LTIMindtree BlueVerse

Whitepaper

Unlocking Agentic AI

Transforming Processes and Operations at Scale

LTIMindtree BlueVerse

Contents

Executive Summary

Agentic AI is redefining how organizations approach process automation and operational efficiency. By enabling autonomous, goal-driven decision-making, it goes beyond traditional AI systems to deliver self-healing capabilities, proactive actions, scalability, cost savings, and improved accuracy.

This whitepaper explores how agentic AI can transform business processes and operations at scale, with a special focus on the European Union's regulatory landscape. As the EU AI Act introduces stricter requirements for transparency, accountability, and compliance, businesses must navigate these complexities to responsibly adopt high-risk AI systems.

Drawing on real-world case studies, the paper demonstrates how organizations across sectors, including IT, sales and marketing, and after-sales service, are applying agentic AI to achieve measurable business outcomes. These examples provide actionable strategies for leaders seeking AI-driven process automation and effective AI in business operations.

By combining regulatory insights with practical applications, this paper provides a roadmap for organizations to unlock the full potential of agentic AI and embrace enterprise AI transformation to remain competitive in an evolving digital economy.

Introduction

AI is undergoing a significant shift, from reactive systems that respond to inputs, to proactive systems capable of autonomous action and collaboration. Generative AI, which focuses on producing content from user prompts, has paved the way for agentic AI, where intelligent agents independently set and achieve goals, interact with other systems, and drive complex workflows without constant human oversight.

In our earlier paper[1], "Architecting Agentic AI", we explored the core principles, implementation challenges, and emerging trends surrounding this technology. Building on that work, this paper focuses on two practical applications: process-based AI, which automates individual tasks and workflows, and operations-based AI, which optimizes broader business functions to deliver AI in business operations at scale.

With the European Union taking a leading role in AI governance through the EU AI Act, organizations face both opportunities and challenges in adopting these advanced systems. The following sections explore this regulatory environment and its implications, setting the stage for how agentic AI can be deployed responsibly as part of an enterprise AI transformation strategy.

European Context for AI use

EU AI Act

The European Union has introduced the EU AI Act; a key regulatory framework aimed at governing the use of artificial intelligence within the European Union. The Act classifies AI systems into risk levels, unacceptable risk, high risk, and minimal or no risk, based on the potential harm they may cause[2].

According to Article 3(2)[3] of the Act, risk is defined as "the combination of the probability of an occurrence of harm and the severity of that harm". This definition has several implications for agentic AI systems:

Due to its autonomous and complex nature, agentic AI is likely to be classified as high risk, potentially affecting how it is adopted and deployed.

In sensitive sectors such as healthcare, transportation, and education, where decisions directly impact human lives and involve personal data, agentic AI will face stricter regulatory scrutiny.

The Act places strong emphasis on transparency and accountability, especially for high-risk systems. Agentic AI solutions that are designed with clear auditability and ethical practices will be better positioned for approval and may face reduced regulatory barriers.

Agentic AI in Action
Automated Insurance Claims Processing

A practical example of agentic AI in action comes from a major Dutch insurer that automated up to 90% of its automobile claims workflow using custom AI agents[4]. These autonomous systems review and assess claims without human intervention, leading to:

Significant operational efficiency gains

Faster claims processing times, and

The ability for employees to focus on creative and strategic tasks rather than routine administrative work.

Cyber Resilience Act (CRA)

The other major regulatory frameworks that can affect adoption of agentic AI systems in Europe are Cyber Resilience Act[5] and NIS2 Directive[6].

The CRA ensures digital products which include hardware and software meet cybersecurity requirements. It promotes a "cybersecurity by design" approach to minimize the likelihood and impact of cyberattacks. Connected technologies, such as IoT devices, AI-driven tools, and analytics platforms, fall within its scope. Manufacturers are required to:

NIS2 Directive

The NIS2 directive expands on the original Network and Information Systems (NIS) Directive, extending its scope to a broader set of sectors, including:

- Energy, transportation, and healthcare,
- Finance, water management, and digital infrastructure,
- Public communications and digital services such as social platforms,
- Waste management, postal and courier services, and
- Critical product manufacturing and public administration.[6]

NIS2 requires enhanced incident reporting and stricter security measures, ensuring that essential service providers maintain a high standard of cybersecurity[7].

Impact of CRA and NIS2 Directive on Agentic AI Adoption

Noncompliance with either regulation carries significant penalties, fines under the CRA can reach 2.5% of global annual turnover[8]. For organizations deploying agentic AI, this means factoring in cybersecurity and compliance from the earliest stages of development. Key considerations include:

- Higher Development Costs: Investments are needed to ensure secure design, proactive monitoring, and robust incident reporting mechanisms.
- Regular Audits: Frequent, rigorous audits must be conducted to identify and mitigate cybersecurity risks.
- Cross-Functional Collaboration: Legal, cybersecurity, and AI teams must work closely from design to deployment to maintain compliance and reduce exposure to risk.

By addressing these requirements early, organizations can reduce compliance hurdles and strengthen trust in their AI systems.

Process-Based AI

The next section examines two core applications of agentic AI: process-based AI and operations-based AI.

Process-based agentic AI focuses on automating individual workflows and tasks. These systems break processes into multiple sub-processes, with specialized AI agents managing each step autonomously.

A comprehensive knowledge base, such as a large language model (LLM), supports decision-making by providing access to historical data, standard operating procedures (SOPs), rules, and workflows. This enables the system to perceive its environment, process information, and act independently to streamline and optimize operations.

Example: IT Service Management (ITSM) Using Agentic AI

Building the Knowledge Base for Ticket Resolution

In ITSM, agentic AI can significantly enhance problem-solving accuracy, speed up ticket resolution, and automate various repetitive sub-processes. The foundation of this system lies in a robust knowledge base, built using:

- Historical ticket data and resolutions,
- Records of previously implemented solutions,
- Categorization of issue and incident types, and
- Knowledge and resource data from connected systems such as CRM, and ERP.

Multiple AI Agents use this information to simplify sub-processes and automate various steps in the ticket management process.

How Agentic AI Resolves Issues

When a new ticket is raised, agentic AI performs several functions autonomously:

- Categorizes the Issue: Determines the type, severity, and potential impact of the problem.
- Searches for Solutions: References historical data, knowledge bases, and predefined rules to identify possible resolutions.
- Acts on Insights: Automatically resolves the issue, suggests potential solutions for user selection, or assigns the ticket to the most qualified resource based on historical performance patterns.

Intelligent Issue Prioritization

Additionally, the agentic AI-based ITSM system has the capability to identify and segregate multiple issues and incidents, ranking them based on priority. This prioritization can be determined by factors such as the user raising the issue, the impact on workflows, and the severity of the problem.

The system then addresses these issues as a priority. Furthermore, it can recognize duplicate service requests and suppress them, focusing only on the original request. This reduces the overall number of tickets and shortens the time required to resolve issues. The below figure shows how agentic AI system can be leveraged for ITSM.

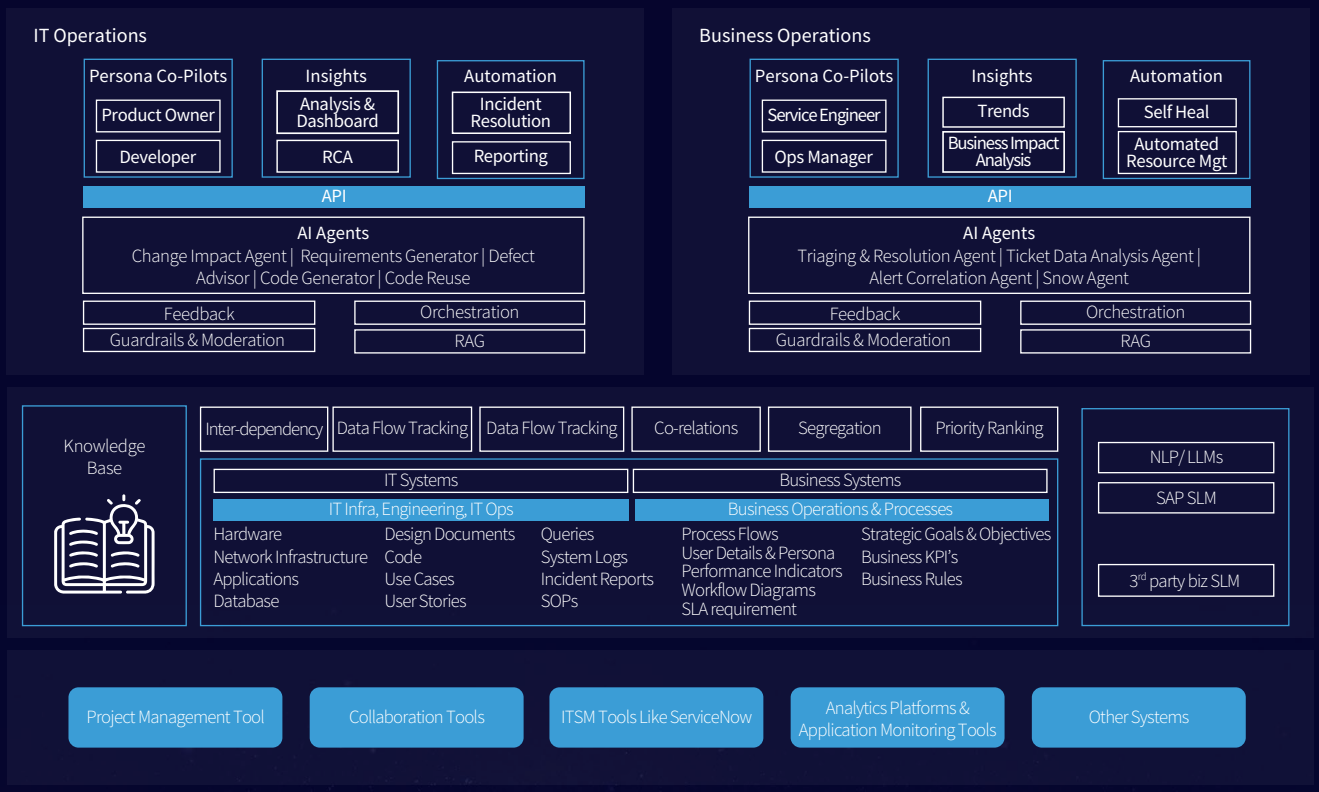

Figure 1: Typical Architecture of Agentic AI system for ITSM Process

SUCCESS STORY 1

Optimizing Configurations and Quotes for CHVAC Products

A global leader in building cooling solutions faced challenges in their sales process. They sought to improve price prediction accuracy for new product configurations while increasing the volume of quotes analyzed.

LTIMindtree deployed an agentic AI-powered multi-agent system that used advanced machine learning and deep learning models, alongside historical customer data, to recommend optimal configurations and generate accurate quotes.

Key Outcomes

- Accurate Price Predictions: Achieved 82% accuracy for new product configurations using 18 ML models and three deep learning models.
- Boosted Sales Potential: Enabled a 7x increase in sales opportunities.
- Enhanced Quote Analysis: Processed nearly three million quotes post-implementation.
- Improved Sales Conversions: Platform incorporated market-based pricing logic and tracked pricing decisions to support better lead conversion.

Operations-Based AI

How Agentic AI Supports Business Operations

Operations-based AI applies agentic AI to optimize and automate broader business operations, rather than focusing solely on isolated processes. This approach enables real-time decision-making, adaptability, and self-healing capabilities, a hallmark of AI in business operations.

Common applications include:

- Autonomous vehicles,
- Predictive asset maintenance, and
- Smart city energy management systems.

By dividing large-scale operations into manageable subtasks, synchronized AI agents can collaborate to ensure continuous optimization and rapid response to changing conditions.

Example: Predictive Maintenance

For predictive maintenance of assets, operations-based AI systems use machine learning, data analytics, and real-time data to forecast equipment failure and take appropriate action to prevent it. This enables the system to shift from a reactive approach to an initiative-taking approach, helping reduce unplanned downtimes, improve equipment lifecycle, and increase operational efficiency.

The predictive maintenance system requires real-time operating data like pressure, temperature, etc., from equipment, which is collected from sensors. The data from the sensors is then cleansed, and appropriate parameters are selected through data analytics.

A machine learning algorithm is used to identify anomalies and compare them with historical operational data to detect patterns and predict equipment failure. Afterward, the agentic AI system selects an appropriate resource based on availability and schedules maintenance to avoid equipment failure.

Figure-2 shows an agentic AI model for predictive maintenance. It consists of multiple AI agents and a data-and-decision Co-Pilot that automates processes, provides insights, and assists key personas in data processing and decision-making in maintenance.

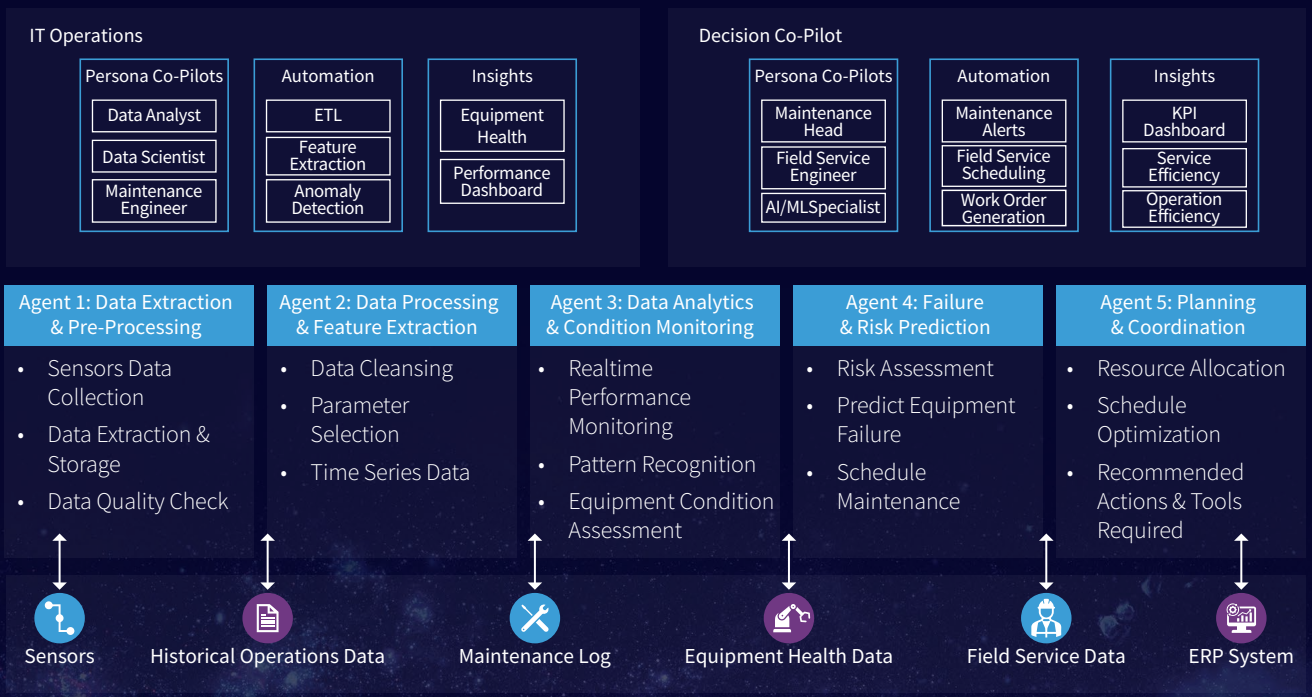

Figure 2: Operations-Based Agentic AI for Predictive Maintenance Operation

SUCCESS STORY 2

Predictive Maintenance for Elevators and Escalators

A leading global elevator company partnered with LTIMindtree to implement a "Transportation as a Service" predictive maintenance model, enhancing both service quality and customer experience.

The solution used IoT sensors and agentic AI to monitor elevator and escalator health in real time and to provide proactive maintenance alerts.

Key Outcomes

- Revenue Growth: Increased service contract revenue by 2%
- Cost Reduction: Reduced field service costs by 33%
- Faster Turnaround: Cut service turnaround time by 50%
- Data-Driven Monitoring: Collected IoT data from 284,000 units globally, deploying 14 AI models per unit
- Proactive Alerts: Real-time monitoring and automated notifications via a mobile app, supported by historical service data for predictive analytics

Necessity of Agentic AI in Process-Based and Operations-Based Systems

The nature of agentic AI systems makes them highly suitable for both process-based and operations-based applications. These systems not only outperform non-AI solutions but also deliver distinct advantages over traditional or standalone AI, which often depend heavily on manual oversight and static, pre-trained data.

The key benefits of agentic AI over traditional AI can be summarized as follows:

- **Self-Healing and Autonomous Nature:** Agentic AI systems operate independently, completing multiple tasks without human intervention. If a defined goal is not met, the system can auto correct and adjust its approach. In contrast, traditional AI requires ongoing human monitoring and manual corrections.
- **Proactive Approach:** Unlike traditional AI, which is reactive and waits for input, agentic AI anticipates issues and initiates solutions proactively. This makes it particularly effective in dynamic, fast-changing environments.
- **Accuracy Improvement:** Goal-driven by design, agentic AI consistently enhances the accuracy of outputs over time. It minimizes human error and is especially valuable in high-precision areas such as financial forecasting and medical diagnosis.
- **Scalability and Cost-Effectiveness:** These systems are built to manage large workloads and datasets without compromising speed or quality. By automating complex tasks, agentic AI reduces reliance on manual effort, cutting down costs, time, and resource requirements.

Challenges in Adopting Agentic AI Systems

While agentic AI offers transformational benefits, its adoption comes with a variety of technical, regulatory, and operational challenges. Addressing these concerns early is essential for smooth implementation.

- **Stringent Regulations:** The EU AI Act and other frameworks emphasize transparency and accountability, especially for high-risk AI. The complex algorithms behind agentic AI may produce outcomes that are difficult to interpret, creating trust and compliance issues.
- **Cybersecurity Risks:** Because agentic AI systems interact with third-party applications and handle sensitive data, they require robust cybersecurity measures to prevent breaches and unauthorized access.
- **Data Readiness:** High-quality, diverse datasets are critical for accurate outcomes. Poor data quality or incomplete datasets can lead to errors or bias, limiting the system's effectiveness. A thorough data readiness assessment is vital before deployment.
- **Operational Integration Challenges:** Implementation can be complex, involving integration with legacy systems, scaling across environments, performance monitoring, and ensuring access to clean, relevant training data.
- **User Acceptance and Trust:** Agentic AI may be perceived as a threat to jobs or as a "black box" system due to its complexity. Building user trust through transparency, education, and pre-adoption strategies is crucial for long-term success.

Among the many organizations working in this space, LTIMindtree offers a structured approach to help businesses adopt agentic AI responsibly and effectively. This includes conducting end-to-end assessments of current IT landscapes to evaluate system maturity, identifying gaps, and enabling seamless integration of agentic AI systems. The AI practice teams at LTIMindtree include over 500 data scientists, and more than ten accelerators to expedite AI rollouts, supported by eight Centers of Excellence (COEs) that serve over 100 customers globally.

Conclusion

Agentic AI represents a paradigm shift in how organizations manage processes and operations. By enabling autonomous decision-making and optimizing workflows with minimal human intervention, it drives adaptability, precision, and scalability at a level traditional AI cannot match.

However, its full potential can only be realized by carefully addressing regulatory, ethical, and operational challenges. Building trust through transparency, ensuring data quality, and planning for seamless integration are essential steps for successful adoption.

As AI continues to advance, agentic AI will play a pivotal role in shaping the future of autonomous systems across industries. Organizations that develop a clear roadmap and well-defined implementation strategy today will be best positioned to lead in this emerging landscape, achieving meaningful AI in business operations, AI-driven process automation, and successful enterprise AI transformation.

References

[1] *Beyond Intelligence: Architecting Agentic AI Enterprise,* Prasanna S, LTIMindtree, August 2025: https://www.ltimindtree.com/blueverse/beyond-intelligence-architecting-the-agentic-ai-enterprise/

[2] *Governing Agentic AI: Addressing Security and Compliance Gaps in the EU AI Act,* Sandeep Purewal, May 2025: https://hiddenlayer.com/innovation-hub/governing-agentic-ai/

[3] *EU Artificial Intelligence Act:* https://artificialintelligenceact.eu/article/3/

[4] *40+ Agentic AI Use Cases with Real-life Examples,* Cem Dilmegani and Mert Palazoğlu, AI Multiple Research, August 2025: https://research.aimultiple.com/agentic-ai/

[5] *Cyber Resilience Act, Policy and Legislation,* European Commission, September 2022: https://digital-strategy.ec.europa.eu/en/library/cyber-resilience-act

[6] *NIS2 Directive: securing network and information systems, European Commission:* https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

[7] *The NIS 2 Directive | Updates, Compliance:* https://www.nis-2-directive.com/

[8] *Final Text, European Cyber Resilience Act: https://www.european-cyber-resilience-act.com/Cyber_Resilience_Act_Article_64.html#:~:text=Non%2Dcompliance%20with%20the%20obligations,operator%20for%20a%20similar%20infringement;*

LTIMindtree BlueVerse

About the Author

Vishal Vadabhat

Senior Manager-Program and Project Management, CTO Europe

Vishal is a Senior Manager with over 10 years of experience in Engineering and IT Consulting. At LTIMindtree, he combines deep domain expertise with strategic insight to drive technological innovation. With a strong background in Manufacturing, Vishal has led transformative initiatives that leverage emerging technologies to solve complex business challenges. He currently supports business development for the Manufacturing portfolio in Europe.

Prasanna S

Chief Technology Officer, Europe

Prasanna is an accomplished IT leader with over 25 years of experience in designing solutions and platforms that help organizations navigate the digital landscape. He has been instrumental in driving digital transformation, managing large teams, and fostering innovation across industrial manufacturing, intelligent transportation, and telecommunications. Prasanna emphasizes understanding the business context to build technology solutions that deliver real impact.

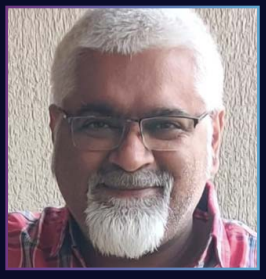

Sagar Mulay

Director - Program & Project Management, CTO Europe

Sagar Mulay leads the Solution Architect tower in the Europe Technology Office. With over 25 years of experience, he has successfully delivered digital transformation initiatives, managed large teams, and driven innovation across Banking, Manufacturing, and Hi-Tech industries. Sagar is passionate about creating customer-centric solutions that prioritize innovation, transformation, and modernization.