MUKUND KALMANKER
Global Head, Data, Analytics and AI, Apexon

The next enterprise edge is responsible autonomy at scale

AI is clearly moving beyond assistance and into execution across a widening set of enterprise layers, from data engineering and IT operations to core business processes and customer-facing workflows. That transition is already visible in finance and operations, where AI systems do more than identify discrepancies. They can apply business rules, trigger reconciliations, and update downstream systems without manual intervention. In customer and domain workflows, agents are handling end-to-end journeys such as claims processing, service request resolution, and onboarding, making contextual decisions in real time and completing transactions across enterprise platforms.

The same shift is visible in pricing optimisation, supply chain management, and risk monitoring, where AI is beginning to execute decisions within pre-approved thresholds, accelerating response times and helping organisations react faster to change. In IT and platform operations, AI-driven systems are managing incidents, optimising workloads, and triggering remediation. Even internal functions such as HR and finance are moving from conventional application-led workflows towards agentic systems operating under strong governance.

The harder question is not whether AI can act, but whether enterprises are ready to trust it doing so. That requires robust trust and responsibility frameworks rooted in traceability, explainability, governance, AI-based validation, and human oversight in critical scenarios. AI should be empowered to execute, but only inside an environment that is governed, observable, and aligned with business intent.

Inside the Autonomous Enterprise

By Shrikanth G

For years, enterprise AI lived at the edge of work, surfacing insights, suggesting actions, and helping people move faster. That phase is now giving way to something more consequential. Across industries, AI is beginning to move inside the workflow itself, validating, triggering, routing, resolving, and, in some cases, completing work with minimal intervention. The shift is not universal, and it is far from risk-free. But it is real enough to demand a closer look.

In the pages that follow, Dataquest brings together voices from across the technology ecosystem to examine how this transition is unfolding on the ground. From banking, cybersecurity, and healthcare to software engineering, manufacturing, and customer operations, these perspectives explore where AI is already functioning as an execution layer, and what enterprises must solve before they can trust it at scale. Together, they reveal that the autonomous enterprise is not defined by AI alone, but by the systems, guardrails, and judgement that make AI-led execution possible.

PARITOSH ANAND
Chief AI and Digital Officer, Wadhwani Center for Government Digital Transformation

The autonomous enterprise raises the value of human judgement

The rise of the autonomous enterprise is not a sudden leap, but the latest stage in a long evolution of work. At the beginning sat pure human labour, where every judgement, every action, and every keystroke belonged to a person. That gave way to organised automation, fixed macros, rule-based bots, and robotic process automation that could execute repetitive tasks quickly, but only within narrow and fragile boundaries. The next meaningful shift was cognitive rather than operational, with AI surfacing insights, flagging anomalies, and offering recommendations while humans remained firmly in control of decisions and action.

Most enterprises today still sit in that intermediate phase, where AI can draft actions and suggest responses but a person must still approve them. The more consequential pivot comes when AI begins to act directly. Routine decisions move into autonomous handling, while human roles shift from operator to overseer, stepping in when confidence is low or the edge cases become too complex. That is the first point at which AI starts to inhabit what was once a human job.

The next stages are even more structural. AI co-workers begin adapting to live data, policy rules, and risk signals. Multiple agents coordinate work across silos, share goals, and eventually manage entire business processes end to end. In that future, the enterprise itself becomes legible to AI. What rises in value is not routine execution, but human judgement at a higher level: deciding, governing, imagining, and defining the ethical limits within which the system evolves.

JASPREET BINDRA
Co-Founder, AI & Beyond

Execution is spreading faster than enterprise trust

AI-driven execution is no longer confined to niche enterprise experiments. It is already visible across customer service, manufacturing, supply chain, and finance, in exactly the kinds of environments where action matters more than recommendation. Chatbots and virtual assistants are handling customer issues autonomously rather than merely routing them. In manufacturing, predictive maintenance and quality control systems are increasingly taking corrective action instead of just flagging conditions for review. Supply chains are being optimised in real time as AI adjusts logistics and inventory decisions, while finance continues to rely on automated trading and fraud systems that can execute transactions at speed.

Taken together, these shifts point to the same larger conclusion: AI is no longer just assisting enterprise decisions. It is beginning to execute them. That makes the governance question unavoidable. Once AI acts inside live workflows, every decision needs to be explainable, auditable, and compliant with the regulatory environment in which the business operates.

The central risk, therefore, is accountability and transparency. Enterprises need to know not just what the system did, but why it did it, how it arrived there, and whether the action can be defended. AI becomes an execution layer only when action can be trusted at the same level as any other critical business process. That means control, traceability, and responsible design have to scale alongside capability. Without them, autonomous execution remains a technical possibility, not a trusted operating model.

SINGARAVELU EKAMBARAM
SVP and Global Head of Delivery, Americas, Cognizant

Onboarding trust at scale: Over 30% of Cognizant's code now AI-generated

Across global enterprises, AI is no longer stopping at recommendation. It is already executing consequential work in production, though selectively and in the most value-dense parts of the enterprise. Customer operations, IT services, finance, and supply chains are leading that shift because these are domains where workflows are repeatable, policies are codified, and outcomes are measurable. Agents are closing service tickets, completing the last mile of KYC, processing claims, adjusting prices, triggering replenishment, resolving incidents, and initiating recovery actions end to end. In software engineering, the same trend is visible, with over 30% of Cognizant's code now AI-generated.

Cognizant leans in AI Builder principles: Do not just build models. Build accountable systems. The autonomous enterprise will be won not by more AI, but by engineered trust at scale. Ekambaram underscores four forces have unlocked this stage: higher model reliability, maturing interoperability protocols such as Model Context Protocol and agent-to-agent frameworks, enterprise-grade orchestration tooling, and, most importantly, treating institutional knowledge as an engineering problem. Context engineering, the work of supplying agents with the right operational logic, exceptions, and tribal knowledge, has become the real differentiator.

The biggest risk is runaway autonomy. The challenge is no longer whether AI can act, but whether its actions remain bounded by explicit decision rights, policies, confidence thresholds, identity controls, and fail-safes. Trust breaks when actions cannot be explained, audited, paused, or reversed. Cognizant's answer is governed autonomy by design, with telemetry, rollback, human escalation, and kill-switch controls built in from the start.

ALBERT NEL
Senior Vice President, Asia Pacific & Japan, Genesys

From scripted support to AI that resolves customer journeys

Customer experience is moving beyond AI that answers queries or automates simple tasks. The more important shift is towards AI that can orchestrate and execute end-to-end customer workflows across enterprise systems, safely, predictably, and at scale. That change is most visible in high-friction moments such as billing disputes, service changes, and delivery rescheduling, where AI agents can understand intent, coordinate actions across systems, trigger workflows, and confirm resolution within a single, connected experience.

This marks a clear departure from earlier automation, which depended on rigid, predefined paths and often faltered as journeys grew more complex or moved across systems. With agentic AI, and increasingly with large action models, enterprises are beginning to use AI not just to respond, but to reason, decide, and act in real time within clearly defined guardrails. The value, in turn, shifts from incremental efficiency gains to end-to-end outcome delivery: faster resolutions, greater consistency, and more personalised, empathetic experiences at scale. Human roles do not disappear in this model. They become more focused on moments that require judgement, creativity, and emotional intelligence.

The real challenge now is governance at the point of orchestration. Once AI begins coordinating actions across systems, channels, and touchpoints, the risk no longer sits in a single output. It sits in the entire experience. Trust will depend on governance by design, with transparency, auditability, accountability, and control embedded directly into the orchestration layer.

DAN SCHIAPPA
President, Technology and Services, Arctic Wolf

Security operations are becoming the proving ground for execution

Let us focus on cybersecurity. It is emerging as one of the first enterprise domains where artificial intelligence is clearly moving from support to execution. In security operations, speed, scale, and consistency matter more than human-only decision-making can realistically deliver, and the gap between signal and action is shrinking fast. At Arctic Wolf, that transition is most visible in the Aurora Agentic SOC, where AI agents do more than surface alerts or suggest remediation. They plan, manage, and execute security workflows end to end, taking on triage, investigation, prioritisation, and response orchestration that once depended on manual effort and scarce specialist expertise.

This is not simply a rules-based automation layer with a new label. The model is built around what the company calls a Swarm of Experts, where different agent types validate findings, initiate actions, and adapt workflows while humans remain involved where judgement, escalation, or higher-order interpretation is needed. That matters because cybersecurity is one area where defenders cannot afford to move slower than attackers, many of whom already operate at machine speed.

The real barrier, however, is trustworthy execution. Early AI systems have been undermined by hallucinations, brittle reasoning, drift, and weak accountability, making them risky in environments where errors carry operational consequences. Trust depends on continuous validation against outcomes, clear governance, disciplined human oversight, and proof that AI can outperform traditional workflows safely. Until that layer is engineered properly, AI will remain stuck at recommendation.

SID UGRANKAR
CEO, Qila.io

Autonomous agents could make software behave like an operating system

A more consequential shift in artificial intelligence is taking shape beyond chat interfaces and co-pilot tools. The next phase is being built around agents that act, not just models that respond. In that emerging architecture, a master intelligence sits at the top, orchestrating a network of specialised autonomous agents beneath it. These so-called drone agents are not simply better prompt responders. They interpret, plan, and execute specific tasks with precision, giving software a much greater degree of agency than traditional applications were designed to handle.

That is why the comparison with Software as a Service matters. SaaS transformed business by making tools such as customer relationship management, finance, and communications widely accessible without requiring enterprises to build their own infrastructure. But each SaaS product still delivered a defined service and a bounded outcome. Autonomous agents push that model further. Instead of navigating to a fixed destination, they decide the route themselves, adapting the product to the customer's intent rather than forcing the customer to adapt to the product's design.

In that model, AI stops being a feature embedded inside the product and becomes the operating system of the product itself. The implications are significant. Businesses can think beyond what service they offer and instead focus on what outcome the customer wants delivered. For users, the experience becomes simpler and more fluid. They state the intent, and the agent handles the work behind it.

SUNIL JOSE
President, Workday India

Workday shows where enterprise AI becomes operational muscle

The move from AI assistant to AI execution layer is becoming tangible in workflows where precision, volume, and business impact intersect. At Workday, that shift is most visible in payroll, contract review, and high-volume frontline hiring. In each of these areas, AI is no longer limited to generating insights and waiting for someone to act. It is scanning for anomalies, suggesting corrections, routing work, and in many cases carrying a workflow through to completion with oversight built into the system.

Payroll offers one of the clearest examples. AI agents can identify missing or inconsistent data before a run, propose fixes, and guide administrators through resolution, sharply reducing rework and employee escalations. In contract management, AI can review large volumes of documents, flag risky clauses, compare terms with approved templates, and prepare a first pass of redlines for legal teams to approve. In high-volume hiring environments such as retail or shared services hubs, AI can manage screening, scheduling, and candidate communication, freeing managers to focus on final decisions rather than coordination.

The value is direct and measurable: shorter cycle times, fewer errors, lower operational overhead, and better employee and candidate experiences. But that only scales when governance is treated as a design requirement. Enterprises need role-based permissions, policy controls, audit trails, and human checkpoints embedded from the start. AI becomes reliable only when it inherits and respects the same guardrails as any other critical system handling people or money.

VANYA SETH
Technology Head, India & Middle East, Thoughtworks

In life sciences, AI is moving from research aid to lab action

One of the clearest signs of AI crossing from support into execution is emerging in life sciences, where the work is data-intensive, highly regulated, and deeply dependent on context. Thoughtworks points to projects with global pharmaceutical companies such as Bayer and Pfizer, where AI research assistants are doing far more than summarising documents. They are pulling from decades of toxicology and study data, comparing similar trials, proposing next experiments, and drafting early versions of study reports and regulatory documents. That matters because it places AI directly inside the flow of scientific work rather than at the edge of it.

The same transition is visible in biologics. Organisations such as Gilead are using digital twins and AI to create more responsive lab environments, where systems continuously re-plan experiments, balance equipment utilisation, and flag process deviations in real time instead of waiting for humans to piece everything together at the end of a batch. The implication is clear: AI is no longer simply informing the scientist. It is helping carry the scientific process forward.

Managing autonomy is the key. Traditional governance before deployment is not enough when AI can change a price, approve a payment, or alter production capacity inside live workflows. Enterprises need runtime control, full audit trails, embedded policy, explainability, and human checkpoints inside the execution layer itself. Without that operational governance, AI may remain useful, but it will not become trusted.

ARUN RAJARAMAN
Senior Director, Software Engineering, Epsilon India

AI becomes useful when action stays observable and reversible

Artificial intelligence is moving most effectively into execution where enterprise decisions can be clearly framed, monitored, and reversed if needed. That is why some of the strongest use cases are appearing in automated campaign optimisation, dynamic pricing, offer orchestration, real-time fraud response, and IT operations. In these settings, AI is no longer sitting outside the workflow as a recommendation engine. It is embedded inside the process, taking action within defined guardrails and learning continuously from outcomes.

That shift matters because execution changes the stakes. Once AI starts acting directly, small errors can scale much faster than human teams are able to catch or correct them. The critical issue, then, is governance at scale. Enterprises need clarity on accountability, transparency into how decisions are made, and secure, resilient controls around data quality, bias, and drift. Without that, automation can amplify problems rather than reduce them.

Trust comes when AI systems remain observable, explainable, and designed with human override built in. That is what allows organisations to delegate execution without surrendering control. The goal is not to automate as much as possible for its own sake. It is to ensure that AI can act confidently in the areas where speed and responsiveness matter, while the enterprise retains a clear line of sight into what the system did, why it did it, and how the action can be corrected if needed.

MUKUND KALMANKER
Global Head, Data, Analytics and AI, Apexon

The next enterprise edge is responsible autonomy at scale

AI is clearly moving beyond assistance and into execution across a widening set of enterprise layers, from data engineering and IT operations to core business processes and customer-facing workflows. That transition is already visible in finance and operations, where AI systems do more than identify discrepancies. They can apply business rules, trigger reconciliations, and update downstream systems without manual intervention. In customer and domain workflows, agents are handling end-to-end journeys such as claims processing, service request resolution, and onboarding, making contextual decisions in real time and completing transactions across enterprise platforms.

The same shift is visible in pricing optimisation, supply chain management, and risk monitoring, where AI is beginning to execute decisions within pre-approved thresholds, accelerating response times and helping organisations react faster to change. In IT and platform operations, AI-driven systems are managing incidents, optimising workloads, and triggering remediation. Even internal functions such as HR and finance are moving from conventional application-led workflows towards agentic systems operating under strong governance.

The harder question is not whether AI can act, but whether enterprises are ready to trust it doing so. That requires robust trust and responsibility frameworks rooted in traceability, explainability, governance, AI-based validation, and human oversight in critical scenarios. AI should be empowered to execute, but only inside an environment that is governed, observable, and aligned with business intent.

MUTHUMARI S
Senior Director, AI, Brillio

How AI becomes the system, not support

AI has moved beyond the role of a support layer and into end-to-end execution across customer service, IT, HR, procurement, sales, marketing, and engineering. It's already visible in internal IT operations, where close to 100% of Level 1 work has been automated through AI. What once followed a familiar sequence of alert, human triage, ticket, and fix is now handled autonomously through detection, root cause analysis, remediation, and closure. That is the moment when AI stops helping the system and starts running it.

The same pattern can be seen in customer service, where deployments for a large telco and an insurance provider have moved beyond agent assistance into full query resolution. AI is handling routine requests, issuing refunds within guardrails, and managing workflows across systems, reducing wait times, lowering cost per interaction, and cutting agent workload so people can focus on complex or sensitive cases.

In HR operations, a global retailer is using AI to orchestrate onboarding, generate documents, trigger approvals, provision access, and maintain compliance. In procurement, AI agents are negotiating within thresholds, routing approvals, and updating downstream systems in real time. In software engineering, multi-agent pipelines are now planning, building, testing, and deploying with engineers stepping in mainly at checkpoints. The decisive issue, however, is decision provenance. Enterprise trust depends on the ability to reconstruct exactly why the system acted, what inputs shaped the decision, and where human oversight applied.

ARVIND ARAVAMUDHAN
Senior Director, Software Engineering – Platform, Harness

In software delivery, AI now acts inside production workflows

AI is no longer confined to generating suggestions for developers to review manually. It is triggering pipelines, generating and running tests, managing infrastructure changes, and taking part in incident response. In many organisations, this activity now sits directly inside the flow of software being built, shipped, and operated, which means AI is increasingly influencing reliability, security, and cost in live production environments rather than in controlled side experiments.

That makes governance far more urgent than it was in the co-pilot phase. The real problem is not the absence of guardrails in principle, but the fact that many of them remain fragmented, unevenly enforced, or too easy to bypass across disconnected tools. Once automation deepens, governance cannot remain an outer layer around the system. It has to be built into the system itself. Every action taken by AI inside the delivery chain needs to be governed through embedded policies, permissions, and real-time validation at the point of execution.

The trust challenge, then, is not whether AI can help teams move faster. It is whether execution can scale safely without sacrificing control. Platforms that standardise how AI operates across the lifecycle, and enforce authorisation, validation, and auditability through a single system of control, will define the difference between rapid automation and responsible automation. The future lies in enforcing decisions at the platform level, not merely surrounding them with guardrails.

ARCHANA VENUGOPAL
Senior Vice President and Chief Information Security Officer, NCDEX

For CISOs, the execution shift is already operational reality

From the CISO's chair, the AI execution layer is already taking shape in cybersecurity operations, where speed, accuracy, and response time directly influence enterprise risk.In security operations centres, AI-driven platforms are now autonomously triaging alerts, correlating signals, and launching containment actions such as isolating compromised endpoints, revoking access, or blocking malicious traffic, often within seconds. Similar patterns are emerging in identity and access management, where AI can enforce access decisions dynamically based on contextual risk, and in fraud systems that automatically stop suspicious transactions without waiting for manual approval.

The same trajectory extends into IT operations through AIOps, where AI is resolving incidents and optimising systems with minimal manual effort. What defines this stage is the emergence of closed-loop, policy-driven execution systems that do not merely interpret data but act on it within predefined governance boundaries. Most enterprises, however, still operate with a human-on-the-loop model, keeping oversight in place while gradually widening the scope of AI autonomy.

That makes the central risk painfully clear: loss of control and accountability. Once AI moves from advisory output to live action, the consequences become operational and systemic. Trust depends on control frameworks that make every AI-driven action traceable, auditable, and reversible, backed by strong policy guardrails, continuous monitoring, and fail-safe mechanisms. In cybersecurity, where a wrong move can affect business continuity, trust will come not from raw accuracy, but from the ability to govern, constrain, and intervene.

AMIT SHARMA
CTO, Ionic Wealth

Wealth management is where AI gets personal, proactive, and faster

Wealth management offers a strong case for what well-executed AI can become when it moves past summarisation and into action. With assets under management expected to expand sharply in the years ahead, the pressure is on firms to serve more clients with greater speed, context, and precision. AI is increasingly being used not simply to digest information, but to act as a high-fidelity execution partner capable of handling high-frequency queries, supporting client engagement, and delivering a far more responsive omnichannel experience.

At Ionic Wealth, it reflects in two distinct but related layers. At the business layer, proprietary algorithms embedded within the platform's database go well beyond templatised advice, enabling personalised financial planning based on each client's portfolio, goals, and risk appetite. At the engineering layer, tools such as Claude Code and Cursor are being used as autonomous agents that understand entire repositories, make multi-file changes, run tests, and iterate through tasks with minimal human input. That has helped drive close to 80% of new code generation, improve code coverage, and accelerate release quality. AI is also unlocking more proactive wealth workflows by detecting liquidity events, market shifts, and changes in client cash flow so that actions can be triggered within pre-approved mandates.

The deeper challenge is verifiability at execution time, what Sharma describes as the intent-to-output alignment gap. The real danger is not simply that AI may be wrong, but that enterprises may not realise it quickly enough. Without governance and verification built in from day one, AI execution becomes less a leap forward than a faster way to make mistakes at scale.

ABHINAV PARASHAR
Co-founder & CEO, Digio

AI inside compliance-heavy workflows where speed matters

Artificial intelligence is increasingly being embedded directly into enterprise workflows. That shift is can be seen in systems designed to improve user journeys, strengthen trust, raise productivity, and simplify due diligence. This is not AI sitting outside the process and offering guidance. It is AI interacting with live data, validating inputs, and enabling faster decisions across critical touchpoints inside regulated environments.

The most tangible examples come from compliance-heavy use cases. Digio has built AI-enabled offerings around intelligent character recognition, impersonation checks, and contact point verification, all of which reduce manual effort in verification and shorten turnaround times for regulated entities and their end customers. The same pattern is visible in compliance solutions covering AML, countering the financing of terrorism, politically exposed persons screening, and transaction monitoring. Here, AI is not just identifying potential risks. It is helping organisations curate rules and scenarios in line with customer risk appetite and operationalise those frameworks in real time.

That makes explainability and governance the essential trust layer. In regulated sectors, black-box execution is not acceptable. Enterprises need visibility into how decisions are made, traceability across actions, boundaries around where AI agents can operate, approval layers for higher-risk moves, and human oversight where necessary. Continuous training, testing, and validation are equally important as fraud patterns and regulatory frameworks evolve. In real-world enterprise settings, trust is built not on capability alone, but on control, transparency, and accountability. Only then does AI move from assistive tool to trusted execution layer.

AKSHAY GUPTA
Associate Director – GenAI, Navi

The future is controlled autonomy, not blanket automation

Two areas stand out. The first is customer experience, where AI-powered voice bots and chatbots are no longer just supporting human agents but resolving customer queries from start to finish. That includes understanding the issue, guiding the user through resolution, and completing the required action in high-volume support environments. The second is engineering productivity, where AI is moving beyond code suggestions and into the active execution of changes across repositories, taking over repetitive or surface-level work so engineers can focus on more complex problem-solving.

The broader point is that AI is no longer limited to generating insight. It is now embedded within workflows and taking action as part of the enterprise execution layer itself. Yet this is precisely where the biggest risk emerges: loss of controllability at scale. Once AI begins executing decisions, errors are no longer isolated. They can spread rapidly across systems and processes, turning small slips into larger operational failures.

That is why strong guardrails, structured quality control, traceability, observability, and fail-safe escalation mechanisms matter so much. Enterprises need systems that let AI act autonomously, but never without visibility, reversibility, and accountability. Navi also makes an important contrarian point. The future is unlikely to belong to fully autonomous enterprises everywhere. In practice, AI will remain selectively autonomous, handling high-volume, repeatable tasks while humans stay involved in ambiguous, high-stakes, or heavily regulated scenarios. The goal is controlled autonomy, not total automation.

Techkriti 2026: Forging Futures, Fueling Innovation

Techkriti 2026 wasn't just a fest. It was drones in the sky, robots in combat, generals talking strategy, AI talking medicine, and music shaking the nights. Four days where tech, war rooms, code, and concerts collided. A campus turned into a mini future.

BY DQ BUREAU

Techkriti, the annual flagship festival of IIT Kanpur, continues to stand tall as Asia's largest technical and entrepreneurial festival. Since its inception in 1995, it has embodied innovation and excellence. The 32nd edition, held from March 19–22, 2026, revolved around the theme NeoNousSingularita, bringing together technology, intellect, creativity, and electrifying experiences under one umbrella.

The name "Techkriti," a fusion of "tech" and "kriti" (creation), perfectly reflects its spirit: a celebration of ideas brought to life. Over four action-packed days, the campus transformed into a hub of discussions, competitions, and unforgettable performances.

A GRAND BEGINNING: DAY 0 HIGHLIGHTS

The festival commenced with an inauguration at the Main Auditorium, featuring the ceremonial lamp lighting and addresses by institute dignitaries. GVG Yugandhar, Director General of the Aircraft Accident Investigation Bureau, delivered a keynote on aviation safety and aerospace innovation, followed by Rahul Mazumdar's talk on making India a global drone hub. The day also featured events like Hovermania, Atlassian Career Café, and aerial showcases like Multirotor.

DAY 1: TECHNOLOGY, LEADERSHIP, AND ELECTRIFYING NIGHTS

Day 1 blended intellect and energy. Major General C. S. Mann spoke on indigenous defense technology and self-reliance. Global experts like Dr. Eric Grimson discussed medical technology and AI advancements. Group Captain Subhanshu Shukla inspired audiences with insights on India's return to space. Competitions like Bridge Design Challenge, Atlassian and Eightfold AI Hackathons, Manoeuvre, IARC, and CNC showcased technical brilliance. The day ended with a high-energy EDM Night by Nucleya.

DAY 2: DIVERSITY OF IDEAS AND UNSTOPPABLE ENERGY

Day 2 focused on inclusivity and innovation. The Women in Tech Panel featured Shailaja Donempudi, Dipti Saudagar, and Taranjeet Kaur, sharing leadership journeys. Mouna Neelkanta delivered the AWS keynote, followed by Air Marshal Balakrishnan Manikantan's defense insights. Competitions like Sky Sparks, Multirotor, and TGP continued, while Band Night by Omkar energized the evening.

DAY 3: STRATEGY, STRENGTH, AND A GRAND FINALE

The final day featured the impactful panel "Operation Sindoor: Tri-Service Operations in the 21st Century," with Lt. General D. P. Pandey, Air Marshal Philip Thomas, and Rear Admiral Ajit Thekkepat discussing modern warfare and joint operations. Robowars thrilled audiences with intense battles of custom-built machines. The festival concluded with a spectacular Bollywood Night by Sachin-Jigar.

A CELEBRATION BEYOND BOUNDARIES

From insightful talks and intense competitions to electrifying pronites, Techkriti 2026 was more than a festival; it was an experience. With participation from across the country and visionary speakers, it reinforced its role as a platform where ideas come to life.

As the curtains fall, Techkriti continues to inspire the next generation of innovators and leaders, pushing boundaries and shaping the future.

DEEPAK VISWESWARAIAH
Senior Vice President and Managing Director, Pegasystems India

When AI stops suggesting and starts acting

Enterprise AI has reached an inflection point where the central question is no longer what AI can tell the business, but what it can do inside the business process itself. For years, AI largely operated as an insight layer, augmenting human judgement with predictions and recommendations. That model is now giving way to something more operational: AI as an execution layer capable of taking action autonomously within live workflows.

This is already visible in customer service, fraud detection, and workflow orchestration. In these environments, decisions are no longer merely surfaced for human review and delayed action. They are being triggered and executed in real time. At Pegasystems, that means AI can recommend the next best action and automatically launch the associated case-resolution workflow. In financial services, it can detect potential fraud and immediately initiate an investigation or block a transaction. In operations, AI-driven decisioning can orchestrate workflows dynamically across systems without manual intervention, turning insight directly into action.

That is why governance now sits at the centre of the autonomous enterprise conversation. The real barrier is not intelligence alone, but trust. Enterprises need transparency, explainability, and control over how AI-driven decisions are made and executed. Without that foundation, the risk of unintended consequences and compliance gaps becomes too high. AI will succeed as an execution layer not because it is clever, but because it is governed responsibly and consistently at scale.

SAMIT SHETTY
Country Leader, Automation Platform, IBM India & South Asia

AI scales when control stays with the enterprise

The shift towards the autonomous enterprise is already visible in workflow-centric automation, where AI is being embedded directly into business operations across HR, IT, customer service, and finance. The model is no longer limited to support or recommendation. It is increasingly built around systems that can detect, decide, and act in real time across end-to-end workflows.

IBM points to its own "client zero" transformation as proof of that shift. By combining AI and automation, the company says it has delivered USD 4.5 billion in productivity gains globally and designed more than 155 AI use cases across core functions. In HR, its AskHR AI agent now handles 94% of employee queries autonomously, reducing support tickets by around 75% against historical levels, while overall transformation efforts have lowered operating budget by 40% over the past four years. In IT operations, AI-led automation has cut standard support tickets by 56% between 2022 and 2024, with AI agents resolving about 86% of employee queries. IBM also cites work with API Holdings, where AI-driven observability using IBM Instana helped reduce mean time to resolution by up to 30%, and with Karnataka Bank, where a modernised API platform improved scalability by 50% and reduced operational costs by 30%.

As AI starts making decisions and triggering actions, trust depends on sovereignty across data, technology, and operations. Without that layer of control, autonomous AI may scale experimentation, but not enterprise confidence.

PREMALAKSHMI RAMAKRISHNAN
MD & Area Vice President, India and SAARC, NetApp

AI execution will only be as strong as the data beneath it

Many organisations still treat AI as a decision-support layer, but it is already executing operational decisions across the enterprise. In banking and financial services, AI is approving or blocking transactions in real time, detecting fraud, and automating compliance. The same shift is visible in telecom, manufacturing, and the public sector, where AI is helping manage networks, optimise supply chains, and respond to cyber threats as they happen. The bigger change now is that AI is moving beyond applications and into the data and infrastructure layer. Once it begins managing the systems that run the business, including data, cloud, and cyber resilience, it stops being a tool on the side and starts becoming part of the operating model itself.

That is where infrastructure becomes decisive. NetApp underscores the high-throughput AI infrastructure, automated data pipelines, and hybrid multi-cloud data connectivity and governance as the foundation that allows autonomous AI operations to scale securely. In this view, the autonomous enterprise will ultimately be built on data infrastructure because AI can only execute as fast, and as safely, as the environment behind it allows.

The central risk is data trust. Fragmented data across on-premises systems, multiple clouds, and business platforms can lead AI to act on incomplete or inconsistent information. Enterprises need unified, governed, and secure data environments, with clear visibility, access control, cyber resilience, and recoverability, before AI can be trusted to execute at scale.

RAHUL MAHAJAN
VP & CTO, Digital Business Transformation, Nagarro

Governed execution begins with shared context, not hype

The shift from AI as passive adviser to AI as active executor is accelerating in organisations that treat AI initiatives as enterprise-grade products rather than experiments layered onto old systems. The most successful approaches are being built for governed execution from the outset, with guardrails, explainability, evaluation frameworks, confidence thresholds, and carefully designed human handovers embedded directly into operations. That shift is especially visible in agentic platforms, multimodal intelligence, and context-aware architectures where multiple specialised agents coordinate across workflows rather than operating in isolation.

What is changing here is not just model quality, but the system around the model. Shared context is becoming the real breakthrough. When every system operates with the same business understanding and under the same rules, fragmented automation begins to turn into coordinated intelligence. That requires looking beyond conventional AI engineering and investing in specialised agentic governance skills and semantic modelling that make execution more reliable across the enterprise. The harder challenge, however, is change management. It is an even bigger issue than model risk alone. Trust in execution depends on whether enterprises redesign skills, structures, accountability models, and continuous-improvement practices around agentic operations. It also depends on much wider AI literacy across leadership teams, business functions, and front-line workers. Even the most capable autonomous system will fail to scale if people do not understand when to trust it, when to challenge it, and how to improve it. Technological maturity matters, but organisational readiness decides whether AI can truly move into execution.

VINOD KUMAR
Chief Digital Officer, Shriram Finance

Shriram Finance sees AI becoming the operating engine of decisions

Artificial intelligence is rapidly shifting from an advisory instrument into a proactive execution layer across financial services. In the front office, it is already managing large parts of the customer lifecycle, from automated onboarding built on document verification, data extraction, and real-time eligibility decisioning, to conversational systems that not only answer queries but also trigger backend actions. It is also enabling proactive personalisation by deploying pre-approved offers without manual intervention.

The deepest change, however, is visible in credit and risk, where AI increasingly functions as a core decisioning engine. Credit underwriting models are beginning to influence approval and rejection outcomes directly. Fraud systems can autonomously identify, flag, or block suspicious transactions. Alternative data is being pulled straight into automated decision engines. In operations and the back office, AI acts as a digital operator, handling KYC, compliance checks, anomaly detection, workflow orchestration across legacy and modern systems, and collections strategies that adapt in real time to borrower behaviour. The shift is no longer from question to recommendation. It is from question to action executed within defined governance boundaries. That makes explainability and accountability the defining challenge. In regulated environments, black-box decisioning cannot be the basis for trust. Every action must be auditable and justifiable, ownership of outcomes must be clear, and drift or bias must be monitored continuously. Human-on-the-loop governance, explainable AI frameworks, and real-time overrides are what separate scalable AI execution from unmanaged risk.

JAYANTH SEKAR
Sr. Director, Data Science & AI – Customer Experience, GE Aerospace

GE Aerospace puts AI to work across the engine lifecycle

At GE Aerospace, artificial intelligence is no longer an experimental layer sitting beside engineering. It is already operating at scale across the lifecycle of an engine, from design and production to inspection, maintenance, and workforce productivity. One of the clearest examples comes from engine design, where researchers have demonstrated generative AI capable of producing hundreds of design iterations in seconds, compressing the development cycle for future engine technologies and moving AI from design support into active participation.

The same shift is visible in fleet readiness. GE Aerospace monitors its global fleet of roughly 50,000 commercial engines around the clock using AI to identify predictive maintenance issues. That has led to a 45% increase in issue detection, 60% faster lead times in identifying maintenance needs, and a 50% reduction in false alerts. AI-powered inspection tools, including the Blade Inspection Tool developed at GE Aerospace's Bengaluru Technology Centre, have halved inspection times. The company is also using AI to forecast work requirements for engine maintenance so the right parts are available before an engine even arrives at a Maintenance, Repair, and Overhaul facility. Its internal generative AI platform, AI Wingmate, is used by around ~11,000 employees daily.

ARUN RAMCHANDRAN
CEO, QBurst

The real shift is from chatting to doing

The real inflection in enterprise AI is not better conversation. It is the move from reactive chat interfaces to agentic workflows that execute multi-step tasks across systems. That distinction matters because chat-based AI, however useful, still depends on humans to drive the process. Agentic AI changes that by allowing digital workers to plan, react, and collaborate inside business environments with far greater autonomy, provided they operate within clear guardrails.

That is no longer confined to knowledge-intensive work. QBurst points to emerging use cases in sectors such as mining and manufacturing, where agentic AI, and in some instances Physical AI, can perceive the physical environment, verify raw materials against delivery standards, and trigger supply chain alerts automatically. In such settings, AI is moving from information and insight towards action. More importantly, it is beginning to take on what Ramchandran describes as the intent layer of the organisation, translating business goals into coordinated, system-level execution.

The central risk, however, is data accuracy and readiness in the enterprise context. If the data feeding the system is not current, contextual, or consumable, AI-led execution becomes unreliable at best and dangerous at worst. That is compounded by the fact that AI agents do not fail like traditional software. They can fail silently and inconsistently. Trust, therefore, requires a managed-agents approach built on observability, full execution tracing, human-in-the-loop controls, and systems designed for recovery rather than blind automation.

VIJAY VIJAYASANKAR
Global Agentic AI Officer, Genpact

Agentic operations as the next enterprise operating model

Genpact's view of the autonomous enterprise is grounded in what it calls Agentic Operations, a model that shifts work from being human-processed and human-validated to machine-processed and human-validated. The difference sounds subtle but significant. AI is no longer only helping workers move through a process. It is increasingly executing the process itself, resolving exceptions, orchestrating actions, and driving measurable outcomes within defined guardrails.

Accounts payable offers a clear example. Genpact's AP Suite uses a network of pre-trained, self-learning AI agents to autonomously ingest and validate invoices, trace missing supplier data, streamline inquiries, and resolve exceptions while escalating only the most complex cases for human review. The result is greater precision and more straight-through processing, but it also points to a larger truth: AI only becomes truly useful as an execution layer when the enterprise is architected for responsible action at scale, not when agents are simply bolted onto broken workflows.

That is why Genpact ties trust to AI-first redesign rather than isolated deployment. Enterprises need process intelligence at the core, with autonomy designed across data, architecture, orchestration, and governance. If the underlying data layer is not ready, agents become polished interfaces sitting on top of flawed processes. The organisations pulling ahead are the ones deliberately designing AI-led workflows where agents behave like a governed workforce, scaling autonomy without sacrificing trust, auditability, or human judgement.

PRAVEEN OJHA
CTO, EPAM India

The real issue is not the model, but the system

Enterprise AI is moving closer to execution in a number of high-impact environments, but the transition remains uneven. Financial services, manufacturing, retail, and healthcare are among the sectors where AI is being plugged into live operations to improve workflow efficiency and support faster decisions. In financial services, AI is becoming part of transaction flows, enabling real-time risk decisions and more adaptive customer interactions. In manufacturing, it is powering closed-loop systems that learn continuously from sensor data to improve production, quality, and energy use. In retail, it is helping orchestrate end-to-end commerce, while in healthcare it is finding a place inside clinical pathways under strong human oversight.

The common thread is proximity to the system of action. AI is no longer sitting entirely outside workflows. Even so, most enterprise deployments are still not mature enough to be fully embedded into production execution layers. They remain assistive, decision-support driven, and tightly bounded by guardrails. That caution reflects a deeper issue. The real risk is not the sophistication of the model, but the systemic reliability of the environment in which it operates.

In fragmented enterprises with uneven governance, AI often fails subtly through drift, misalignment, and inconsistent behaviour that is hard to detect at scale. That makes trust a matter of predictability, traceability, and control rather than intelligence alone. The answer lies in deliberate re-architecture, where data is standardised, systems interoperate cleanly, and governance is embedded throughout the AI lifecycle rather than added after the fact.

ARUN BALASUBRAMANIAN
Managing Director, India & SAARC, Dynatrace

How AI is already acting inside digital pressure points

In fast-scaling digital environments, AI is already moving beyond analysis and into live execution. That is especially visible in sectors such as digital commerce and banking, where performance, uptime, and customer experience are business critical. During peak events such as festive traffic surges, AI is no longer only flagging anomalies for teams to investigate. It is autonomously scaling infrastructure, optimising workloads, and resolving performance bottlenecks in real time.

What makes this possible is not automation alone, but a combination of full-stack observability and precise, context-aware understanding of how complex systems behave. He points to a wider regional pattern as well. Across APAC, many organisations are already pushing agentic AI into limited production and departmental deployments, with a growing share moving towards enterprise-wide integration. The acceleration is particularly visible in customer-facing applications, where teams are prioritising AI to improve satisfaction while also monitoring agent performance and data quality more closely.

That makes the central risk easy to define: lack of trust in how AI makes decisions. In tightly connected environments, an action taken without the right context can create downstream consequences that quickly become costly. Enterprises need AI grounded in real-time context, guarded by clear policies, and supported by full traceability of actions. Without that foundation, autonomy at scale becomes a liability rather than an advantage. With it, AI can move from being a support layer to a trusted engine for resilience and digital experience.

DEB DEEP SENGUPTA
Area Vice President, South Asia, UiPath

AI delivers value only when execution is orchestrated end to end

The shift underway is not just from insight to automation, but from insight-led systems to execution-led systems. That marks a fundamental change in how work gets done. Traditional AI largely focused on surfacing insights and automating discrete, rules-based tasks. It could recommend, predict, or assist, but the responsibility for execution still sat with people. Agentic AI changes that boundary. It is now taking action inside workflows, resolving cases, enforcing compliance, and in some situations completing transactions end to end.

This is already visible in high-volume functions such as customer operations, finance, and compliance, where consistency, speed, and process discipline matter more than isolated intelligence. The real value does not come from insights alone. It comes when decisions are connected to systems, processes, and actions across the enterprise. In that sense, execution depends on orchestration, not intelligence by itself. Work is no longer routed to people by default, with automation added later. It is increasingly being orchestrated dynamically across humans, agents, and systems based on capability, allowing long-running and complex processes to flow end to end.

The main risk is assuming AI is ready for execution when the enterprise is not. Most organisations still run on fragmented systems, inconsistent data, and workflows not designed for autonomous action. That is why pilots often stall in production. Trust will depend on governed workflows with orchestration, auditability, policy enforcement, and human oversight built in from the start.

RITWIK BATABYAL
CTO and Innovation Officer, Mastek

AI has already crossed into action

The enterprise AI shift is no longer about better answers. It is about systems that can plan, act, and coordinate work across functions with far less human intervention. That change is already visible across customer service, operations, supply chain, finance, compliance, and enterprise IT. AI agents are moving beyond handling frequently asked questions and into operational action, whether that means processing refunds, rebooking services, escalating complex cases, dynamically rerouting shipments, or adjusting production schedules in response to changing supply, weather, or demand conditions.

The same progression is visible in finance and compliance, where autonomous trading and risk systems can monitor live conditions, execute trades, and enforce thresholds simultaneously. Inside enterprise IT, internal agents are beginning to deploy software updates, resolve tickets, and balance workloads on their own. The broader pattern is unmistakable. Enterprises are re-architecting workflows around AI agents that do not just advise humans on the next step, but carry that step forward themselves.

That makes accountability and transparency the foundational trust requirement. Before organisations delegate execution, they need governable intelligence, AI that can be monitored, understood, and constrained within clearly defined human boundaries. Explainability, auditability, and regulatory alignment are no longer optional features. They are the conditions under which AI can move from useful automation to trusted enterprise execution. Without that layer of control, faster action simply means faster exposure to risk.

JAYAPRAKASH NAIR
Global Head of Data and AI – Lab and Capability Center, Altimetrik

Caution, not speed, defines serious enterprise AI

Enterprise interest in AI-led execution is rising, but the movement remains cautious, especially in regulated sectors such as banking, financial services, and life sciences. Altimetrik's view is that companies are willing to reimagine workflows with AI, but only when autonomy can be tightly controlled. Deterministic agents may occasionally be granted full autonomy after extensive due diligence, yet the moment an AI system becomes probabilistic or stochastic, enterprises become much more hesitant. The reason is simple: these decisions are not binary, and even highly experienced humans can struggle with them, let alone models operating with narrower context windows.

That uncertainty is worsened by the fact that AI outputs can change over time. The same system may produce different, even conflicting, answers under different conditions, making the decision path harder to trust. Probabilistic engines still provide real value, but they are far better suited to a model where a human remains in the loop and judgement is preserved at critical moments.

That is why Nair places such strong emphasis on Responsible AI. Fairness, reliability, accountability, explainability, transparency, inclusiveness, privacy, and security are not interchangeable buzzwords. They are distinct disciplines that together determine whether AI execution can be trusted. In practice, most organisations may not wait until every condition is perfect. They will prioritise the most important tenets first and move forward with supervision. But trust will only deepen when building the right product is matched by building it the right way.

CHANDRASEKAR RAMAMOORTHY
Co-founder, Mozark.AI

How AI is reshaping product, testing, and sales together

AI is already deeply embedded in our own product, engineering, testing, and sales workflows, making it a strong example of how execution begins inside the operating core of a business. In product development, AI has changed how work gets initiated. Product managers can now create functional prototypes directly instead of drafting long traditional requirements documents, reducing the back-and-forth that often slows collaboration between product and engineering teams. In development itself, AI agents are generating nearly 50% of the company's code, helping accelerate feature rollouts, shorten release cycles, and increase organisational agility.

As a quality assurance platform, we also use AI to raise test productivity and coverage, enabling better outcomes with less manual effort. The commercial side shows a similar pattern. In sales, AI helps teams instantly generate customised prototypes, dashboards, and analytics based on each prospect's needs, making demonstrations more relevant from the first interaction and helping shorten the journey from lead to order.

The trust question, however, is equally practical. Before AI can be relied on as an execution layer, agents must be rigorously validated for security, accuracy, and performance. As systems gain access to more tools and data sources, the testing burden only increases. The effectiveness of AI ultimately depends on the quality, security, and privacy compliance of the tools and contextual data it relies on. Autonomous execution can only scale when that foundation is assured.

AVEG AGARWAL
India Business Head, Bidgely

India's power sector move from analytics to action

India's power sector offers a compelling view of how AI is moving from insight to execution in environments where decisions must happen at scale and in real time. With more than 250 million smart meters being rolled out under the Revamped Distribution Sector Scheme, utilities are beginning to move beyond analytics towards automated action. That includes dynamic load balancing, outage prioritisation, and far more personalised consumer engagement, all of which make distribution companies more responsive while freeing human teams to focus on higher-order decisions and customer outcomes.

In this setting, AI is no longer simply recommending what utilities should do next. It is helping them act faster, with more granularity, and across a much larger operational footprint. Bidgely argues that this is where vertical AI becomes important, because it allows utilities to deploy specialised machine learning models inside their own cloud environments rather than relying on generic black-box systems outside their security perimeter.

Utility leaders remain wary of moving sensitive infrastructure data beyond their established boundaries, and with good reason. Data sovereignty and model brittleness are major concerns when critical infrastructure is involved. The answer lies in systems that preserve sovereign control while still enabling much richer operational insight. If that balance is achieved, the shift from analytics to execution could unlock not just better utility operations, but very significant economic value at scale.

SAJITH NAMBIAR
Head of Solutions, UST

AI execution where workflows are repeatable and bounded

AI is moving most credibly into execution in environments where workflows are repeatable, data is available, and the operating boundaries are already defined through business rules. That is the pattern UST sees across service operations, document-led workflows, and process orchestration. In IT and service operations, AI is no longer limited to summarising incidents or recommending next steps. It is increasingly classifying tickets, enriching cases with business context, routing work to the right teams, triggering remediation workflows, and in some cases carrying out low-risk actions automatically.

The same progression is visible in procurement, finance operations, compliance, and customer support, where agentic systems can extract information, validate it against business rules, and push the workflow forward without waiting for human intervention at every stage. In such settings, AI stops being a point solution and becomes part of the operating fabric of the enterprise. The key enabler is not just a better model, but a stronger enterprise foundation, one where business context, process logic, policies, and metadata are cleanly connected.

That is why UST defines the main risk as governance of action. A poor recommendation can still be reviewed, challenged, or ignored. An executed action changes the equation completely. Trust depends on guardrails, role-based permissions, policy boundaries, escalation thresholds, and the ability to explain, observe, and reverse what happened. The future will belong not to more powerful AI alone, but to AI that is more controllable, transparent, and safe to act.

SRIVIDHYA SRINIVASAN
Co-founder & CTO, Amagi

When AI becomes the workflow

A fundamental shift is taking place in high-velocity industries such as media and digital publishing, where AI is moving from advisory support to becoming the execution engine itself. Earlier, AI functioned as a digital assistant, summarising data, generating text, or suggesting tags while human operators still navigated fragmented interfaces and manually pushed the work forward. Agentic AI changes that equation by collapsing the interface itself. In this model, AI is no longer recommending a clip selection or draft. It is identifying the narrative, carrying out spatial reformatting, applying brand packaging, and publishing the final asset directly to endpoints.

That matters because AI has moved from advising the workflow to becoming the workflow. In sectors where speed and output variation are high, that shift has major implications for how content is produced, packaged, and delivered. The productivity upside is obvious, but so is the risk. Once AI becomes the execution layer, a hallucination or logic error is no longer an internal misstep. It becomes a live, public-facing mistake.

The real challenge, then, is operational control at scale. Enterprises need to solve the black-box problem by making AI actions traceable and explainable. That means moving away from opaque systems and towards deterministic, policy-driven guardrails governed by plain-English business rules, compliance requirements, and brand standards. Human-in-the-loop approval gates must remain available where needed. Trust will not come from making AI smarter alone. It will come from making its boundaries transparent, explicit, and enforceable.

AARTI KAPUR
Leader, Platform and Solutions, Tiger Analytics

How data platforms are becoming action platforms

The move from AI as support tool to AI as execution layer is already visible in how enterprises are running specific workflows directly on their data platforms. In demand planning and compliance, AI is no longer stopping at insight generation. It is monitoring data, updating forecasts, generating reports, and triggering the next action within the same workflow.

Retail offers a clear example. Demand is being recalibrated in real time, with replenishment decisions triggered directly from the system. In compliance functions, the same pattern is taking hold. Anomalies are flagged, reports are generated, and follow-up actions are initiated without waiting for manual stitching across separate tools. What makes this possible is not only the model, but the architecture. Intelligence is being embedded within the data layer itself, with agent-driven workflows handling specific tasks and moving processes forward.

That also explains the real risk. Trust breaks down when AI acts on inconsistent definitions, fragmented data, or an unreliable understanding of the business. Speed without guardrails is not acceleration. It is liability. The answer lies in a strong data and semantic foundation, shared definitions of metrics and rules, clearly bounded execution rights, and audit trails that make every decision traceable and reversible. AI may be ready to execute, but enterprises still need to become structurally ready to let it do so responsibly.

YOGESH JADHAV
Group CTO, Choice International & Choice Techlab

Why AI confidence can be deceptive

One of the sharpest warnings in the autonomous enterprise conversation is not about hallucination, model drift, or context limits alone. It is about over-trust. Yogesh Jadhav describes the missing capability as calibrated distrust, and the phrase resonates, because it captures a deeply human problem that technology by itself cannot solve.

When an AI agent produces polished output, clean error handling, sensible logging, and apparently strong test coverage, reviewer vigilance tends to fall. Cognitive load drops, and with it, scrutiny. The result is what Jadhav calls quality debt at velocity. Decisions are approved because they look convincing, not because they have been properly evaluated. Over time, an organisation starts accumulating actions that were signed off but never fully understood.

That risk becomes existential in regulated sectors such as financial services. A single unreviewed assumption about data handling, compliance interpretation, or execution logic can create serious regulatory exposure. The danger is subtle because the system may appear competent, even elegant, while quietly embedding flawed judgement into live workflows.

The implication for enterprises is uncomfortable but essential. Trust in AI execution cannot be built on fluency alone. It requires disciplined scepticism, review mechanisms that remain active even when output looks impressive, and operating cultures that resist mistaking polish for proof. Without calibrated distrust, autonomous execution creates not just efficiency, but hidden institutional risk.

TANU GARG
AI Service Delivery Transformation Leader, EY GDS

Embedding AI into the execution fabric of service delivery

At EY GDS, the shift from AI as assistive layer to AI as execution fabric is already taking shape across engineering, risk and compliance, and data operations. The transition is not being driven by isolated tools, but by a structured operating model. A governed marketplace of AI assets and accelerators, aligned to prioritised use cases and cleared through rigorous information security and compliance checks, allows teams to reuse proven solutions at scale. That reduces duplication, shortens time to value, and turns AI from one-off experimentation into an execution capability.

A key differentiator is the codification of domain knowledge into ontology-driven knowledge cartridges. These allow AI to generate outputs that are far more contextual and industry-specific, pushing it beyond generic assistance and into execution-grade intelligence embedded within workflows. Trust is being treated in the same structured way. Evaluation and governance are built into the model lifecycle from the outset, covering quality, security, and responsible use. Live telemetry from engagements feeds standardised evaluation pipelines so performance can be monitored continuously as solutions scale.

The real risk here is not the technology itself, but fragmented and ungoverned adoption. Organisations will struggle if AI spreads unevenly across functions without shared controls, context, or discipline. The ones that succeed will be those that embed AI into how work is done, govern it rigorously, and scale it through reuse rather than reinvention.

KANAKALATA NARAYANAN
Vice President, AI and GenAI, Ascendion

Execution succeeds only when the operating model changes

Software development is emerging as one of the clearest proofs that AI has moved from insight into execution. Agentic systems are now autonomously writing, testing, and deploying code inside continuous workflows, producing measurable gains in productivity, cost efficiency, and time to market. The important point, however, is not simply that AI is augmenting developers more effectively. It is that AI is being embedded as an execution layer within core operations.

That shift changes the conversation from capability to operating model. The biggest blocker is no longer technological readiness. It is whether enterprises have redesigned decision rights, accountability, governance, and performance metrics to match this new reality. Many are still attempting to layer AI onto legacy structures, which is why so many programmes remain stuck in pilot mode despite serious investment. Execution at scale demands more than better models. It demands system-level redesign.

In that sense, the winning pattern is human-AI collaboration by design. AI drives execution where speed, repetition, and pattern recognition matter. Humans retain judgement, control, and accountability where interpretation, ambiguity, or risk rise. Organisations pulling ahead are the ones aligning AI directly to business outcomes and rethinking delivery itself, moving away from effort-based models and towards value-based execution. The lesson is clear. The autonomous enterprise is not built by sprinkling AI on old processes. It is built by redesigning the enterprise around what AI can now do reliably.

VENKAT SITARAM
Senior Director & Country Head, Infrastructure Solutions Group, Dell Technologies India

Dell sees the execution shift spreading across India's enterprise stack

As enterprise AI matures in India, the shift from assistive intelligence to execution layer is becoming visible across core business processes. That includes IT operations, cybersecurity, supply chain, and customer experience, where AI is moving beyond insight generation and into autonomous decisioning and action. Self-healing AIOps and real-time, AI-optimised logistics are among the strongest examples of this broader evolution.

The underlying drivers are equally important. This change is being powered by more capable data platforms, stronger edge computing architectures, and deeper integration of AI into enterprise applications. In other words, the execution shift is not happening because AI suddenly became useful in isolation. It is happening because the surrounding infrastructure is finally maturing enough to support action at enterprise scale.

Dell Technologies positions this as an operationalisation challenge rather than a model showcase. Scalable infrastructure, edge-to-core architectures, secure data foundations, and trusted environments are what allow AI to move safely from recommendation to execution. That makes trust, transparency, and governance essential, not optional. Autonomous execution can only scale if enterprises know how decisions were made, what systems were touched, and where boundaries remain in place. The autonomous enterprise, from this perspective, is not a futuristic slogan. It is the next stage of enterprise architecture, where AI becomes part of how the business runs, provided the infrastructure and governance are strong enough to carry that weight.

RAKESH RAVURI
Chief Technology Officer & Senior Vice President, Engineering, Publicis Sapient

AI execution creates decision debt if left unguided

For years, AI sat in the enterprise as a layer of insight, analysing data, producing recommendations, and supporting human decision-making. That balance is now changing. AI is moving into the core of execution itself, especially in software engineering and enterprise operations. In software delivery, it is not just assisting developers but helping build, test, refactor, and modernise systems. Delivery is becoming more continuous and machine-assisted. At the same time, workflows such as claims processing and customer onboarding are shifting from suggested next steps to AI-driven decisions that are actually executed inside the process.

This marks the arrival of a new architectural layer. AI as a decision engine embedded within enterprise workflows. The deeper risk, however, is not simply hallucination or isolated model error. It is what happens when AI-driven decisions scale without shared context, visibility, and control. Ravuri describes this as decision debt. As organisations deploy multiple agents across functions, each one may optimise locally while creating fragmentation globally. One agent approves what another might reject. Over time, the enterprise accumulates inconsistent decision logics that are neither fully traceable nor aligned.

That is why trusting AI as an execution layer requires more than better models. It demands governance structures where decisions are observable, traceable, and continuously evaluated, supported by shared context layers and clear boundaries between autonomous action and human oversight. The challenge is no longer whether AI can decide. It is whether enterprises can govern those decisions at scale.

GANESH GOPALAN
Co-Founder & CEO, Gnani.ai

Gnani.ai is turning voice AI into an enterprise execution engine

Voice AI is emerging as one of the clearest examples of how artificial intelligence can move from support into execution at scale. In high-volume, decision-intensive functions such as customer support, sales operations, lead qualification, and collections, AI is no longer limited to agent assistance. It is increasingly owning end-to-end workflows, resolving queries, qualifying leads, processing requests, and launching follow-up actions without depending on human intervention at every step.

Gnani.ai's Inya VoiceOS is designed around that transition. It enables voice AI agents not just to understand and respond, but to take action by integrating with backend systems, executing tasks, and completing entire interactions end to end. This has clear operational implications for functions such as collections, outbound engagement, and customer service, where consistency and response speed directly shape business outcomes. The scale of impact is already visible. One deployment with a top-three bank handled over one crore customer interactions, increased call-handling capacity by 270%, and improved customer satisfaction by 50%, all without adding headcount.

The real risk is not whether AI can drive decisions, but whether it can do so predictably and accountably. Trust depends on clear decision boundaries, human oversight where needed, robust monitoring, and the ability to trace why an agent acted the way it did. In this model, conversation analytics becomes a governance layer, not just a reporting one. It gives enterprises the audit trails needed to trust autonomous voice systems in live operations.

SANJAY AGRAWAL
Head Presales and CTO, Hitachi Vantara India and SAARC

Execution begins where data can move at action speed

AI is moving into execution most credibly in environments where real-time action matters and the data infrastructure is mature enough to support it. That pattern is already visible across healthcare, manufacturing, and financial services. In healthcare, AI is no longer just flagging anomalies. It is triggering workflows such as case prioritisation and data routing across clinical systems. In manufacturing, the transition is even more direct. AI has moved from predictive insight to autonomous control, with systems ingesting machine data and adjusting production parameters in real time. In financial services, AI is increasingly blocking transactions, adjusting risk thresholds, and initiating compliance workflows inside highly governed environments.

What ties these examples together is infrastructure. Agentic AI, in this view, is not just another application layer trend. It is meaningful because it plans, coordinates, and executes multi-step actions across many systems and agents. That makes contextuality critical. The system must know which agent should do what, when, and why. None of that is possible without robust hybrid data infrastructure, low-latency integration between information technology and operational technology, and secure environments where action can be grounded in reliable data.

The larger shift, then, is from AI as co-pilot to AI as operator. Outcomes increasingly depend on how quickly and accurately enterprises can turn data into action. Where the infrastructure is resilient and context-rich, AI moves into execution. Where it is not, the system remains stuck in assistance mode.

RAHUL LODHE
Global Vice President, Head of SAP Copilot Joule, and Head, SAP Artificial Intelligence India, SAP

How AI closes the loop inside systems of record

The clearest evidence of AI moving from insight to execution lies not in what it can analyse, but in what it can now complete inside systems of record. The shift is most visible in workflows that are high-volume, rule-rich, and tightly connected to core enterprise platforms. In procurement, for example, AI has moved well beyond extracting data from invoices and purchase orders. It is now validating documents against policy, routing approvals, and writing results back into enterprise resource planning systems without waiting for a human to initiate each step.

The same logic applies in supply chain operations, where production-planning agents can perform prerequisite checks, identify workarounds, and trigger downstream workflows autonomously. In IT operations, systems are moving from anomaly detection to real-time investigation and response. In engineering, an automotive supplier used a squad of AI agents to generate initial test case descriptions from historical requirements, cutting the time required by 50% for certain requirement types. The common thread is architectural. These are agents with defined permissions, access to live systems, and the ability to close the loop by updating records and changing states directly within the workflow.

That makes controllable autonomy the central issue. Enterprises need end-to-end accountability for what an autonomous system did, why it did it, and whether the action can be reversed. Auditability must be treated as a design constraint, not a compliance afterthought. Until then, the rational response will be to keep humans as the final executors.

RANGA JAGANNATH
Senior Director – Growth, Agora

AI execution depends on real-time interaction that never breaks

AI is moving fastest into execution where automation and real-time interaction intersect. Customer experience is a strong example. Voice AI agents are now handling inbound queries, resolving routine issues, qualifying leads, and routing complex cases without human intervention. With platforms such as Agora's Conversational AI Engine, developers can deploy real-time voice agents that understand speech, respond instantly, and manage multi-turn conversations with ultra-low latency. That means AI is no longer simply providing suggestions. It is actively participating in operational workflows.

The same shift is visible in outbound engagement. AI agents are now capable of conducting sales outreach, running surveys, and scheduling appointments at scale. Multimodal AI agents, powered by real-time application programming interfaces, can combine voice, text, and contextual signals to take action during the interaction itself. Conversational AI is also expanding into robotics and connected devices, bringing voice-driven execution into smart products and Internet of Things environments. That promise, however, rests on a fragile condition: reliable real-time performance at scale. Once AI becomes an execution layer, the tolerance for error drops sharply. Latency spikes, audio drops, response lag, or misread intent can directly damage user experience and business outcomes. Trust therefore depends on more than model intelligence. It depends on ultra-low latency, high availability, resilient communication layers, strong oversight, and clear escalation paths. Without that infrastructure, AI remains a tool. With it, AI starts to behave like a true execution system.

RENGARAJAN THIRUVENGADAM
Senior Director – Head of Operations, TransUnion GCC India

Accelerating AI execution on unified platforms and governed trust

AI begins to move from insight into day-to-day execution when the operating model is unified. Across our GCC network, AI is being applied to how platforms run, how products scale, and how decisions are operationalised. In India, this shows up across a wide stack, from contact centres and technology transformation to advanced analytics, risk decisioning, fraud prevention, and identity-driven insight. The common thread is that AI is not being treated as an isolated tool. It is being embedded into common platforms under shared engineering standards and governance.

That convergence model matters because it allows AI solutions to move more smoothly from design to deployment. Engineering, research and development, information technology, and business process management operate as a single enterprise engine rather than as disconnected silos. As a result, explainable and accountable AI can be built directly into operations and customer-facing outcomes at scale. The defining requirement remains trust. AI-driven action must be secure, explainable, and aligned with enterprise policy, governance, and regulatory expectations. That trust should be anchored in a security-first innovation approach, with 'Data Risk Committee' oversight, structured AI risk assessment, and ongoing tracking of global regulatory developments. The organisation's AI principles, focused on fairness, safety, transparency, accountability, and data protection, are aligned with recognised frameworks such as the National Institute of Standards and Technology AI Risk Management Framework. That is the core lesson here. AI earns its role as an execution layer only when governance is built in from the start.

GANESAN KARUPPANAICKER
Chief Technology Officer, Birlasoft

Birlasoft makes the case for AI that runs the business

Artificial intelligence is moving decisively from insight generation to direct execution across core enterprise functions. In areas such as IT operations, supply chain management, and customer experience, it is no longer confined to recommendations. It is resolving incidents, orchestrating workflows, and triggering real-time decisions within clearly defined guardrails. What is powering this shift is a new generation of agentic AI models that combine contextual awareness with enterprise-grade controls, allowing systems to act rather than merely advise.

The most meaningful progress is happening where AI is embedded directly into execution layers across cloud, data, and enterprise resource planning environments. That is where faster decision cycles, stronger operational resilience, and measurable productivity gains begin to show up. At Birlasoft, this evolution is reflected in initiatives such as Optimus, which moves beyond co-pilots towards AI-orchestrated, end-to-end processes where intelligence is built into how work actually gets done. The company has backed that with an AI-first process reimagination framework, reimagination workshops with business teams, defined agent responsibilities, human-in-the-loop checkpoints, and a validated agentic technology stack.

The core risk is governance at scale. Once AI moves from insight to action, the risk shifts from a flawed recommendation to an undesirable outcome. Trust depends on actions remaining auditable, bounded, and accountable, with end-to-end traceability of what the agent did, why it did it, and who owns the outcome.

AVANI PRABHAKAR
Chief People and AI Enablement Officer, Atlassian

AI becomes part of the system of work

Our view is that AI is a cultural transformation first and a technology shift second. That hypothesis matters because the move from support tool to execution layer only works when AI is woven into strategy, ways of working, and day-to-day decision-making across teams. One of the clearest examples is NORA, the Newlassian Onboarding Rovo Agent, built not by engineers but by members of the People team using our no-code AI platform. In just two weeks, the team went from idea to live, embedded agent.

What makes NORA significant is that it is not sitting on the edge of the workflow answering questions. It applies policy, routes requests, and triggers actions across Jira and Confluence directly inside Atlassian's system of work. Since launch, it has answered more than 19,000 onboarding questions and saved over 2,000 hours of manual operations work. It has also accelerated AI fluency from day one, with 96% of new hires becoming weekly active AI users and 45% emerging as AI super users.

The real risk is speed without boundaries. Agents must behave like well-managed digital teammates, operating only on trusted data, within explicit guardrails, and inside an auditable framework. Trust is not only a technical architecture problem. It is a cultural one, requiring responsible technology principles, clear guidance, leadership modelling, and safe spaces for teams to experiment and learn.

RHYS OXENHAM
VP & General Manager for AI, SUSE

Agentic AI turning infrastructure into active execution

The shift from AI support to AI execution is not being built inside closed ecosystems, but through the open source community. In its view, open source provides the transparency, flexibility, and vendor-agnostic architecture enterprises need to handle complex, data-intensive, agentic systems at scale. The recent rise of the OpenClaw project is one example of how quickly that shift is becoming real, showing that highly complex, multi-step workflows can be automated through autonomous agents.

Two early execution zones stand out. The first is physical AI, robotics, and the edge. In factory settings such as food production plants, agentic workflows can use computer vision to detect anomalies or labelling errors in real time and trigger immediate corrective action locally, where latency matters most. The second is IT infrastructure and operations. Agentic workflows can now detect vulnerabilities, assess impact on production clusters, test patches in sandboxes, and deploy remediation in line with company policy. We are also working on native Model Context Protocol extensions across its Linux and Kubernetes portfolio to provide a governed bridge between large language models and enterprise tooling.

The trust challenge is inseparable from governance and digital sovereignty. Once agents gain access to proprietary data, critical infrastructure, and enterprise tools, organisations need least-privilege access, rate limiting, continuous auditability, revocation controls, and clear escalation paths. Trust will depend on how securely and transparently agents operate across the full stack.

ANAND (JUDE) KANNABIRAN
Vice President, Delinea (APJ-Asia)

Identity is becoming AI's execution control plane

In India, one of the clearest signs of AI moving from insight to execution is emerging in identity and security control planes. In sectors such as banking, financial services, and fintech, AI can potentially block Unified Payments Interface transactions automatically while enforcing step-up authentication in real time. In large enterprises, it can dynamically grant and revoke privileged access, rotate credentials, and enforce least privilege with little or no human intervention. In this model, identity itself becomes part of the execution layer. AI is not merely detecting risk. It is controlling access and containing threats at operational speed.

That is especially important in a market like India, where transaction volumes are massive, systems are complex, and regulatory scrutiny is intense. Autonomous enforcement, therefore, cannot be treated as a futuristic possibility. It has to be operationally ready, deeply governed, and aligned with compliance frameworks from the start.

The risks are equally clear: uncontrolled privilege, weak audit controls, and over-automation without governance. In identity environments, these risks are amplified because one poor decision can ripple across systems at scale. Trust depends on AI acting within strict identity controls, enforcing least privilege, maintaining full audit trails, and preserving human oversight in higher-risk scenarios. AI execution only becomes viable when identity, compliance, and governance frameworks move together rather than in parallel.

YADI NARAYANA
Field CTO, Asia-Pacific & Japan, Datadog

Observability is the trust layer beneath execution

The shift from AI insight to AI execution is further along than many enterprises admit, though it remains more nuanced than the hype suggests. In several environments, AI has already moved beyond surfacing anomalies for human review and is beginning to close parts of the operational loop within defined guardrails. One of the strongest examples is incident response, where AI can identify root causes, recommend actions, and, in more mature settings, execute predefined remediation steps such as rolling back deployments or isolating services.

As AI gets closer to live execution, the decisive requirement becomes observability over the AI itself. Datadog's view is that observability is not a passive monitoring function sitting alongside the system. It is the intelligence layer underneath it, ensuring that the signals AI acts on are accurate and correlated, surfacing drift and unexpected data flows, and creating the audit trail that makes the system governable. Without that visibility, trust becomes fragile very quickly.

The biggest risk is the absence of observability over AI behaviour in production. Models drift. Prompts can be manipulated. Workloads can run away and create costs or operational disruption. The more authority enterprises delegate to AI, the more important real-time telemetry becomes. The organisations that scale execution confidently will not be the ones moving fastest, but the ones that can see clearly what the system is doing and intervene decisively when needed.

HARI ATMAKURI
GVP, Data, AI & Products, Providence India

AI becomes part of healthcare execution

At Providence, one of the clearest changes is that AI is no longer confined to clinical decision support. It is starting to reshape the workflows themselves. Clinical early-warning systems are a strong example. AI agents continuously monitor patient data, detect early signals, validate them against clinical criteria, and trigger timely alerts for intervention. That means AI is not simply helping clinicians interpret information. It is actively shaping how the workflow moves.

The same pattern appears in more complex coordination settings, such as preparing cancer cases for specialist review. A process that once required heavy manual effort to bring together patient histories, diagnostics, clinical guidelines, and trial options is now being organised into structured, decision-ready summaries. In clinical registry abstraction, AI can extract data from unstructured records, validate it against reporting standards, and generate structured submissions while sending exceptions for human review. Providence's registry management platform has already been deployed across 38 facilities, delivering up to twice-faster submissions, 25–30% cost efficiency, and 85–90% improvement in data quality. The organisation also points to its work on central line-associated bloodstream infections, where a machine-learning model built on data from over 83,000 patients helps identify risk within a 24–72 hour window across 56 facilities. Trust in this environment depends on context, compliance, consistency, and reliability. In healthcare, AI execution has to remain observable, auditable, aligned to clinical and ethical boundaries, and strengthened by human oversight where judgement matters most.

VIJAY GUMMADI
CEO, Autorox

How AI is fixing the garage workflow

The auto repair industry has long struggled with a trust gap, and much of that comes from the way work is still carried out. Handwritten job cards, diagnostics stored in a technician's head, and endless messaging for approvals may seem ordinary, but they create structural chaos that makes it hard to scale workshops with consistency. We are tackling that problem by embedding AI directly into the workflow rather than treating it as a separate advisory tool.

Its AI Service Advisor acts as a real-time guide for technicians. Based on a customer's complaint, it walks the technician through a structured diagnostic checklist, suggests the required parts and labour, and turns the job card into a living workflow rather than a static note. Estimates are effectively locked in, approvals can be triggered instantly, and billing errors are caught before invoices are generated. The result is that a junior technician can begin to perform with the consistency and precision of a far more experienced professional because the process now carries the expertise within it.

The trust question is straightforward. Speed without accountability is dangerous. If AI acts on messy or incomplete data, the business does not just get a flawed report. It gets a broken process. That is why every AI-driven action needs a visible, auditable trail, along with constraints and guardrails that preserve human control over critical decisions. The real opportunity is not speed alone, but radical transparency.

PHIL LEWIS
Senior Vice President, Solution Consulting International (EMEA & APJ), Infor

AI is trusted only when it lives inside operations

AI is moving into execution most effectively when it is embedded within operational systems rather than sitting outside them as a separate analytics layer. Infor sees that transition clearly across distribution, food and beverage, and manufacturing. In distribution, the shift is visible in warehouse and order management, where AI is not just forecasting demand but actively orchestrating fulfilment by prioritising orders, optimising pick-pack-ship workflows, and dynamically repositioning inventory across locations.

In food and beverage, the execution stakes are even higher because variability affects margins, product quality, and compliance. Here, AI is being embedded within industry-specific enterprise resource planning systems to adjust production runs based on ingredient availability, optimise batch yields, and maintain real-time traceability. On the manufacturing shop floor, the most advanced use cases are appearing in 'Manufacturing Execution Systems', where AI synchronises schedules with live machine data, triggers maintenance before failures occur, and responds to line variability without disrupting throughput.

The common thread is context. Purpose-built platforms that give AI the operational awareness to act within real business constraints. That is also why trust breaks down when AI is treated as an add-on layer that generates recommendations disconnected from how the enterprise actually works. Enterprises will trust AI as an execution layer only when it is grounded in real workflows, real-time trusted data, and the operational interdependencies that shape actual outcomes.

STEVEN SCHNEIDERMAN
VP, Technology and Forward Deployed Context Engineer, Emids

The biggest risk enterprises must address is not the model itself, but the lack of context

The strongest production-level use cases for AI execution are emerging in operational environments where decisions are repeatable, rules-driven, and time-sensitive. In healthcare, that includes prior authorisation intake, documentation review, claims triage and routing, payment integrity workflows, provider data validation, and member service processes. In these settings, AI is not simply surfacing an insight for someone to act on later. It is assembling documentation, validating data, triggering next steps, routing cases, and in some situations initiating transactions within carefully defined guardrails.

What makes this shift real is that enterprises are no longer deploying AI as a standalone model. They are embedding it into operational systems that combine data, workflow context, decision logic, and governance. Once those pieces come together, AI can take on execution responsibilities for specific tasks while humans retain oversight for exceptions, clinical judgement, and high-risk decisions. That is where the most credible adoption is taking shape.

The biggest risk is not the model itself, but the absence of context and governance inside the workflow. If AI does not understand the operational, regulatory, and data environment in which it acts, it can execute quickly but incorrectly, creating compliance and operational risk at scale. Trust depends on clear decision boundaries, full auditability, and governance frameworks that treat AI-driven actions with the same seriousness as human-led operational decisions.

ASHISH MODI
President – India & APAC, Honeywell

Industrial AI moving from insight to controlled response

The move from AI insight to AI execution is becoming visible in environments where large volumes of operational data must be analysed and acted upon in real time. Industrial settings offer some of the clearest examples. In manufacturing, AI is no longer sitting on top of dashboards as a passive layer of analysis. It is helping teams make small but critical adjustments continuously, stabilising processes, managing variability, and sustaining operations so that human teams can focus more on higher-value decisions.

The same shift is visible in logistics and warehousing, where routing, inventory movement, and order prioritisation are becoming more dynamic because systems can react to changing conditions rather than follow static rules. In buildings, AI has moved beyond tracking energy use and occupancy to actively controlling heating, ventilation, air conditioning, lighting, and overall building performance. Much of this happens quietly in the background, but it is already changing how efficiency and sustainability are delivered in practice.

The decisive risk is data reliability. In industrial settings, AI is only as reliable as the operational data beneath it, and poor-quality or fragmented data can affect safety, uptime, and business continuity. Honeywell's position is that enterprises need trusted, explainable, domain-trained AI supported by strong data integrity, cybersecurity, and human oversight. The shift to execution can scale only when the underlying operational technology environment is understood deeply enough for AI to act predictably and safely.

SREE BALAJI
Co-founder and Group CEO, iLink Digital

Execution begins when data, context, and action align

The clearest enterprise movement from AI insight to AI execution is happening where companies have already built a unified, real-time data foundation. In those environments, AI is embedded directly into the core data and application stack rather than operating as a separate analytics layer. That allows it to act on business events as they happen instead of simply interpreting them after the fact.

Balaji points to a strong example in commercial bid intelligence. What was once a manual review of requests for quotation and blueprints has been turned into an AI-powered optical character recognition and decisioning system. The platform now ingests documents, evaluates specifications, determines feasibility, and ranks opportunities automatically, allowing a single user to manage tens of thousands of requests while prioritising the most valuable ones in real time. In another case, iLink helped unify data across multiple enterprise systems into a central platform. Once that foundation was in place, AI could move from generating insights to enabling real-time decision-making and workflow triggers across the business.

The biggest risk is allowing AI to act without a shared definition of business context. In fragmented enterprises, data is interpreted differently across teams and systems. That may be manageable in analysis, but it becomes dangerous in execution. Trust depends on strong semantic foundations, cybersecurity by design, clearly defined execution boundaries, and controlled autonomy that expands only as reliability and confidence grow.

SUMEET AGRAWAL
Vice President, Product Management, Informatica at Salesforce

Trusted data is what turns AI from advice to action

The shift from "AI told me" to "AI did" is becoming visible where enterprises have built a trusted and governed data foundation. Informatica's Global CDO Insights 2026 study points to AI increasingly moving beyond identifying issues and into autonomously resolving them. One of the clearest examples is inside data pipelines, where AI is now being embedded to detect and fix quality problems in real time, preserving data integrity and context rather than leaving those issues for human teams to clean up later.

That same pattern is spreading across IT, customer service, and sales. AI is beginning to close tickets, trigger follow-up actions, and move pipeline workflows ahead without manual intervention. The broader significance is that AI is starting to act confidently inside production environments instead of remaining confined to pilots and proofs of concept.

The real risk is weak governance. Once AI starts updating records, triggering workflows, or engaging customers directly, every decision must be transparent, auditable, and controllable. Without that, enterprises risk costly errors and erosion of trust that can be difficult to repair. Informatica's position is straightforward: the question is not simply whether AI can execute, but whether the enterprise's data infrastructure and governance framework are strong enough to support responsible, trustworthy execution with full accountability. Trusted data is not an input to AI execution. It is the condition that makes that execution possible.

JAIDEEP VIJAY DHOK
Chief Operating Officer – Technology, Persistent Systems

AI only scales when governance moves inside the workflow

AI is rapidly moving from insight and recommendation into real-time execution inside core enterprise workflows. The most visible impact is appearing where AI is integrated into operating models rather than deployed as a standalone intelligence layer. That is the transition from systems of record to systems of action. Software development and engineering operations are one strong example, where AI is being used across code generation, testing, and deployment to improve delivery at scale. Process automation is another, with rule-based approaches evolving into agentic, outcome-led workflows that can initiate, manage, and optimise tasks with minimal intervention. Predictive analytics is also changing shape, turning into real-time decisioning systems that act on signals rather than merely reporting them.

A particularly telling example comes from a large United States commercial bank that worked with Persistent to embed AI into its lending workflow. The real breakthrough came not at pilot stage, but when the bank operationalised AI within the process itself, standardising data quality, codifying financial parameters, and embedding human oversight to meet risk and regulatory requirements. That allowed the bank to move from AI-assisted insights to execution-ready credit memo generation at scale. The central risk that hampers AI execution is weak data foundations, especially gaps in quality, lineage, and governance. Trust depends on governance-by-design, with explainability, accountability, security, and compliance built into the workflow itself. Automating an existing process as is will only take AI so far. Real returns will come when businesses redesign products and processes to make AI native to the experience.

SUHAIL GULZAR
Senior Manager, Solutions Engineering, Neo4j

Execution becomes trustworthy when explanation is structural

AI has already crossed into execution in ways many enterprises underestimated. Across Neo4j's customer base, the pattern is clear. AI is no longer sitting on the sidelines, handing people a report. It is inside the workflow, making decisions, triggering actions, and closing loops. Uber's ConfigGraph is a case in point. Built with Neo4j and GraphRAG, it validates configurations in real time across domains, catching misconfigurations before they spread. Product launches that once took weeks of careful coordination can now happen in minutes because the system is executing with graph-based guardrails built in.

Other examples make the same point. Walmart's Cross-Channel Insights initiative combines a knowledge graph with a ReAct agent to turn employee feedback into actions that programme managers can use immediately. In life sciences, Pfizer is using knowledge graphs and GraphRAG to run drug manufacturing at a scale that would be almost impossible to orchestrate manually across factories worldwide. In each case, the knowledge graph is not just storage. It is the reasoning substrate that lets AI act consistently across complex dependencies. That is why explainability at the moment of action becomes the decisive risk issue. When AI is executing, post-mortem explanation is not enough. Enterprises need decisions that are explainable by construction. Neo4j's argument is that explainability cannot be bolted on. It has to be structural, with every decision traceable through meaningful relationships that humans and regulators can inspect and challenge.

RAGHAVENDRA CHINHALLI
CIO, HGS

AI is becoming the execution layer of experience and operations

Across the technology landscape, AI is moving quickly from insight generation to Intelligent Experience-led execution. In customer experience operations, generative and agentic AI are no longer confined to analysing journeys. They are beginning to predict intent, orchestrate the next best action, and autonomously execute responses across voice, chat, and digital channels. That ability to deliver predictive, personalised, and outcome-driven experiences at scale is becoming a genuine differentiator.

A similar shift is visible in HR and employee experience services. AI platforms are being used to assess skill gaps, recommend personalised learning pathways, automate HR interactions, and proactively nudge employees. The value lies not only in reducing manual effort, but also in strengthening engagement and workforce readiness for client delivery. In internal support functions such as finance, HR, and supply chain operations, AI is increasingly executing rather than advising. Embedded into core enterprise resource planning and operational systems, it can ingest transactional data, demand signals, and external market inputs to forecast requirements, flag anomalies, automate approvals and replenishment, and trigger interventions before problems escalate. The single biggest risk is opaque decision-making combined with unresolved data ownership and privacy concerns. Once AI starts acting directly in sensitive domains such as customer experience, pricing, or service resolution, lack of transparency becomes a control problem. Trust depends on explainable AI, strong data governance, clear usage rights, auditable logs, privacy safeguards, and human oversight in high-impact decisions until AI has earned the right to operate more independently.

KISHAN SUNDAR
Senior Vice President & Chief Technology Officer, Maveric Systems

Banking AI moves from review to regulated action

Artificial intelligence has moved beyond assisting banking workflows and is beginning to operate them in real time, especially where execution was once slowed by manual checks and fragmented compliance processes. The clearest shift is visible in Know Your Customer and Anti-Money Laundering during onboarding, long considered among the biggest bottlenecks in financial services. AI is now validating documents, running sanctions checks, assessing risk, and triggering onboarding decisions in minutes rather than days, with human intervention reserved largely for exceptions.

The same movement is visible in fraud and payments. AI is no longer just identifying suspicious patterns for teams to inspect later. It is blocking transactions, launching investigations, and enforcing compliance actions instantly. The shift extends into software delivery as well, where AI agents, grounded in enterprise requirements, architecture, and policies, can generate code, create test cases, execute validations, and prepare releases. In customer operations, they can close complaints, update records, and trigger workflows without escalation.

That changes the trust equation completely. The real concern is not just bias or hallucination in isolation, but their combination with regulatory exposure and weak accountability. In a sector governed by model risk, enterprises need clear guardrails around what AI can execute, full traceability of actions, and human oversight in higher-risk scenarios. AI earns trust only when it remains governed, auditable, and accountable.

RAJAN SETHURAMAN
CEO, LatentView Analytics

Execution scales where outcomes are measurable

The movement from insight to execution is becoming most visible in parts of the enterprise where decisions happen frequently, outcomes can be measured, and quick wins create the confidence to scale. Supply chains provide a clear example. AI is no longer only forecasting demand. It is triggering replenishment, rerouting shipments, and dynamically pricing inventory. In product innovation, it is continuously folding real-time signals into portfolio decisions, compressing cycles that once took quarters into weeks. In customer operations, AI agents are resolving queries end to end rather than merely assisting human teams.

What stands out is that the organisations making genuine progress are not chasing autonomy for its own sake. They are starting where decisions are repeatable, measurable, and reversible. That disciplined approach matters because the biggest risk is not the absence of model sophistication, but the absence of robust data governance. Enterprises need to know where their data comes from, who is using it and for what purpose, whether it meets regulatory obligations, and whether the models built on top of it continue to behave as intended over time.

These are not separate governance problems. They form a single discipline that many organisations still have not developed at the pace that AI deployment now demands. Governance is also what protects against drift and bias, particularly in high-impact use cases such as fraud detection where small shifts can have outsized consequences. In the autonomous enterprise, the real bar is not merely whether the model was right. It is whether the organisation can understand, justify, and stand behind what it did.

AMIT YADAV
Vice President and Head – Global Delivery, Kellton

The rise of systems built to act

The most meaningful shift is happening where Agentic AI moves beyond suggesting and begins operating. The transition is from traditional systems of record to systems of action, where enterprise platforms do not just store or analyse information but actively carry work forward. In logistics and manufacturing, that means AI is no longer simply predicting a supply chain bottleneck. It is autonomously rerouting shipments and adjusting inventory orders directly inside the enterprise resource planning system, delivering faster response times and sharper operational efficiency.

A similar pattern is emerging in digital engineering. AI is moving from code assistance to autonomous, self-healing systems that can detect production failures, trace root causes, and deploy patches automatically. This reduces downtime and lowers reliance on manual intervention. By integrating AI agents into enterprise middleware, organisations can execute complex, multi-step workflows such as hyper-personalised customer journeys and automated claims processing at a speed and scale traditional insight-led models cannot match.

The biggest risk is uncontrolled autonomy. Once AI is allowed to act, even a small mistake can turn into a high-stakes business liability. A hallucination in a live workflow that can move money, alter inventory, or modify records is no longer a technical nuisance. It is an operational risk. Trust depends on strict guardrails, strong access boundaries, auditability, explainability, reversibility, and a safety-first architecture that keeps humans in the loop for critical decisions.

NITIN CHANDEL
GVP & India Country Manager, UKG

AI takes the shift-planning burden out of frontline operations

Artificial intelligence is moving beyond recommendations and into execution, especially in structured, compliance-heavy environments where speed and accuracy matter. One clear example is frontline workforce scheduling. What was once a manual, time-consuming task is increasingly being handled by AI agents that do more than suggest possible replacements. They can identify qualified workers based on skills, certifications, availability, and preferences, coordinate shift swaps, and update schedules with minimal human intervention. That marks a meaningful move from AI as a support tool to AI as an execution layer inside day-to-day operations.

The real barrier to trust, however, is governance and accountability. Once AI starts acting autonomously, enterprises need to know how decisions were made, whether those decisions were fair, and who is responsible when something goes wrong. Without that clarity, the risks multiply quickly: opaque decision-making, bias from flawed or incomplete data, compliance failures, and uncertainty over ownership of outcomes.

That is why governance has to be built in from the start. Enterprises need clear policies on AI use and accountability, cross-functional oversight across technology, legal, compliance, and HR, and systems that make decisions traceable and explainable. Bias monitoring, fairness checks, and alignment with local and global regulations also become essential. Human oversight must remain in place for critical decisions. AI can scale execution, but it cannot replace accountability.

SUMEET MATHUR
Senior Vice President and Managing Director, ServiceNow India

HR is where AI starts to execute

One of the clearest examples of AI moving from support to execution is now emerging in HR service delivery (HRSD), where workflows are structured, repetitive, and closely tied to employee experience. ServiceNow points to Coforge as a live example of what this looks like at scale. The company rolled out an AI-powered HRSD platform across 35,000 employees globally, with AI agents handling predictive onboarding, automated goal-setting, and real-time performance analytics. That is no longer a trial run or a proof of concept. It is AI embedded directly into enterprise operations.

What makes the example more telling is that Coforge is itself a leading ServiceNow partner, deploying similar systems for clients while also choosing to run the platform internally first. That reveals something important about enterprise confidence in AI execution. Trust is rarely built in the demo. It is built when organisations use the system themselves, watch it hold under pressure, and see that it can operate within real business conditions.

The bigger obstacle, however, is fragmentation. Many enterprises have pockets of AI that work well inside one system but fail the moment they need to operate across two. The problem is not model intelligence alone. It is that the AI cannot see the full workflow, access connected data, or complete actions across systems with different owners and rules. Until the enterprise becomes more connected, AI execution will remain partial, not truly autonomous.