



TECH NAVIGATOR: SECOND EDITION APPLYING AGENTIC AI TO INDUSTRIES





CONTENTS

Executive summary	4
Unified operational control in oil and gas	6
Inventory discrepancy in the retail industry	12
How to manage discretionary risk in financial services	20
Agentic AI is redefining the manufacturing value chain	30
How agentic AI is re-engineering telecom business sales	44

Executive summary



In our [last edition of Tech Navigator](#), we explained the rise of agentic AI and how it amplifies business. AI agents are already enabling organizations to reimagine and re-engineer business processes — when the ROI is justified — and on a smaller scale, drive process automation and simplification. Beyond re-engineering, ongoing experiments pave the way for vibe coding and exploration for business users to continuously innovate.

Many large enterprises have found that implementing intelligence at scale is easier said than done. They need an architecture-first approach, and they need to treat AI as a transformative program. At the same time, it is important to set up the right hub and spoke operating model, backed by a value- and outcome-driven AI investment approach.

On the technology front, building AI runways that power platform and democratization capabilities across all layers of the architecture is proving key to effective AI scaling.

Agentic AI is already proving its value in broad domains, such as IT, business operations, and information intensive business processes. However, in this new Tech Navigator report, we shine a spotlight on how Infosys is using agentic AI to solve real-world challenges in five distinct industries.


Agentic AI now has the power to transform mission-critical domains: unified operational control in upstream oil and gas companies; inventory discrepancy in consumer, retail, and logistics organizations; discretionary risk management at global banks; aircraft and

shop-floor maintenance in manufacturing; and business sales process in telecom companies. In so doing, we take the reader on a whistle-stop tour of what's really needed to take agentic AI from a point solution to the nexus of digital and AI transformation.

But technology isn't everything. Highly autonomous, most agentic AI implementations also rely on humans for governance and control. Harnessing the requisite talent is therefore not just important but critical to any agentic AI change management strategy. Organizations shouldn't just hire more data scientists and AI experts but look at how this new technology will redefine current roles and require training in the continuous, adaptive oversight of agentic AI systems.

The goal? Higher levels of efficiency, adaptability, creativity, and business insight, all needed in a world where threats multiply — from economic volatility to supply chain disruptions.

Both pragmatic and highly relevant, this edition has been designed to give insight into how agentic AI could rise to be one of the most important tools in an organization's strategic arsenal — depended on by executives and consumers alike.



Unified operational control in oil and gas



Agentic AI does not always play well with traditional, sometimes antiquated, software that is often critical to operations in asset-heavy industries. However, the new model context protocol forms a bridge between these systems, maintaining unified operational control while allowing for the synthesis of real-time information.

Oil and gas industry executives increasingly agree that artificial intelligence is essential for understanding the vast data generated by Internet of Things (IoT) devices that are multiplying in asset-heavy industries. Advanced technologies can harness this data to optimize operations and enhance decision-making. In fact, [59% of global oil and gas executives](#) recently told Infosys that AI will make a significant contribution to their revenue in three years. And [75% say](#) that AI investments will deliver a measurable competitive advantage in this same time frame.

At Infosys, we're supporting our clients' use of AI across the energy value chain — embedding the technology into oil fields, drilling rigs, compressor stations, and distribution networks, among other critical upstream, midstream, and downstream resources (Figure 1). As a result, these smart assets lead to gains in efficiency, resiliency, and operational integrity. And through the use of generative AI, IT and operational technology (OT) data streams provide insights into plant configurations in real time, reduce the downtime of vessels, and support maintenance planning. Further downstream,

AI creates a smarter workplace through persona-based copilots, including personas in smart inspections and field operations, marketing, and energy trading.


In this chapter we concentrate on upstream operations — as shown on the left side of Figure 1 — which includes exploration, drilling, and extraction of crude oil and natural gas. We look specifically at how agentic AI can use the

model context protocol (MCP) to query real time and historical data from proprietary software management systems, leading to autonomous, efficient, and fail-safe upstream operations.

Agentic AI for upstream operations

Upstream energy operations still depend largely on **fragmented commercial off-the-shelf (COTS) software and proprietary**

Figure 1. AI blueprint across the energy value chain



Source: Infosys

systems for drilling, production, and reservoir management. These specialized tools operate in isolation with limited interoperability — a manageable problem in the past but now a limiting factor.

With these constraints, companies will not be able to take advantage of the benefits of agentic AI — the next evolution of AI — where software systems operate autonomously with minimal human intervention. In our [Business Value Radar report](#), we found that agentic AI is a highly successful and popular use case across industries and can be used as a low cost way to fundamentally change how companies operate. In fact, orchestration of AI agents is the most commonly pursued AI use case type in our survey of over 3,200 companies, and it has a high success score of 1.16 (where a success score of 1 is average), placing it in the top 10% of successful AI use cases.

New systems, meet old systems

But to be effective, AI agents will have to be able to go where they are needed and interact with large swathes of the company's technology. However, older software systems weren't built with interactivity in mind. These often archaic COTS systems can't easily expose their capabilities and tools to AI agents, and so the agents themselves can't work as intended.

Even after [migrating to the cloud](#), energy companies face a [variety of challenges](#), such as vendor lock-in, incompatible data formats, fragmented data sources, proprietary data formats, and inconsistent APIs across

platforms. These limitations keep AI agents from accessing important data sources needed for AI-driven optimization and advanced decision-making capabilities.

To stay competitive, upstream energy operations must find a way for legacy software and multiplatform — often closed — cloud ecosystems to easily interact with agents and orchestrate core business processes.

The power of MCP

Asset-heavy industries have found that MCP server architecture provides an intelligent tooling layer between AI agents and existing systems. MCP servers can provide standardized connectors for major platforms, enabling the necessary unified operational control and increasingly valuable synthesis of real-time information.

The MCP architecture, introduced by [Anthropic](#) in November 2024, enables agents to access legacy systems and other COTS and software-as-a-service solutions that have traditionally been siloed and locked in.

MCP is an open standard that enables developers to build secure, two-way connections between their data sources and AI-powered tools. Think of it as a USB-C port for AI applications. Just as USB-C standardizes device connections, MCP standardizes how AI models connect to external data and tools. Before MCP came along, organizations labored with fragmented AI integrations, where every application required custom connections to databases, APIs, and

enterprise COTS systems. Made specifically for agentic AI, MCP solves the integration issue and democratizes the use of advanced AI in the enterprise.


By exposing these COTS systems as an array of “skills” to agents — such as data skills for enterprise data stores or package skills for SaaS or COTS packages — MCP enables agents to operate intelligently across multiple functions. In this context, a skill is a self-contained ability embedded in the agent, allowing it to perceive, reason, or act toward

a specific goal, such as sending emails, retrieving weather info, or recognizing intent in conversation. Exposing the skill makes that skill — essentially a discrete, callable capability — available for the agent to use within its operational environment.

These agents can then allow for partial, or even full, automation in a wide range of business processes (Figure 2).

The protocol has evolved rapidly with pre-built MCP servers for popular enterprise

Figure 2. A version of the MCP architecture



Source: Infosys



systems like Google Drive, Slack, GitHub, Git, Postgres and Puppeteer, making it easier for developers to integrate AI into existing workflows.

And there are custom solutions too. Infosys created an MCP server for [Aveva PI](#), packaging its COTs features as skills to expose them to AI agents. We then created an ITOps agent using the PI MCP server to resolve common IT tickets and another business agent to enable users to query historical and real time data from Aveva PI systems for analysis. At Aveva PI, we were able to use the ITOps agent to reduce IT ticket volumes by 50%, and reduced turnaround time for asset configuration from four days to four hours.

To ensure [security](#) and control, classification of MCP servers as [OAuth Resource Servers](#) (using the June 2025 specification) — which enables resource indicators to prevent token misuse, and much clearer security best practices — ensures agents get only the data and capabilities they're allowed from the COTs systems in the background.

This connected ecosystem increases interoperability among fragmented, COTS-heavy, proprietary systems and brings it into the agentic AI era — a boon for traditionally slow-moving energy and utilities businesses.

How to reduce implementation risk

The tech industry is a fast adopter of MCP protocols, connecting common products to agents, including databases like Postgres and SQLite; collaboration tools like Slack; and development tools like GitHub and Git. On

the other hand, the upstream energy industry is further behind, and it will take time before MCP is widely used.

While the value is clear, companies should proceed cautiously when implementing MCP for business-critical areas. Agentic AI systems are nondeterministic and introduce vulnerabilities into business processes given the nature of the underlying generative AI.

These vulnerabilities include giving different outputs for similar inputs (you never know if the answer is going to be the same the second time around) and actually changing the nature of the environment in which they operate. Prioritizing safety, security, and reliability during the implementation of MCP is critical.

At Infosys, we recommend a three-phase implementation of MCP.

1. Choose an area of the business that has high process predictability and automation potential, where outcomes are more deterministic. A good example of this is IT operations, where processes such as resolving tickets and answering customer queries are bounded and agents can work within well-defined parameters.
2. Introduce agents with tools and skills that provide read-only information access. This phase ensures the agents aren't changing

data in the underlying COTS software and reduces business risks.

3. The MCP registry provides access to agents with tools for analysis and decision-making. Throughout the process, the focus should be on proving the patterns of deployment of MCP servers, and the creation of a secure and dependable MCP registry with appropriate access control.

Once agents are up and running, and safeguards are in place, businesses can extend MCP to other areas, deploying agents to other parts of the enterprise.

Agentic AI as a strategic imperative

Agentic AI is a great advance for the complex needs of asset-heavy industries. In the upstream oil and gas market, enterprises need intelligent systems to prevent costly downtime and to optimize production. However, they can't simply shed decades of technology that is critical to operations.

MCP servers that connect to industry COTS packages create a flexible infrastructure by establishing standardized interfaces. This allows AI agents to orchestrate actions across existing technology without disrupting workflows. This enables autonomous intelligence, while maximizing safety and efficiency and ensuring AI investments deliver a measurable competitive advantage.

Inventory discrepancy in the retail industry



Global retail depends on delivering what customers want, when they want it, and where they want it. Inventory discrepancies can lead to overselling, underselling, and even bans in some marketplaces. Agentic AI can quickly identify inventory discrepancies, correct them, and safeguard critical operations.

Retail has become a fraught industry. The emergence of the “zero customer” — zero boundaries, zero patience, zero loyalty, and demands for net zero emissions — has made the introduction of artificial intelligence a significant tool in the retailer’s software suite. Without it, retail organizations would lack the predictive capabilities to meet customer expectations.

In this climate, not only AI but technologies like robotics and more traditional process automation promise more efficient and effective operations. Many executives are

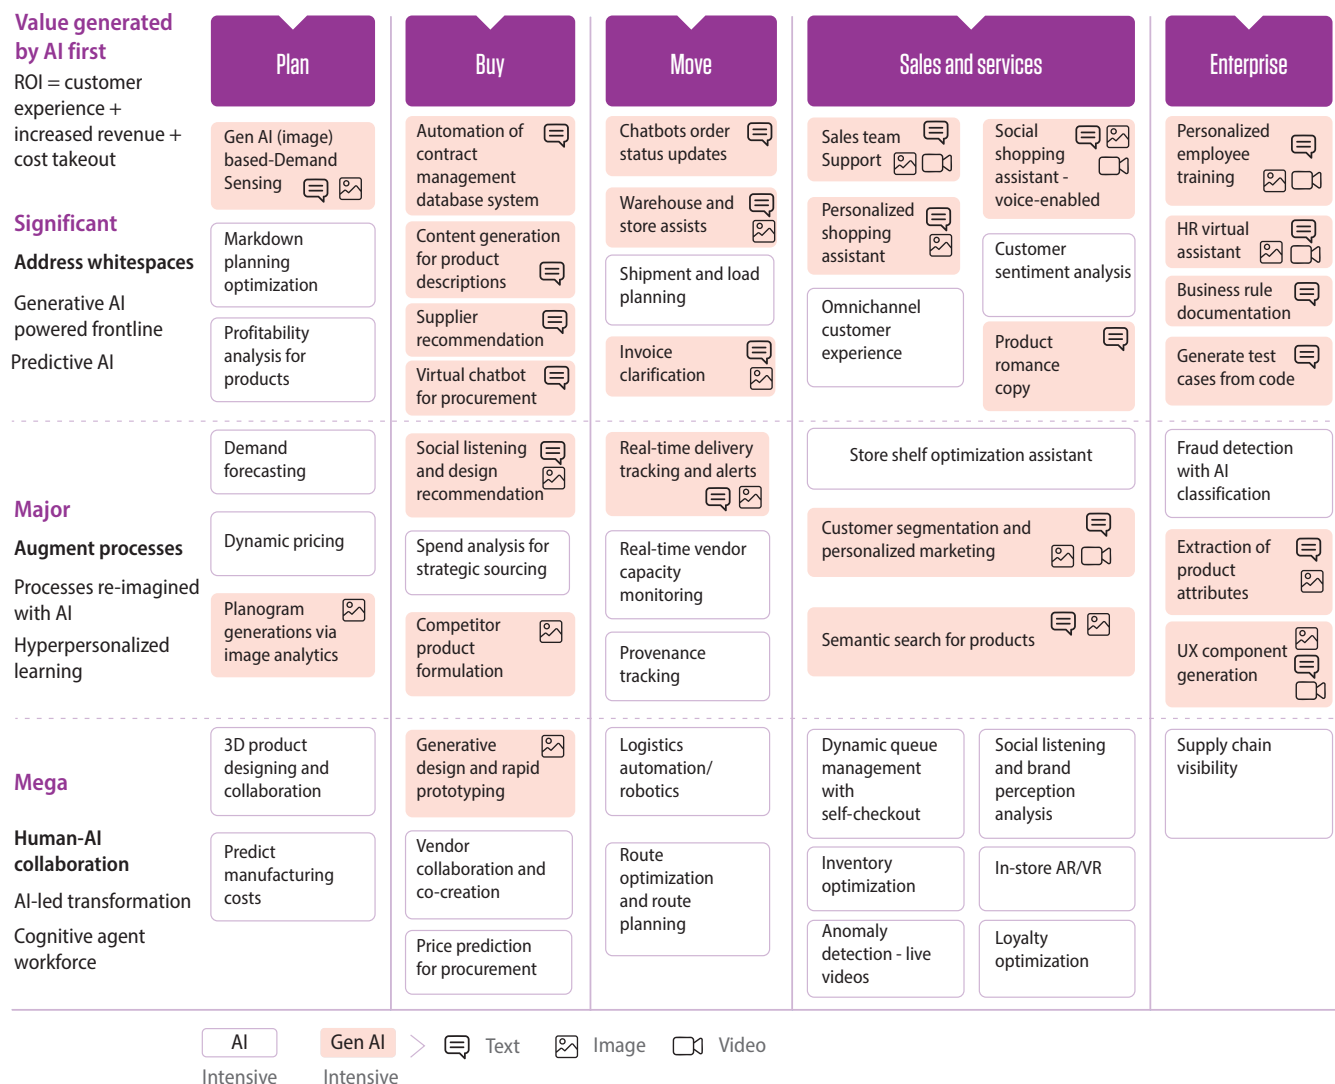
also looking to reimagine roles, enhance workforce skills and capabilities, and increase customer experience scores through 24/7 advisors — AI agents that do the heavy lifting so service professionals can focus on empathetic and strategic customer conversations.

The future of retail is a shift from AI augmentation to humans managing an increasingly cognitive agentic workforce (Figure 1). To advance towards this future, organizations need to unlock data from various systems. Use cases like logistics

automation, loyalty optimization, and more transparent supply chains can then take advantage of AI agents — autonomous software driven by goal-seeking algorithms. They will also need to make partnerships a part of their AI strategy, fostering relationships with tech players in the direct-to-consumer (D2C), business-to-consumer (B2C), business-to-business (B2B), and quick commerce arena.

In this chapter, we look at one of these agentic AI use case's, inventory optimization, as we've found that tackling this problem sets up retailers to advance quickly with other agentic solutions (Figure 1). However, there's a caveat. Retailers must also transform their operating models and their talent at the same time, as we discuss in our [AI Business Value Radar 2025](#) report.

Figure 1. Infosys AI blueprint for retail organizations



Source: Infosys

The need for discrepancy management

Modern global retail is a complex and challenging environment. Today's leading retailers must fulfill orders across a fragmented set of channels, systems, and partners.

Retail sales now occur through online and physical stores, each of which has its own systems and processes that are further fragmented by region. On top of this, they need to serve a partner network that is also segmented by online and physical stores, with multiple systems and processes. Straining things further, these systems are often incompatible with each other, with older, legacy systems struggling to share data with newer cloud-based or omnichannel solutions.

Globally, e-commerce retail sales [only account for one-fifth](#) of sales — yet [three-fifths of these sales](#) are estimated to come through online marketplaces, a challenging environment where retailers often face price wars, reduced discoverability, and downward pressure on margins.

Further complicating the situation, shoppers increasingly demand to have orders fulfilled through channels different from where the sales occurred — such as buying online and picking up in store, or buying in store and having the product delivered to their home. With nearly as many permutations, returns reverse the process and can be a significant drain on resources, particularly in retail sectors such as fashion and apparel. The annual cost of returns reached [\\$890 billion in 2024, with](#)

[16.9% of annual sales returned](#), and this is only set to increase unless retailers can meet customer expectations for convenience, while reducing fraud and increasing accuracy.

Balancing these challenges is essential to remain competitive in a market that values speed, accuracy, and low prices. Customers now expect an Amazon-like buying experience with product delivery dates and final prices displayed early in the order cycle. Previously, a retailer could present the delivery timing and price at checkout without worrying about how these factors might affect sales.

Managing all these elements requires an effective and transparent inventory system that provides accurate visibility into each product's availability and location in real time.



The risk of not getting it right

Large retailers generally build well-integrated supply chain inventory systems to match their online, physical, and partner networks. Yet it is inevitable that discrepancies and mismatches arise among these channels and their underlying systems and processes.

These discrepancies can result in oversell, when a product is promised and sold for a particular delivery time but is not fulfilled within that time range, or — in the worst cases — not fulfilled at all. In other instances, undersell happens when a product is sitting in a store or warehouse with no record of its existence.

Even if a retailer manages to get its discrepancy rate as low as 1%, this can still be a problem that hurts the bottom line. A small proportion of disappointed customers can have ripple effects on the brand and even put the company at risk in critical marketplaces.

Globally, **69% of shoppers** say their perception of a brand would worsen if an item were out of stock in-store after being shown online to be available. In addition, **(70%) are likely to shop** at a competitor to find an item that is out of stock at their preferred brand's store. Retailers risk having shoppers switch to another brand or platform and never return.

A small number of poor reviews online can have a significant effect on a product's sales. But in important markets such as China, marketplaces can place high penalties on — **and in some cases, even ban** — retailers for



cancelling orders or having fulfillment gaps.

What all this amounts to is a significant mismatch between promise and provision. In 2023, the **total cost of inventory distortion** was projected at \$1.77 trillion, a phenomenally large number (though down \$172 billion from 2022, with out-of-stocks accounting for \$1.2 trillion and overstocks totaling \$562 billion). Out-of-stock items accounted for 68% of the total cost of inventory distortion, and overstocks accounted for the remaining 32%.



Fix discrepancies the agentic way

Of course, automated solutions for solving discrepancies have been around for a while.

Inventory management software detects, flags and even auto-corrects anomalies when stock movements deviate from expected patterns, and radio frequency identification scanning of all stock movements has helped improve real-time accuracy and traceability. But given the nature of the underlying systems and incompatibility with cloud-based platforms, **very few** have been able to use this inventory management software effectively.

Indeed, despite the high stakes, tracking and correcting inventory discrepancies is still a manual process. It requires staff to check data in multiple applications, such as order management systems, website, and inventory visibility systems, to locate discrepancies, determine root causes, and then take corrective actions.

This process distracts supply chain teams from their core mission of operating and continuously improving the existing system, rather than sweeping up after its mistakes.

Enter agentic AI — continuously learning, autonomous, goal-driven, and adaptable software that can directly solve complex problems. An agentic AI-based solution can monitor discrepancies, identify reasons for failure, and invoke application programming interfaces (APIs) in different applications to correct issues automatically. Over time, the system learns to anticipate discrepancies and correct them.

A multi-agent solution

Infosys' work with retail clients has found that agentic AI-based solutions are effective at identifying inventory discrepancies. Once problems are detected and understood, agentic systems can then execute corrective actions.

Unlike standard generative AI solutions that require a command as its input ("create text" or "translate text"), the agentic AI solution starts with a goal. It identifies the tasks needed to achieve the goal and then executes them autonomously. By leveraging short-term and long-term memory, the system improves its decision-making as it gains more experience.

For this solution, "resolving inventory discrepancy" is the goal. The system splits the goal into executable tasks and identifies the relevant and available tools and skills needed for each.

Multiple agents are created as part of this solution, including, but not limited to, the following:

- **Available to promise (ATP) monitor agent:** Tracks ATP messages regularly being published to the channels. ATP represents the quantity of products that a retailer can confidently commit to deliver to customers based on the current stock and expected supply.
- **Demand verifier agent:** Compares the demand across systems.
- **Supply verifier agent:** Compares the supply across systems.

- **Decision agent:** Decides if supply or demand needs to be updated in any systems.
- **Mismatch resolver agent:** This agent updates the discrepancy between the recorded inventory data and the actual physical stock on hand, as required.
- **Orchestrator agent:** This meta-agent summons the other agents to connect to various systems, such as order management systems and the inventory visibility systems, and compares the values for the given stock-keeping units (SKUs).

In this architecture, task-level agents such as the demand verifier agent or decision agent activate APIs or connect to the database, to fetch the required details, with the orchestrator agent acting as the cognitive center.

When discrepancies occur, the solution identifies which system needs to be corrected based on history and error logs, and then resolves the issue. The system that Infosys clients have found most beneficial also learns the significance of quantity differences based on product type. For example, a one unit difference in footwear is more significant than a one unit difference in socks.

For added business value, the solution also includes a dashboard, providing a view of the number of discrepancies identified and the types of actions taken. This supplies business teams with high-level telemetry that they can use in further sales conversations.

This solution, the architecture of which is described in Figure 2, can also have a human

in the loop, with final recommendations executed by support engineers.

More bang for their buck


Agentic AI is now being applied across a wide range of business processes, as noted in Figure 1. What matters the most is its effective use, rather than treating it as a panacea for all business problems.

Organizations should focus on using this technology to mimic repetitive human processes, even as it also demands discernment, a dynamic problem to solve, and learning. An effective process requires cognitive decision-making after each step to

identify the next step. A useful rule of thumb is to apply AI agents to problems that are too big to ignore but not large enough to justify hiring more full-time employees. Reducing inventory distortion falls into both these buckets.

Indeed, inventory discrepancy management is a common and persistent problem for all retailers and requires huge manual effort to correct. The proposed agentic solution significantly reduces manual effort, while ensuring a highly effective and transparent inventory system offering visibility into a specific product's availability and real-time location. One client found that these agentic AI systems led to at least a 50% reduction in

Figure 2. High-level view of inventory discrepancy agents and orchestrator



Source: Infosys

inventory-related order cancellations. The solution also significantly reduces the effort for the support team, which was correcting more than 1,000 SKUs inventory differences each week.

While this solution is designed specifically for inventory discrepancy management, similar solutions can be built for other omnichannel CRL use cases, including managing stuck orders, payment processing, and even fraud management.

Of course, given that these agents are accessing important enterprise commercial off-the-shelf systems and databases, having a good level of human in the loop is crucial. We recommend solutions that enable robust access control, security, and authentication, with the final decisions as well as some intermediate steps still executed by trained inventory management specialists.

With global inventory distortion totaling a trillion dollars annually, agentic AI provides a significant bang for the buck. Our [Business Value Radar 2025](#) concluded that agentic AI should be at the center of any retail organization's enterprise AI transformation, delivering high value at low cost. In time, the number of disappointed buyers will trend toward zero, and the ripple effects on brand reputation will only go north — taking the whole enterprise with it.



How to manage discretionary risk in financial services

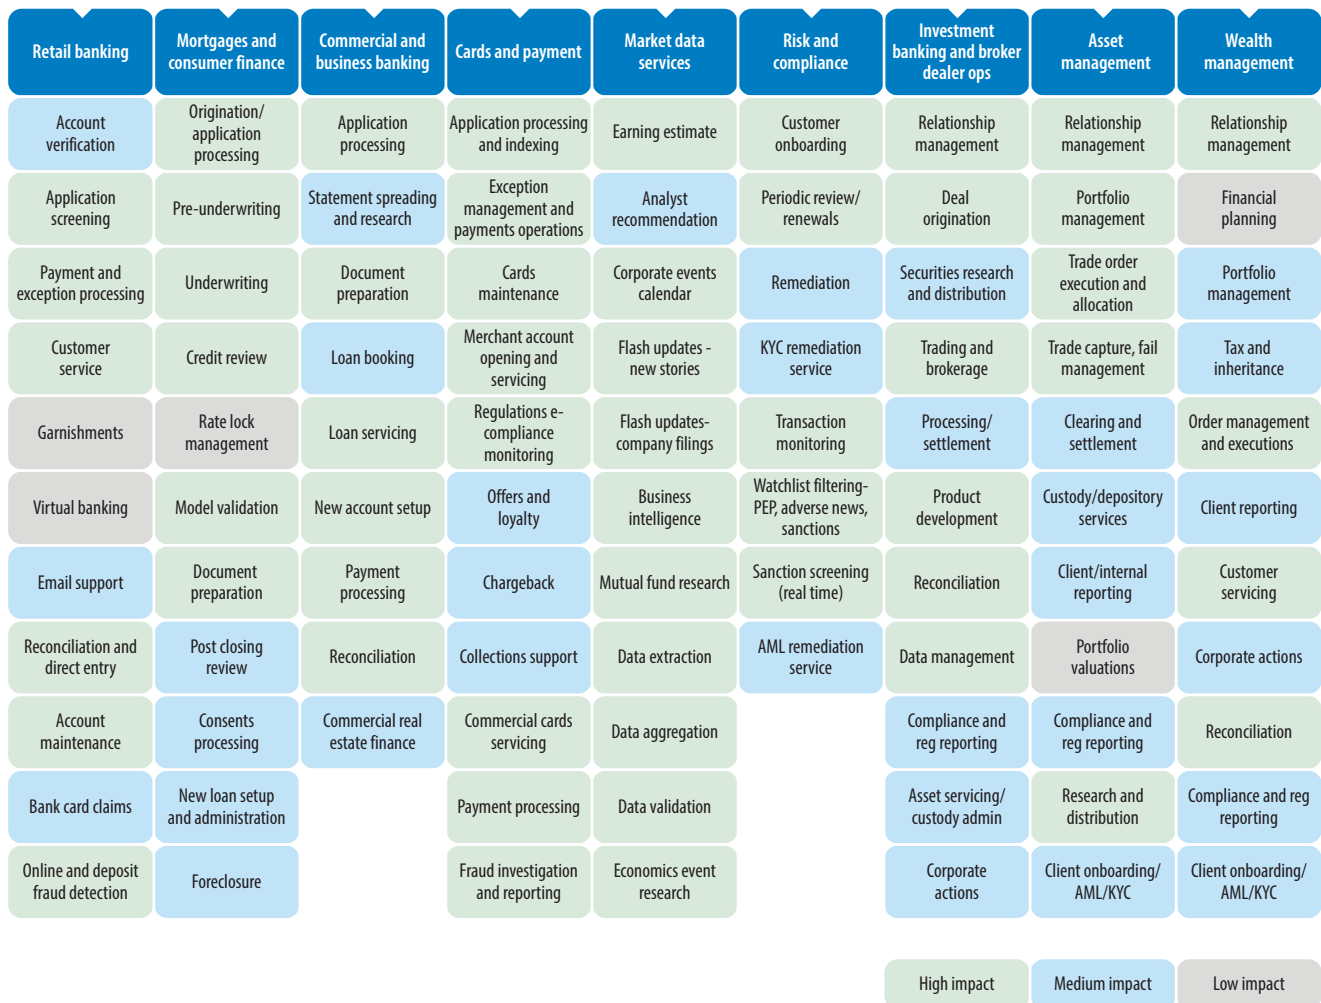
Managing risk effectively is essential for financial institutions, although bias and inconsistent interpretations of data can result in unreliable assessments and suboptimal decisions. Agentic AI systems, with their real-time access to a range of data and market intelligence, can drive more accurate risk assessments.

Banks are now placing innovation and growth at the top of the C-suite agenda, according to new research by the Infosys Knowledge Institute. Financial institutions are scaling artificial intelligence (AI), and looking carefully at their data estate and change management processes. In their efforts to become AI-first organizations, companies expect agentic AI — autonomous systems that can adapt to subtle changes in business environments — to take control of more core processes. High impact use cases across the banking value chain include application screening in retail banking, payment processing in cards and

payments, and relationship management in investment banking (Figure 1).

In addition, banks count on AI to help manage risk, which is at the heart of all financial services products. It's critical not only for growth and competitive differentiation but also for responsible and safe governance, whether it's a \$200 consumer loan or a multimillion-dollar global trade. Agentic AI can help by automatically diagnosing exposures, adapting controls, balancing trade portfolios, and simulating high stake scenarios — areas where there is a high level

Figure 1. Infosys AI blueprint for financial services



Source: Infosys

of subjective judgment and discretionary decision-making.

Risk is fundamentally subjective and discretionary. It's subjective because each institution can decide what amount of risk it wishes to take in a certain market or with certain products. It's also discretionary because even with clear risk guidelines, there are always customers, trades, or situations where a more refined understanding of the risk is needed to approve the transaction.

In this chapter, we look at the use of agentic AI in managing discretionary risk at global banks, areas where Infosys has already made a big leap with clients and partners.

The trouble with discretionary risk

A simple consumer loan highlights many of the complexities and challenges of discretionary risk. Credit assessments for loans to salaried employees are relatively structured and rules-driven. However, many customers

are businesses in the micro-, small, and medium-sized enterprises (MSMEs) segment, which feature a high level of unstructured data and not enough documentation to accurately assess creditworthiness.

Such businesses often secure lines of credit from informal, noninstitutional lenders at high interest rates. In most of these cases, customers have both the capacity and willingness to repay, thereby leaving a big opportunity untapped for other lenders. Depending on the appetite for risk, the underwriting officer might subjectively evaluate each loan on a case-by-case basis, often using proxies like goods and services tax statements or expense statements for assessing turnover or cash flow.

In this instance, the decision isn't policy-driven; it relies on both subjective and discretionary decision-making, necessitating more complete information before making a definite decision.

Similar difficulties are present in real-time trading when there are novel market or trade conditions. In these cases, the change in market conditions breaks the guidelines for risk. Manual intervention is needed to determine whether the trade should continue or whether the overarching trade strategy should change based on new information. Discretionary trading, as it's called, relies on the trader's judgment and intuition, and the evaluation of fast-evolving data. During market volatility, decisions about whether to pull out of a stock or invest more are subjective and can override preexisting rules.

The benefit of this scenario is that a trader can take advantage of stocks that have been written off due to profit warnings, with little upside priced in. This can offer a highly attractive risk/reward ratio. On the downside, traders bring their own emotions to the game; greed, fear, and overconfidence can lead to poor decision-making. Further, this sort of trading is excessively time-consuming, since even experienced traders must be aware of the state of the market and search for other factors that affect prices.

The benefits of using agentic AI

As noted, agentic AI can help financial services organizations in many areas, from personalized financial advisory to automated customer service. However, it is particularly suited to discretionary and subjective risk management, as in the previous two examples. AI agents work well in these scenarios by navigating large amounts of unstructured, natural language data and interpreting it appropriately based on real-time changes in the agent's environment. Other types of AI or automation struggle with these scenarios because such systems are typically rule-based or rely on predefined models and lack the autonomy and adaptive reasoning needed in highly context-dependent situations.

Below, we highlight three client examples where AI agents are already at work in discretionary risk management: regulatory and business control compliance, loan underwriting, and stock trading. In each case, agentic AI turns what were manual processes that lack discretionary insight into processes



that use diverse, real-time data to make decisions with a strong understanding of the actual risk.

Case 1: Compliance orchestration

Banks manage thousands of regulatory updates and compliance audits annually. Many of these could [cause reputational harm and severely damage](#) bank operations if not addressed effectively.

However, manual compliance processes are error-prone, costly, and unable to keep pace with real-time regulatory changes and updates in business controls — the structured policies and processes designed to manage risk.

Agentic AI acts as a centralized compliance orchestrator, interpreting changes in business controls alongside regulatory updates. It works to update internal policy references and run audits of change transactions. Agentic AI can also measure compliance levels across teams against prevailing regulations and business controls and then suggest changes and modifications to the required leads for approval.

One example is JPMorgan's LAW (legal [agentic](#) workflows), which automates and streamlines the review and analysis of complex credit agreements, contracts, and other legal documents. LAW uses agentic AI to interpret, classify, and extract information from vast numbers of documents that would typically require extensive discretionary and manual review by legal and compliance teams. The system [reviews commercial](#)

loan agreements in seconds instead of the several hours, or even multiple days, it took previously.

A similar solution that Infosys created for its clients is a business control validation system. Banks increasingly need to validate complex business controls across operational, financial, regulatory, and risk management domains — and do it with speed and discretion. The complexity and interconnectedness of risks with accelerated digital transformation, and the sheer pace of organizational change, have made this prescient. This validation exercise ensures that business controls are properly designed, consistently implemented, and effective at achieving their risk mitigation objectives. The Infosys solution deploys specialized autonomous agents that collaborate across the business control validation life cycle.

The control discovery agent continuously scans internal systems, policies, and regulatory frameworks to identify all applicable business controls, automatically mapping dependencies and relationships. The validation orchestration agent acts as the central coordinator, intelligently prioritizing control testing based on risk scores, regulatory deadlines, and business criticality while managing resource allocation across validation activities.

Working in tandem with these agents, the evidence collection agent gathers supporting documentation from multiple sources, including ERP systems, transaction databases, audit logs, and external regulatory feeds. This agent employs advanced pattern



recognition to identify relevant evidence and cross-reference data points across systems. Finally, the risk assessment agent continuously evaluates control effectiveness, using machine learning algorithms to detect anomalies, predict potential control failures, and assign dynamic risk scores based on real-time operational data.

The ability of this autonomous system to make good decisions comes from the solution's integration of diverse data sources. Internal systems provide core banking transaction data, customer information, risk management metrics, and audit trails, while regulatory feeds include updates from central banks, financial authorities, and international standards bodies.

According to work with our clients, the benefits of this approach over traditional methods are an increase in automation of manual processes of between 15% and 20%, and a 40% reduction in processing effort through real-time validation and automated reconciliation. Further, regulatory control response times are 30% faster, and predictive analytics reduce compliance violations by 25%.

The solution covers all control categories, from IT general controls to compliance controls, ensuring failsafe operations with holistic coverage.

Case 2: Agentic AI in loan underwriting

Customers without regular salaries — often with seasonal income and profits — typically lack structured documentation to assess their

creditworthiness and repayment capacity. In this situation, the underwriting process for approving loans and credit lines is purely subjective.


In the automotive industry, one option is an agentic AI system known as proactive personalized product offering (P3O), which transforms how lenders extend unsecured car loans. P3O uses rich transactional data to replace what was once a subjective, discretionary risk-based decision (Figure 2).

Imagine if a customer spent about \$400 each month on ride-hailing services like Uber consistently for 18 months. This shows a clear, regular spending pattern. This individual also boasts an impeccable record of zero credit card delinquencies and a steady income from their small business, even though the business itself has seasonal revenue dips and peaks.

Agentic AI processes these data points, recognizing the consistent monthly spending as a proven capacity to repay a fixed amount. The flawless credit history further signals a high willingness to pay. By connecting these patterns, the AI understands the customer's financial discipline and ability to manage consistent payments, even with the business's income fluctuations.

This allows agentic AI to autonomously generate a personalized, preapproved, unsecured car loan offer with an equated monthly installment that directly aligns with the customer's demonstrated spending habits. This approach uses real-world behavioral data to provide accessible, tailored

Figure 2. P3O delivers data-driven, personalized financial products



Source: Infosys

credit, moving beyond the limitations of traditional lending.

Another example is when an MSME business experiences seasonal revenue swings but maintains a predictable, upward trend in its annual income. Banks can use the P3O agentic system here too.

Consider a small business crafting woolen winter wear. Its sales naturally spike in one quarter, dip in another, and normalize for the rest of the year. Despite these fluctuations, its yearly growth remains steady. Traditionally, securing extra working capital during its busy season has been tough. The increased demand, while a sign of growth, might cause it to breach its conventionally allocated credit limits, as traditional lenders often rely on rigid credit checks.

This is where agentic AI proves valuable.

It can analyze extensive past transactional data, identifying consistent revenue streams and upward trends. It then combines this information with publicly available data, such as weather patterns. If winter arrives sooner or later than usual, agentic AI can spot this change. By understanding these subtle dynamics, the solution can offer additional working capital loans when the business needs them the most, helping it maximize peak demand and foster growth, even when it means temporarily exceeding a traditional credit threshold.

Traditional credit models often fall short due to their strict dependence on collateral, extensive paperwork, and inflexible repayment schedules that penalize fluctuating incomes. In contrast, agentic AI can create a repayment plan that matches the business's revenue cycle, allowing for larger payments during high-revenue periods

and smaller ones during off-peak seasons. This adaptable approach empowers MSMEs to manage their finances more effectively. In our implementations with various clients, we've found this solution increases share-of-wallet by between 20% and 25%, increases credit approval rate by 40%, and reduces risk significantly, with a reduction in probability of default of as much as 20%.

Case 3: Trading and investment

Investment companies continuously analyze vast amounts of market data across diverse asset classes to optimize portfolio performance.

However, human traders cannot process the volume and velocity of market data in real-time, while traditional algorithmic trading lacks adaptability to unexpected market conditions.

Agentic AI autonomously monitors markets 24/7, detects nonobvious correlations, executes multistep trading strategies, and dynamically adjusts portfolios based on changing conditions without constant human oversight.

Of course, financial institutions should implement agentic AI trading systems with robust governance frameworks that balance this level of autonomy with appropriate risk controls and regulatory compliance. This governance must establish clear boundaries for when AI is needed and when it isn't, including stock position limits and exposure thresholds. Regulatory oversight mechanisms are needed too. Adherence to evolving

regulations like the Markets in Financial Instruments Directive II (MiFID II), the Financial Industry Regulatory Authority (FINRA), and emerging AI-specific requirements is essential, while also maintaining transparent audit trails for regulatory examination.

Given the nature of autonomous trading, controls should be implemented for excessive pre-trade risks, market disruptions, and human-in-the-loop escalation.

This multilayered approach ensures that while agentic AI systems operate with necessary speed and efficiency in market environments, they remain within acceptable risk parameters and regulatory boundaries that protect both institutional interests and market integrity.

One Infosys solution, which is currently in the proof-of-concept stage, uses autonomous market agents to monitor global market conditions, price movements, news sentiment, and alternative data sources. Strategy execution agents dynamically adjust portfolio positions, with the capability for intentional goal-driven strategy creation that anticipates and adapts to market conditions. Then risk-management agents assess counterparty risks, monitor position concentration, and execute pre-emptive hedging strategies when risk thresholds are approached.

Though not yet at production scale, this architecture can simultaneously process complex market scenarios across multiple asset classes, delivering speed, accuracy, and adaptability compared to traditional trading

approaches. In fact, in preliminary trials, we've found that the solution can reduce hedging costs by 10%, with an increase in alpha (return on investment compared to a relevant benchmark index) of as much as 100 basis points.

The agentic AI imperative

Agentic AI systems are driving high levels of interest, and banks should be alert to the opportunities. These agents can synthesize information quickly, monitor for emerging risks, adapt to global markets, and incorporate real-time intelligence into underlying AI models.

To get ahead and make the most of this new technology in discretionary risk management, banks should take three steps:

1. Use APIs to integrate legacy systems

with AI agents: Existing infrastructure remains critical for core operations, and a complete system replacement is usually not feasible. System APIs act as a secure bridge to legacy technology, while process APIs orchestrate workflows, and experience APIs deliver user-friendly interfaces. Gartner projects that agentic AI systems will face a 40% failure rate by 2027, mostly due to the complexity of legacy systems rather than flaws in the AI technology.

2. Ensure all decisions made by agents

are transparent and explainable: This is critical given the complexity of some black-box systems, where agents set their own goals and execute multistep strategies. Often, their actions are difficult

to audit. In addition to reputational harm, financial institutions face mounting regulatory pressure from the General Data Protection Regulation (GDPR), the US Equal Opportunity Act, and Australia's Privacy Act (scheduled for December 2026) to act more transparently. Lenders are sometimes fined when they make decisions without a sufficient explanation.


3. Prioritize training: Compliance in the agentic era requires entirely new skill sets. Compliance and risk teams must develop the talent for continuous and adaptive oversight of agentic systems. Training must address tiered decision-making frameworks, where different autonomy levels require different levels of oversight, from automated monitoring to full and disclosed human intervention.

Doing all three things well is required to ensure responsible and safe governance, and these steps act as a foundation for many other agentic AI use cases that require discretion, speed, and real-time decision-making.

The result is significant mitigation of human-made bias, increased ability to spot subtle correlations, and generation of dynamic context-aware risk assessments that are more accurate than traditional static or rule-based methods. As we move into 2026, and an age of increased global market tension, banks will need to adapt with more dynamic, tech-enabled, and resilient strategies. Agentic AI may just rise to be one of the most important tools in the strategic arsenal — depended on by banking executives and consumers alike.



Agentic AI is redefining the manufacturing value chain



Manufacturing has long embraced digital transformation, but previous advances have been limited to discrete pockets, rather than integrated across the entire business. Agentic artificial intelligence promises to shift this dynamic, ushering in an age of autonomous, goal-driven operations across the enterprise.


The manufacturing of products is becoming more connected, intelligent, and autonomous. At the heart of this evolution is the product value chain — each stage contributing to competitive advantages, cost efficiency, and differentiation.

Artificial intelligence (AI) has played a critical role, evolving from knowledge-based engineering to machine learning and deep learning. More recently, generative AI has helped make sense of product data dispersed among multiple unstructured data stores and document management systems.

Despite its benefits, this digital transformation has only solved discrete, isolated problems in areas such as predictive maintenance, operations, and smart services (Figure 1).

Previous AI systems couldn't drive efficiency across the value chain and collapse entire processes to create competitive advantages. For this to happen, various engineering, operational, and information technologies would have to be integrated. This level of orchestration hasn't been possible — until now.

Figure 1. Blueprint for AI-first physical product engineering



Source: Infosys

Agentic AI can independently make decisions, learn from its environment, and take actions to achieve defined goals. It can orchestrate end-to-end processes across engineering, manufacturing, and services, increasing efficiency (Figure 2).

Three manufacturing use cases

How can agentic AI enable autonomy across engineering, production, and supply chain processes? Four elements make this emerging technology ideal for the job.

1. Modularity: Processes can be broken down, and agents can be tailored to specific tasks.


2. Interoperability: The technology can integrate with various manufacturing execution systems, enterprise resource planning (ERP) systems, and product life cycle management (PLM) systems.

3. Human in the loop: Human oversight is built into the system to ensure strict oversight and control.

4. Scalability: Companies can easily shift implementations from pilots to full ecosystems.

With these features in place, agentic AI marks a shift from isolated AI-enabled point solutions and automation to

Figure 2. Evolution of AI in the product value chain



Source: Infosys

full-scale manufacturing process reengineering. This enables real-time decision-making and adaptation, self-improvement across domains, and significant cost savings as AI scales. In this chapter we consider three use cases for agentic AI in manufacturing that we are currently working on with clients.

These use cases have been selected as they use agentic AI to traverse multiple steps in the manufacturing value chain (Figure 1), reengineering processes along the way.

Case 1: Aircraft cost optimization

Aircraft manufacturers face growing pressure to reduce product costs while managing complex configurations of thousands of parts and balancing frequent design changes. Their primary goal is to identify the right price with the right supplier at the right time.

However, cost estimation, sourcing, and part-grouping (organizing parts with similar characteristics) are manual, siloed, and reliant on spreadsheets. Supplier evaluation is inconsistent and rarely informed by performance data.

Workflows lack real-time insights and cross-functional integration, making it difficult to reuse past knowledge or anticipate cost drivers — such as labor hours or production volumes — early in the process. As a result, cost risks are identified too late, limiting early cost-control opportunities and leading to inefficient decisions in fast-paced programs.

A key component of cost optimization is “should-costing” — estimating what a product should cost under best-practice conditions. This involves analysing each element of the cost, including raw materials, production processes, batch sizes, and

tooling, and then calculating the expected cost based on efficient manufacturing processes. However, these analyses often lack precision and consistency across teams, leading to nonconformance risks, inflated costs, and reactive procurement decisions.

A hierarchical agentic AI system integrates four intelligent modules: part grouping, supplier intelligence, should-costing, and sourcing optimization. Specialized autonomous agents within these modules collaborate using orchestration and human-in-the-loop layers to drive adaptive, real-time, cost-aware decision-making across the life cycle (Figure 3).

A critical facet of this architecture is the aircraft product cost optimization orchestrator. This system acts as a central, AI-powered controller that integrates the entire cost estimation and sourcing ecosystem. It operates across the following three coordinated layers:

- 1. Data integration:** Connects to enterprise systems, including computer-aided design (CAD), PLM, ERP, and procurement platforms, to aggregate structured and unstructured data. This provides real-time access to design, supplier, and cost information.
- 2. Intelligence coordination:** Continuously monitors life cycle events, such as design changes or sourcing requests, and dynamically activates the appropriate agents in the correct sequence. It also manages conflicts, aggregates outputs, and ensures end-to-end workflow integrity.




3. Human in the loop: Delivers actionable insights to engineers, buyers, and cost analysts, allowing them to review and interact with AI-generated recommendations. Users can provide feedback or override decisions, maintaining human control and judgment. Over time, this feedback loop helps refine the orchestration model through continuous learning and reinforcement.

The orchestrator also breaks down silos by coordinating decisions across teams. For instance, when a design modification is made, the orchestrator initially calls the part

grouping agent to find similar components that were used previously. This enables the reuse of historical cost logic (grounding asset valuation in past transaction prices) and sourcing decisions, a capability often lacking in traditional workflows. Next, it engages the should-costing agent, which uses updated design parameters and supplier data to generate cost estimates that reflect real-world manufacturing conditions. Finally, sourcing optimization and supplier intelligence agents re-evaluate preferred suppliers based on cost, delivery lead time, and past performance, which improves negotiations and reduces procurement risks (Figure 4).

Figure 3. Architecture for agentic AI-driven aircraft product cost optimization



Source: Infosys

This AI system can be deployed across engineering, sourcing, and manufacturing. In addition, it integrates with CAD, PLM, and procurement tools.

Infosys clients report up to 50% cost savings by using this system to identify cost risks early; improve should-cost accuracy; reuse past cost logic and supplier intelligence; and enable better-informed and timely supplier negotiations. The system also delivers a 20% to 30% cycle-time reduction from design iteration to sourcing decisions, improved alignment across engineering, sourcing, and procurement, and higher negotiation leverage through data-backed, real-time cost targets.

Case 2: Smart asset maintenance

Maintenance operations face pressure to meet increasingly stringent customer service level agreements. Achieving this requires predictive insights into equipment issues, minimal repair and downtime, and real-time coordination across data platforms. Without these capabilities, manufacturers struggle to optimize asset usage and prevent repeat failures on the shop floor.

Traditional reliability and maintenance (RMT) processes in manufacturing are often reactive and fragmented and rely heavily on manual coordination. Equipment is generally repaired only after failure, which leads to unplanned downtime. Maintenance planners manually align technician schedules, asset criticality, and inventory. Further, data from sensors, computerized maintenance management systems, ERP systems, and supervisory control

and data acquisition systems are siloed and poorly integrated.

Predictive capabilities are further limited by underused historical data and the loss of practitioner knowledge. Safety and audit documentation are often incomplete or delayed.

Agentic AI introduces a network of agents that can collaborate to manage RMT workflows — with minimal human intervention and supported by a conversational chatbot interface. Key capabilities include:

- Autonomous task planning and execution.
- Predictive maintenance and resource optimization.
- Seamless integration of enterprise systems through user-friendly interfaces.
- Conversational interfaces for real-time operator support.
- Continuous learning and adaptation.

Agents use real-time data from internet of things sensors and platforms like SAP Plant Maintenance, Azure Databricks, and Cognite Data Fusion, to monitor, plan, and act on maintenance needs dynamically. The agents form an interconnected network, where each specializes in a critical aspect of the RMT workflow, collaborating seamlessly to create a smart, adaptive maintenance ecosystem.

For example, the predictive maintenance agent analyzes real-time sensor data to forecast potential equipment failures. Once a risk is detected, a work order is created. The canvas agent defines the required steps for

Figure 4. Key functions of aircraft product cost optimization agents

Agent name	Key functions and actions
Part-grouping module	
Part grouping goal agent	Defines the objective of part clustering; orchestrates subtasks such as metadata extraction, embedding, and clustering coordination
Metadata extraction agent	Extracts and standardizes structured metadata (e.g., material, process, weight) from raw parts records
Image embedding agent	Converts part images (2D or 3D) into semantic vector embeddings using deep learning for visual similarity
Similarity search agent	Performs multimodal similarity search, combining metadata and embedded image for semantic part grouping
Human review agent	Allows domain experts to review or refine auto-clustered part groups before passing to supplier analysis
Supplier intelligence module	
Supplier intelligence goal agent	Initiates supplier evaluation; manages supplier data aggregation, scoring, compatibility checks, and ranking
Supplier scoring agent	Scores suppliers based on KPIs like quality, lead time, past performance, and geographic presence
Supplier compatibility agent	Assesses technical compatibility between suppliers' capabilities and specific parts requirements
Supplier quotation analysis agent	Analyzes supplier quotes, detects pricing anomalies using historical benchmarks, simulates sourcing alternatives
Human approval agent	Allows procurement experts to validate or override supplier rankings and flag potential risks before finalizing recommendations
Should-costing module	
Should-costing agent	Oversees cost estimation; manages the two-stage semantic similarity process and historical price interpolation
Cost forecasting agent	Forecasts cost and schedules impacts of late-stage design changes, simulates bills-of-materials variations, and recommends cost-effective alternatives to avoid overruns
Similarity search agent	Identifies semantically similar parts based on metadata and image embeddings from historical data
Cost interpolation agent	Calculates estimated part costs by interpolating prices from closest historical analogs
Cost approval agent	Enables finance or procurement experts to review, validate, or override AI-generated cost estimates
Sourcing optimization module	
Sourcing optimization goal agent	Integrates inputs from supplier intelligence and should-costing agents; evaluates trade-offs such as cost, quality, and lead time to recommend the optimal supplier and pricing for each part
Sourcing decision review	Presents the AI-generated sourcing recommendations to human procurement experts or sourcing managers for review, approval, or override; ensures final decisions incorporate human judgment and risk assessment
Feedback and learning agent	Collects performance feedback from procurement outcomes and human experts; updates agent models and parameters to improve future decision accuracy

Source: Infosys

the solution in the work order. A scheduler agent picks up work orders and dynamically adjusts technician schedules to prioritize urgent repairs. Simultaneously, the inventory agent ensures the required spare parts are available by monitoring stock levels and triggering procurement, if needed.

Then, the safety agent reviews compliance and audit documentation to maintain regulatory standards, while the incident response agent coordinates responses when unexpected events occur. The feedback agent gathers insights from technicians and operators to refine decision-making, and the budget agent tracks and manages costs in real time to optimize resource allocation.


In addition, a document agent handles the retrieval and updating of maintenance

records, ensuring that information is accurate and accessible. The chatbot agent — powered by a large language model (LLM) — is the interface that provides intuitive conversational support to users (Figure 5).

These agents communicate through a centralized platform, enabling automated coordination across the RMT process. Together, they reduce manual effort, streamline operations, and deliver predictive, data-driven decisions — making maintenance more intelligent, efficient, and responsive (Figure 6).

Based on field deployments, we have seen significant improvements among the following key maintenance metrics, depending on system maturity and integration depth:

Figure 5. Key agents in RMT workflows



Source: Infosys

- Mean-time-to-repair was reduced by 25% to 30%.
- Equipment uptime and availability improved by 30% to 35%.
- Technician productivity increased by 30% to 32%.
- Operational efficiency increased by 28% to 35%.

These results are driven by intelligent automation, real-time alerts, and AI-driven support at the point of maintenance planning and execution.

Case 3: MRO decision-making support

Aircraft maintenance, repair, and overhaul (MRO) requires close coordination with

engineers, suppliers, regulators, and original equipment manufacturers (OEMs). Teams manage thousands of documents — from manuals to directives — and must ensure ongoing airworthiness compliance for more than 25,000 commercial aircraft in a global, highly regulated business environment. Minimizing aircraft-on-ground (AOG) time is critical to avoid disruptions and financial losses to airlines and their passengers.

Damage to aircraft can occur anywhere in the world and affect any part of the aircraft, including critical areas that are hard for people to access or inspect. Field damage reports are often inconsistent, with data coming from handwritten and annotated notes, low-quality images, misaligned drone

Figure 6. RMT agents and their capabilities

Agent	Capability
Predictive maintenance agent	Detects anomalies using real-time and historical data
Scheduler agent	Orchestrates maintenance tasks based on priority and availability
Inventory agent	Tracks parts, tools, and resource availability
Safety agent	Enforces maintenance safety and standards compliance
Incident response agent	Responds to emergencies
Feedback agent	Supports continuous learning and adaptation
Budget agent	Calculates and monitors maintenance costs
Canvas agent	Applies insights from historical notifications and work orders
Document agent	Extracts procedural data, instructions, and compliance criteria from manuals
LLM chatbot agent	Provides conversational interface for planners, supervisors, and technicians

Source: Infosys

photos, and incomplete descriptions. These unstructured inputs make accurate damage assessments difficult.

Engineers must interpret this data, stay current with evolving Federal Aviation Administration (FAA) and European Union Aviation Safety Agency (EASA) directives, and make complex repair-or-replace decisions. At the same time, they must coordinate part sourcing, logistics, and work across fragmented systems — leading to delays, errors, and high AOG times.

A solution built on multi-agentic architecture automates key steps — including inspection, damage classification, repair-or-replace decisions, repair planning or part sourcing, compliance validation, and coordination with OEMs and regulatory authorities.

An MRO central orchestrator agent coordinates and streamlines maintenance, repairs, and operations across systems and stakeholders. Specialized agents, including a damage detection agent (perception), document intelligence agent (retrieval), damage severity evaluator (reasoning), and MRO repair agent (action), and OEM coordination agent (planning) collaborate with compliance and supplier agents. Human-in-the-loop checkpoints ensure traceability, safety, and regulatory alignment (Figures 7 and 8).

Infosys estimates that a multi-agentic system can reduce engineering effort by 60% to 70% and cut AOG time, often by 10% to 25%, depending on airport layout, traffic levels, and weather. By integrating with



vision systems, document repositories, and enterprise tools, companies can achieve end-to-end automation with real-time collaboration and compliance assurance across global maintenance operations (Figure 7). This solution lays the foundation for more intelligent and future-ready MRO operations.


This system substantially reduces engineering effort and AOG timelines, while enhancing compliance through near real-time FAA and EASA validation using the most current data. Stakeholders are now coordinated, and the solution is both scalable and traceable — creating a transparent decision-making framework for safety-critical operations.

The agentic AI future in manufacturing

The rapid evolution of AI — particularly agentic AI — is fundamentally transforming how engineering, manufacturing, operations, and maintenance processes are executed.

By overcoming traditional barriers, such as infrastructure, siloed data, and manual workflows, these intelligent systems are unlocking new levels of precision, agility, and scale. Agentic AI is already delivering measurable impact for manufacturers: optimizing product costs, accelerating procurement decisions, and improving asset performance.

Figure 7. Multiagent AI architecture for MRO operations



Source: Infosys



Figure 8. Key agents in the MRO workflow

Agent	Role
MRO central orchestrator agent	Coordinates and streamlines maintenance, repairs, and operations across systems and stakeholders
Shared human-in-the-loop interface	Enables expert intervention at key decision points, including damage classification, escalations, and compliance approvals across systems
Damage detection agent	Identifies and localizes structural damage using visual and textual input; initiates severity assessment
Document intelligence agent	Extracts procedural data, repair limits, and compliance rules from structural repair manuals, aircraft maintenance manuals, and other sources
Damage severity evaluator	Determines repair path — whether in-house at MRO or escalated to OEM — based on damage type and regulatory criteria
Repair-or-replace evaluator	Determines the best action by comparing damage severity and limits against EASA and FAA guidelines and engine manuals to ensure compliance and safety
MRO repair agent	Executes approved repair actions within MRO capabilities, leveraging procedural guidance from documentation agents
OEM coordination agent	Manages complex escalations to OEMs, initiates external workflows, and tracks feedback and repair approvals
Predictive maintenance agent	Analyzes historical failures and sensor data to forecast future maintenance needs and reduce unplanned downtime
Supplier coordination agent	Optimizes part and tool procurement by analyzing availability, lead times, and cost from approved suppliers
Repair verification agent	Confirms that repairs have been correctly executed, flags inconsistencies, and prepares for compliance validation
Compliance agent	Validates final repair outcome against FAA and EASA directives and certifies airworthiness

Source: Infosys

What makes this transformation significant is not just automation, but the intelligence that augments human expertise rather than replacing it. Agentic AI enhances human capabilities through intelligent automation and continuous learning, enabling faster, more informed decisions and driving efficiencies that were previously out of reach.

As AI continues to mature, human-in-the-loop frameworks aligned with [responsible](#)

[AI principles](#) will ensure safe, traceable, and ethical deployment.

Agentic AI offers benefits throughout the manufacturing ecosystem — making it an indispensable tool for driving value, accelerating innovation, and reengineering manual, onerous business processes into responsive, autonomous, efficient workflows. What seems impossible today will become tomorrow's status quo.



How agentic AI is re-engineering telecom business sales



The telecom B2B sales process is strained, with lower sales conversion rates than other industries. Agentic AI can simulate process flows and act as a sales copilot, but only if products, systems, and interfaces are standardized.

In telecom, traditional AI, and generative AI more recently, is helping the business-to-business (B2B) market to grow at a [CAGR of 17.2%, reaching \\$178 billion by 2029](#).

AI is re-imagining customer journeys through generative AI-powered virtual agents and real-time call intent identification. It also makes products more customer-centric and competitive.

AI is improving the software development life cycle through synthetic data generation and code development copilots. In networking, AI is infusing automated decision-making

across the network lifecycle and solving slow response times through self-healing processes (Figure 1).

But it could be doing better, especially in sales, where new research by the Infosys Knowledge Institute found that telecom companies struggle to generate value from AI in sales-specific use cases.

In this chapter, we look at how agentic systems are now addressing this challenge by simulating many steps in the sales, quoting and ordering processes (first column in Figure 1). Instead of looking at simple point

Figure 1. Infosys AI blueprint for telecom

Sales, marketing	Product services	Customers	Networks
AI twins for marketers	Order fallout prediction and remediation	Multistep reasoning-driven assistant	Smart planning
Automated campaign generation	Enhanced service activation and order processing	Semantic search and conversational chat	Intelligent service design
Lead scoring	Reconciling gaps between different user guides	Voice-enabled digital assistants	Smart field operations
Cross selling and upselling	Innovating contract management	Generative AI as a customer for learning	Predictive assurance
Adaptive revenue targeting	Scenario modeling for portfolio evaluation	Call record summarization	Closed-loop assurance and self-heal
Contextual sales outreach	Modeling churn prediction	Omnichannel sentiment analysis	Intent generation
AI humans	Tailored product recommendations	Collaboration and CC analytics	Orchestration and autonomous networks
	Intelligent discount guidance	Cognitive and intelligent nudges	Generate synthetic data and network topology for testing
	Validating fraud detection and response	Conversational AI	Agent-based network operations
	Bill assurance	Agentic AI workflow automation	

Source: Infosys

AI Gen AI and agents

solutions, agentic AI works across sales processes, creating significant business impact by turning the sales funnel from a series of disjointed steps to a coordinated, adaptive system.

Business sales re-engineering

For decades, telecom companies have struggled to take sales leads and turn them into customer orders. Processes are often manual or semi-automated, crossing multiple legacy business support (BSS) and operational support systems (OSS).

In business-to-consumer (B2C) telecoms,

only 54% of calls are answered and just 36% of those calls lead to conversion. This is even worse in the B2B space, where sales call conversion rates for telecom companies are just 12.7%, compared with 20.5% in the insurance industry and 14.4% in financial services.

The B2B market is an important one for telecom providers, as margins can be higher than in the commoditized consumer market, with B2B reporting gross margins of 25% to 35%, whereas in B2C it drops to 5% to 15%. However, this comes hand in hand with more complex and tailored solutions for each client: B2B customers routinely ask for

industry-specific solutions such as secure, private 5G networks bundled with bespoke cloud-based unified communications, or flexible contracts which include dynamic bandwidth scaling.

Strained telecom B2B sales process

The challenges flow throughout the sales process. The problems can start with gathering requirements from a client lead by a sales agent. These are typically through a combination of human conversations, emails, or forms. The information gathered is often patchy, incomplete in some areas, and contradictory in others.

And none of these orders is straightforward. Typically, a B2B order will cover multiple client locations, each with different requirements and different levels of infrastructure readiness. Furthermore, the telecom company itself will also have different levels of capability and resources across geographies that need to be balanced and matched with client needs.

These issues can cause delays and confusion, slowing the sales process and impacting order fulfillment. This can negatively affect service level agreements (SLAs) and customer experience — resulting in penalty fees and in the worst case, the loss of the sale.

Compounding this is the challenge of order transparency. It can be very difficult for a telecom company to analyze a particular order and pull together all of its component parts in order to validate capabilities, pricing, and streamline or accelerate the process.

This is often due to the fact that many telecom companies have multiple overlapping legacy systems. For instance, confirming service availability in one specific region can be a convoluted process, involving the accessing of inventories in multiple OSS platforms.

Agentic AI in the sales process

What's needed is a digital foundation to operationalize the telecom B2B sales process so that sales leads are easy to grasp, and so that SLAs for fulfilling enterprise orders meet or even exceed customer expectations.

To help, the [TM Forum](#) has published APIs that enable agentic AI — intelligent, autonomous systems capable of interacting, reasoning, and performing complex tasks with minimal human intervention.

The APIs are part of the Open Digital Architecture (ODA) standards, and ensure interoperability, scalability, and modular integration of AI agents across the telecom business and network layers. The APIs allow AI agents to access and manage customer data and network telemetry across typically siloed enterprise COTS and cloud platforms.

The agents are able to simulate many steps in the sales, quoting and ordering processes, and can address the immediate pain points due to broken process flows — either because of a lack of order transparency in backend legacy systems, lack of infrastructure readiness, interoperability issues, or an inability for sales personnel to pull disparate data from multiple systems.

AI agents can help avoid manual errors in quoting and pricing, take away the pain of multiple handovers between departments for order management, and expedite order implementation and integration.


Simulation, where agents interact with virtual environments, also enables non-deterministic process flows such as pricing and negotiation possibilities in the quoting stage, or possible future sales requirements in the fulfillment stage. These simulations often involve agents independently decomposing tasks, invoking specialized tools, orchestrating workflows, and refining their actions based on simulated feedback, ensuring robust performance in dynamic, unpredictable situations in the real world.

Agentic copilots are then needed for life

cycle management across lead, opportunity, solution, quote, contract, and order (Figure 2). Copilots in this context operate strictly under human supervision in a predefined workflow, processing information, offering suggestions, and automating routine steps, all in a bid to maximize productivity.

For example, in the quote and negotiation stage, the agentic AI copilot would create, modify, and generate the quote automatically, while in the contract and order stage, the agentic AI system would be able to query the contract and create the order. The agent will also help in suggesting upgrades to the customer over a self-serve portal by validating new promotions and offerings. Over time, the agentic systems would help in achieving target KPIs, ensuring that goals such as cost reduction, customer satisfaction,

Figure 2. Sales copilot using AI




Source: Infosys

uptime, or compliance are systematically pursued.

Six key uses of agentic AI across the sales process are delineated below:

- 1. Generate dynamic bundling, pricing, and discounts:** Products are atomized (see next section), associated pricing is computed, and discounts that are based on cross-product bundling rules can be achieved.
- 2. Produce well-priced quote based on customer segment and region:** Quote generation is based on products with customer segment and region-specific product bundles applied.
- 3. Initiate self-serve portal AI agent engagement for pricing negotiation based on AI-generated revised quote:** An agent assist copilot helps by providing details on products, pricing, and supports further on pricing negotiation using pricing guidelines.
- 4. Agentic AI workflow for order entry and order orchestration process across systems:** This cross-domain agentic workflow handles the orchestration process by connecting various systems in the order journey process, and includes checks like service availability, credit, and address validation.
- 5. Generative AI-driven digital contracts, contract redlining and contract renewals:** These are based on current pricing and promotional schemes.



6. Agentic AI-driven upgrade

recommendation: This is done on a self-serve portal based on the latest product promotions.

How to make agentic AI work at scale

An important facet of these systems is the ODA box architecture, which helps scale the agentic AI solution in the telecom sales process.

By packaging standardized modules such as recommendation management, knowledge management, product catalogue, order management, and resource inventory into an in-a-box solution, integration with AI agents is enabled, reducing manual effort, and enabling service management and intent-based sales operations.

The in-a-box solution also reduces costs through agile integration and increased interoperability, while time-to-market is decreased as these boxes are pre-defined and pretested. In this architecture, prompt engineering crafts the natural language that agentic AI uses to trigger, sequence, and adapt these standardized modules for the sales copilot and simulator to work effectively.

To get the agents to work properly, two steps are crucial.

Step 1: Standardization of products. This

means breaking down each complex telecom product and standardizing it so that more complex and tailored B2B products can be built from these software modules. This allows ODA modules to be assembled like building blocks, simplifying deployment and minimizing custom coding of the agentic system, while setting up the ability for the AI to build a product autonomously.

Step 2: Standardization of systems and interfaces. Agents can then autonomously connect, provision, and manage services across vendor systems because product interfaces, attribute models, and interaction patterns follow industry norms, reducing integration complexity.

According to our estimates working with clients in telecom sales, this solution enhances pre-sales, leads to higher sales conversion rates, and provides early visibility into the outcome of the sales process. The solution also predicts what the final outcome will be, enabling sales associates to confidently relate to clients throughout the process — a boon for client relationships. In fact, the solution can increase revenue by 5% and 15%, and sales conversions by the same amount.

What this all amounts to is increased revenue growth and a greater slice of the [\\$400 billion promised to telecoms](#) as they optimize B2B pre-sales, sales, and post-sales activities.

Authors and contributors

Rafee Tarafdar | Chief technology officer, Infosys

Harry Keir Hughes | Infosys Knowledge Institute, London

Samad Masood | Infosys Knowledge Institute, London

Chapter 1: Unified operational control in oil and gas

Pradeep Sivakaminathan | Unit technology officer, utilities and energy, Infosys

Chapter 2: Inventory discrepancy in the retail industry

Sankarasubbu Arunachalam | Senior principal technology architect, enterprise application integration services, Infosys

Chapter 3: How to manage discretionary risks in financial services

Anmol Jain | Managing partner, Infosys Consulting, APAC

Shreshta Shyamsundar | Distinguished technologist, advanced software engineering and AI, Infosys

Chapter 4: Agentic AI is redefining the manufacturing value chain

Dr. G. V. V. Ravi Kumar | Vice president and unit technology officer, core and digital product engineering services, Infosys

Sundaresan Poovalingam | Distinguished technologist, robotics and autonomous systems, Infosys

Dr. Sunitha Abburu | Principal consultant, advanced engineering group, Infosys

Sachin Shirguppe | Principal consultant, advanced engineering group, Infosys

Chapter 5: How agentic AI is re-engineering telecom business sales

Gnanapriya Chidambaranathan | Associate vice president and unit technology officer, CMT, Infosys

Editorial and production

Vidya Lakshmi Hariharan | Delivery manager, Infosys

Jeff Mosier | Infosys Knowledge Institute, Dallas

Kate Bevan | Infosys Knowledge Institute, London

Pragya Rai | Infosys Knowledge Institute, Bengaluru

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at infosys.com/IKI or email us at iki@infosys.com.

For more information, contact askus@infosys.com



© 2025 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.