-led model, the nature of human oversight evolves. The term human-in-the-loop is often used loosely, but for a successful agentic transformation, a more precise vocabulary is required. The distinction between two key modes of human involvement is critical:

Term	Role and responsibility	Analogy	Typical stage
Human-in-the-loop (HITL)	Active participant, but a human is required in the process. AI cannot proceed past a certain point without explicit human review, correction or approval for an individual transaction.	A copilot who must confirm the pilot's actions before the plane can change course.	Early stages (one and two)
Human-over-the-Loop (HOTL)	Supervisory governor. The human monitors the overall system, reviews aggregated outcomes and manages exceptions. AI operates autonomously on individual transactions, and the human intervenes only when predefined thresholds are breached or for governance purposes.	An air traffic controller who oversees many planes, trusting the pilots to fly their routes, but ready to intervene to prevent collisions or manage emergencies.	Mature stages (three and four)

This evolution from HITL to HOTL is the central pillar of the workforce transformation in an agentic enterprise. It is the journey from being a processor of tasks to a governor of outcomes.

The progressive autonomy model

This does not mean aiming for 100% STP from day one. That is both unrealistic and unwise. Instead, the journey to full autonomy should be a phased progression, with each stage building on the capabilities of the last.

Stage 1: Cognitive assistance (0%–20% STP). In the initial stage, AI acts as a cognitive assistant. It assists a human by automating data gathering, summarization and initial analysis. But every transaction is still reviewed and approved by a human. High HITL oversight is the norm The focus is on building trust. The human is not just processing an exception; they are providing cognitive assistance to AI, and their feedback is captured to train the model.

Stage 2: Supervised autonomy (20%–60% STP). As AI gains confidence, it is allowed to process the simplest, lowest-risk transactions autonomously. AI-driven decisions are made with human validation. A significant portion of transactions are still routed to humans, but AI is now the primary actor for a growing subset of the workload. This stage represents a transition, with some processes under HITL and others beginning to operate under HOTL.

Stage 3: Governed autonomy (60%–90% STP). Autonomous agents handle majority of transactions. Human interventions are reserved for exceptions only. The system has demonstrated its reliability and the focus shifts to expanding the scope of autonomous processing. This is a primarily HOTL environment.

Stage 4: Near-full autonomy (90%+ STP). In the most mature stage, AI handles a vast majority of transactions autonomously. Human oversight is primarily for governance purposes. Humans can move into roles such as strategic oversight and governance, managing the system, handling truly novel exceptions and driving continuous improvement. This is a strategic HOTL role.

The role of a human evolves, moving from an assistant to a supervisor to a strategist. This philosophy of progressive autonomy is the core of the proposed approach. It is a journey, not a destination, that starts with a fundamentally different architectural mindset.

To execute this journey in a structured and measurable way, we use the four-stage maturity model. This model provides a clear roadmap for increasing automation and evolving human roles, which we will discuss further in this white paper.

So, it is important to realize that agent autonomy increases with maturity. While STP will be low and HITL volume high initially, this will progressively shift to high STP with supervisory HOTL as trust builds.

Exception processing: The critical distinction between the old and new models

It is crucial to understand the role of a human in this model. Even in the early stages, it is fundamentally different from traditional exception processing. In the old model, an exception is a failure of the system. The human's job is simply to fix the broken transaction. In the new model, a transaction routed to a human is a training opportunity. The human involved here is not just a processor; they are an AI trainer. Their role is to provide structured feedback, nuanced judgment and contextual understanding that AI needs to learn and improve. This feedback loop is the engine of the progressive autonomy model.

A common misconception is that AI-led, human-enabled model is simply advanced exception processing. It is considered that AI handles routine cases autonomously, while humans handle exceptions that do not fit standard patterns. This view is too simplistic. While exception-handling will certainly be part of the human role, the future BPO workforce will play much more sophisticated and strategic roles, including strategic oversight and direction, complex problem solving, relationship management, AI development and governance, innovation and transformation, and context engineering. These roles leverage capabilities that are distinctly human and likely to remain areas where humans excel compared to AI for the foreseeable future.

HITL escalation workflow

The HITL escalation workflow is not a simple throw-it-over-the-wall mechanism. It is a structured, intelligent process designed to maximize both the quality of human intervention and the learning value for the AI system. It is the primary pattern for stages one and two of the maturity model.

HITL escalation workflow

Primary pattern for Stages 1 and 2 of the Maturity Model

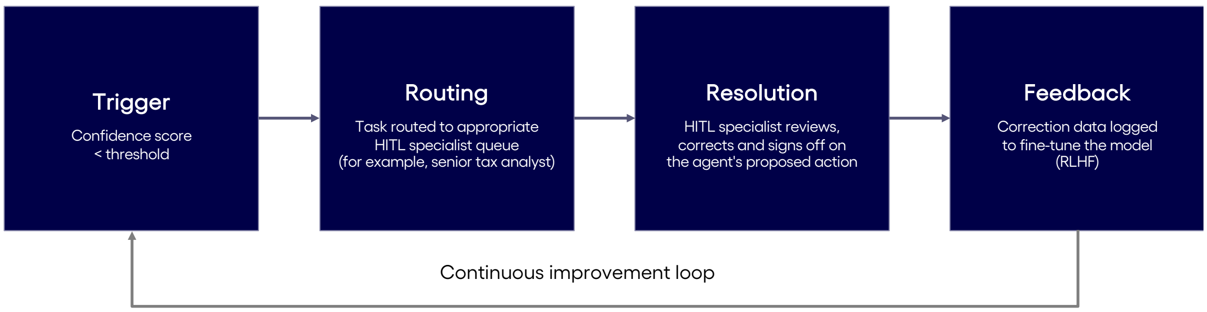

- **Trigger:** The process begins when an AI agent's confidence score for a particular decision fall below a predefined threshold. This threshold is not static. It is calibrated based on the risk profile of the transaction type and the maturity of the agent.
- **Routing:** The transaction is not sent to a generic queue. It is intelligently routed to a specific domain specialist queue. For example, a senior tax analyst for a complex tax-related exception, or a senior underwriter for a borderline mortgage application.
- **Resolution powered by cognitive assistant:** The human specialist does not work in isolation. They are assisted by a cognitive assistant—an AI tool that presents the agent's proposed action, the reasoning behind it, the confidence score and all the relevant contexts. This enables the human involved in it to review, correct and approve the agent's work far more efficiently than if they were starting from scratch.
- **Feedback:** The correction data from human intervention is systematically logged and used to fine-tune the AI model, through reinforcement learning from human feedback (RLHF). This closes the loop, ensuring that the system learns from every human intervention and progressively reduces the need for future interventions of the same type.

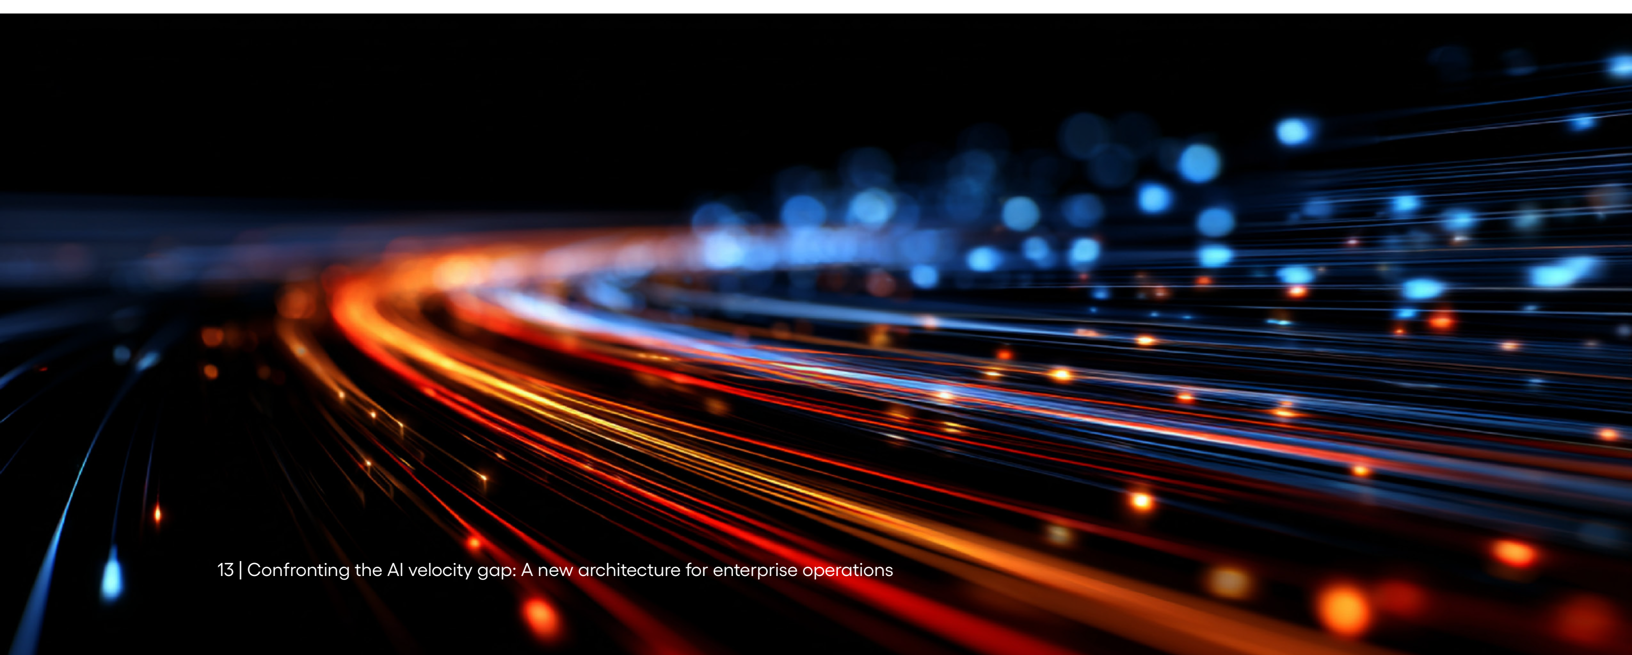

Enterprise agentification: The modern operations architecture

The transition from traditional automation to an agent-first enterprise is not merely a technology upgrade; it is a structural transformation of how operations are designed, governed and scaled. Here, we look at the holistic architecture for enterprise agentification, the blueprint for moving from hyperautomation to truly autonomous agents, which are structurally set up for maturity to come in.

The four-stage maturity model

The transition from traditional automation to an agent-first enterprise is not merely a technology upgrade; it is a structural transformation of how operations are designed, governed and scaled. To navigate this transformation, we use the four-stage maturity model. This model is the practical framework for implementing the philosophy of progressive autonomy.

The journey follows a four-stage model, where each stage represents a distinct operating posture with an increasing level of STP and a corresponding evolution in the human role.

Enterprise agentification: Modern operations architecture

From hyperautomation to autonomous agents: Structurally set up for maturity

Agent autonomy increases with maturity. While STP will be low and HITL volume high initially, this will progressively shift to high STP with supervisory HOTL as trust builds.

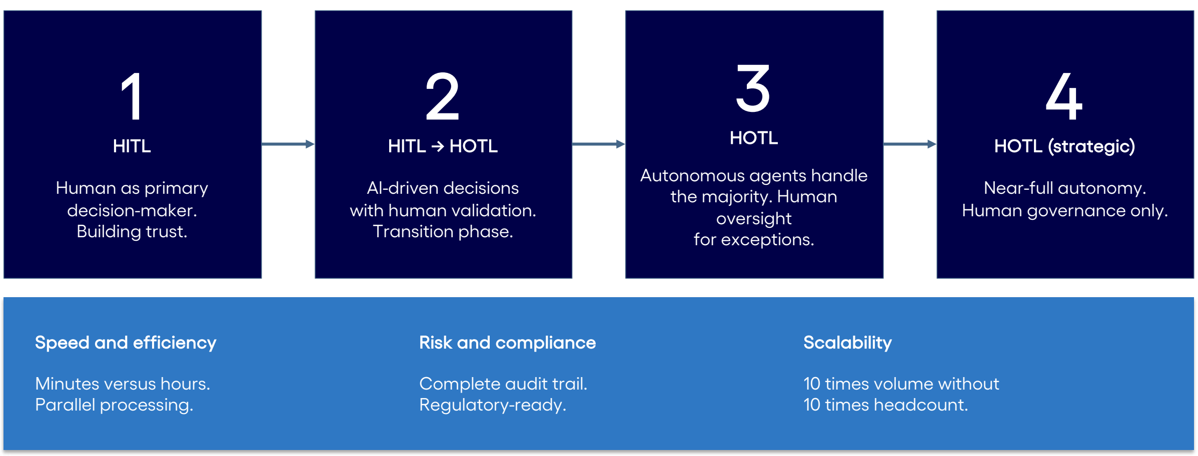

Maturity model

Stage	Operating posture	STP rate	Human role	Key focus
1	Cognitive assistance	0%–20%	HITL: Primary decision-maker	Building trust, capturing training data
2	Supervised autonomy	20%–60%	HITL > HOTL: Validator and reviewer	Expanding scope, calibrating confidence
3	Governed autonomy	60%–90%	HOTL: Exception handler and governor	Scaling, optimizing outcomes
4	Near-full autonomy	90%+	HOTL: Strategic oversight only	Governance, innovation, continuous improvement

The structural foundation of the maturity model

The architecture that enables this progressive maturity rests on three structural pillars.

- **Process reimagination:** This is not about automating existing processes. It is about digital-native redesign built for agent orchestration, not legacy automation. Every process is reexamined through the lens of: "If an AI agent were the primary actor, how would this process be designed?". This often leads to fundamentally different process flows that eliminate unnecessary handoffs, reduce decision points and optimize for STP.
- **Modular architecture:** The system is composed of plug-and-play agents that can be specialized as needs evolve. A modular architecture enables rapid composition of new capabilities, independent scaling of individual agents and graceful degradation when individual components encounter issues. This is the opposite of monolithic AI systems that are brittle and difficult to update.
- **Governed autonomy:** Governance controls are built into the architecture from day one, allowing autonomy to increase as trust metrics are met. These are not arbitrary thresholds. They are data-driven metrics based on the agent's demonstrated accuracy, consistency and reliability over time. The governance framework defines the conditions under which an agent can be promoted from one stage of the maturity model to the next.

Why this structure matters

The structural choices described above are not academic. They have direct, measurable impact on enterprise outcomes.

- **Orchestrated intelligence:** A master orchestrator coordinates specialized vertical and horizontal agents. Rather than operating in isolation, agents work as a coordinated network, with the orchestrator ensuring that the right agent handles the right task at the right time.
- **Confidence-based routing:** Dynamic flow to STP or human review (HITL or HOTL) is based on AI confidence thresholds. This ensures that the system is always operating at its optimal balance of speed and safety, and that this balance evolves as the system matures.
- **Full explainability:** Every decision is traceable with reasoning and confidence scores. This is not just a nice-to-have feature; it is a requirement for SOX compliance, regulatory audits and building trust that is essential for progressive autonomy.

The enterprise benefits of this architecture are substantial—speed and efficiency measured in minutes versus hours through parallel processing, risk and compliance supported by a complete audit trail that is regulatory-ready and scalability that enables 10 times volume without 10 times headcount.

Unlock agentic AI's full potential with robust operational architecture

Agentic AI represents a defining inflection point for modern operations, but realizing its value requires more than incremental automation. As this paper describes, the true opportunity lies in closing the AI velocity gap by rearchitecting operations around an AI-led, human-enabled model with STP as the north star. Organizations must adopt an STP-first, confidence-driven architecture, supported by progressive autonomy, modular agent design and strong governance that creates a clear, scalable path from experimentation to enterprise-grade autonomy. For business and technology leaders, the mandate is clear: success with agentic AI will be determined not by the sophistication of individual tools, but by the strength of the underlying operational architecture that allows intelligence, trust and value to scale together.

Related reading

- "From architecture to intelligence: Building the agentic technology stack"
- "The business of agentic AI: Costs, commercial models and the path forward"

Author

Anoop Nair
Senior Vice President, Global Head of FSI - IOA
anoop.nair@cognizant.com

Follow