THE SILENT INTELLIGENCE

The paradigm shift to intelligent and autonomous application management

Jérôme Fenyo

GOYAVE
CONSULTING

Quoi de plus complet que le silence ?*

Honoré de Balzac

**For what is more complete than silence?*

Why this paper?

Conversations about AI in the IT industry often focus on how consultants – shaping strategies and setting up transformations – and developers – writing code and delivering new solutions – are being augmented or even replaced. But this view ignores financial reality. For decades, around **60–70% of IT budgets have been tied up in maintenance** rather than new initiatives (Aaron Percival's synthesis of CIO benchmarks). This is where the true value and cost reside.

This sector, Application Management Services (AMS), is currently valued at **~$36–42 billion in 2024** (The Business Research Company: $36.1 billion; Grand View Research: $34.4 billion; Precedence Research: ~$42.0 billion). Analysts project a **~15–20% CAGR**, reaching ~$120 billion by 2030 (Grand View Research) and ~$265 billion by 2034 (Precedence Research). With total global IT spend at $5–5.5 trillion (Gartner), AMS-type activities likely represent **$1–1.5 trillion annually**.

This is the "silent" part of IT: less glamorous, rarely on stage, but fundamental. AI, GenAI and Agentic AI are now starting to **rewire this space at scale** – and whoever industrializes AMS with AI first will be **competing for leadership** in one of the largest, most contested markets in the entire IT value chain.

Yet, the most vital aspect may lie elsewhere. I have always held the deep conviction that AMS is far more than an industrial model or technical know-how; it is, above all, a **holistic experience**. Maintaining a legacy estate means knowing it intimately—visualizing it, understanding it, and shaping it daily to drive agility and unlock business value. **AI is nothing less than the heartbeat that allows us to truly feel and live this legacy experience.**

This white paper is first and foremost a **late-2025 state-of-the-art on the transformation of Application Management Services (AMS) into Intelligent Application Management (iAM)**. It draws extensively on analyses and reference frameworks published by major market actors (Accenture, Wipro, Infosys, Cognizant, SmartDev, ISG & HCLTech, Capgemini, Quinnox, IBM, among others), which are listed in detail in the "Inspiring Sources" section.

In particular, our discussion of **NextGen AMS constructs and hyperautomation levers** is informed by publications from Accenture and Wipro; our view on **AI-first and experience-led operating models, Experience Level Agreements (XLAs)** and Total Experience is inspired by work from Infosys, Cognizant and ISG & HCLTech; the catalogue of **AI use cases** and associated performance metrics builds on analyses by SmartDev and several vendor case studies; and the framing of **AIOps, observability and iAM itself** borrows from contributions by Quinnox and IBM. Finally, the positioning around "invisible" or "silent" intelligence echoes the perspectives developed by Capgemini on invisible AI.

Throughout this white paper, these concepts are deliberately synthesized, recombined and extended into a **unified perspective centered on the notion of "Silent Intelligence".**

Scope and Positioning

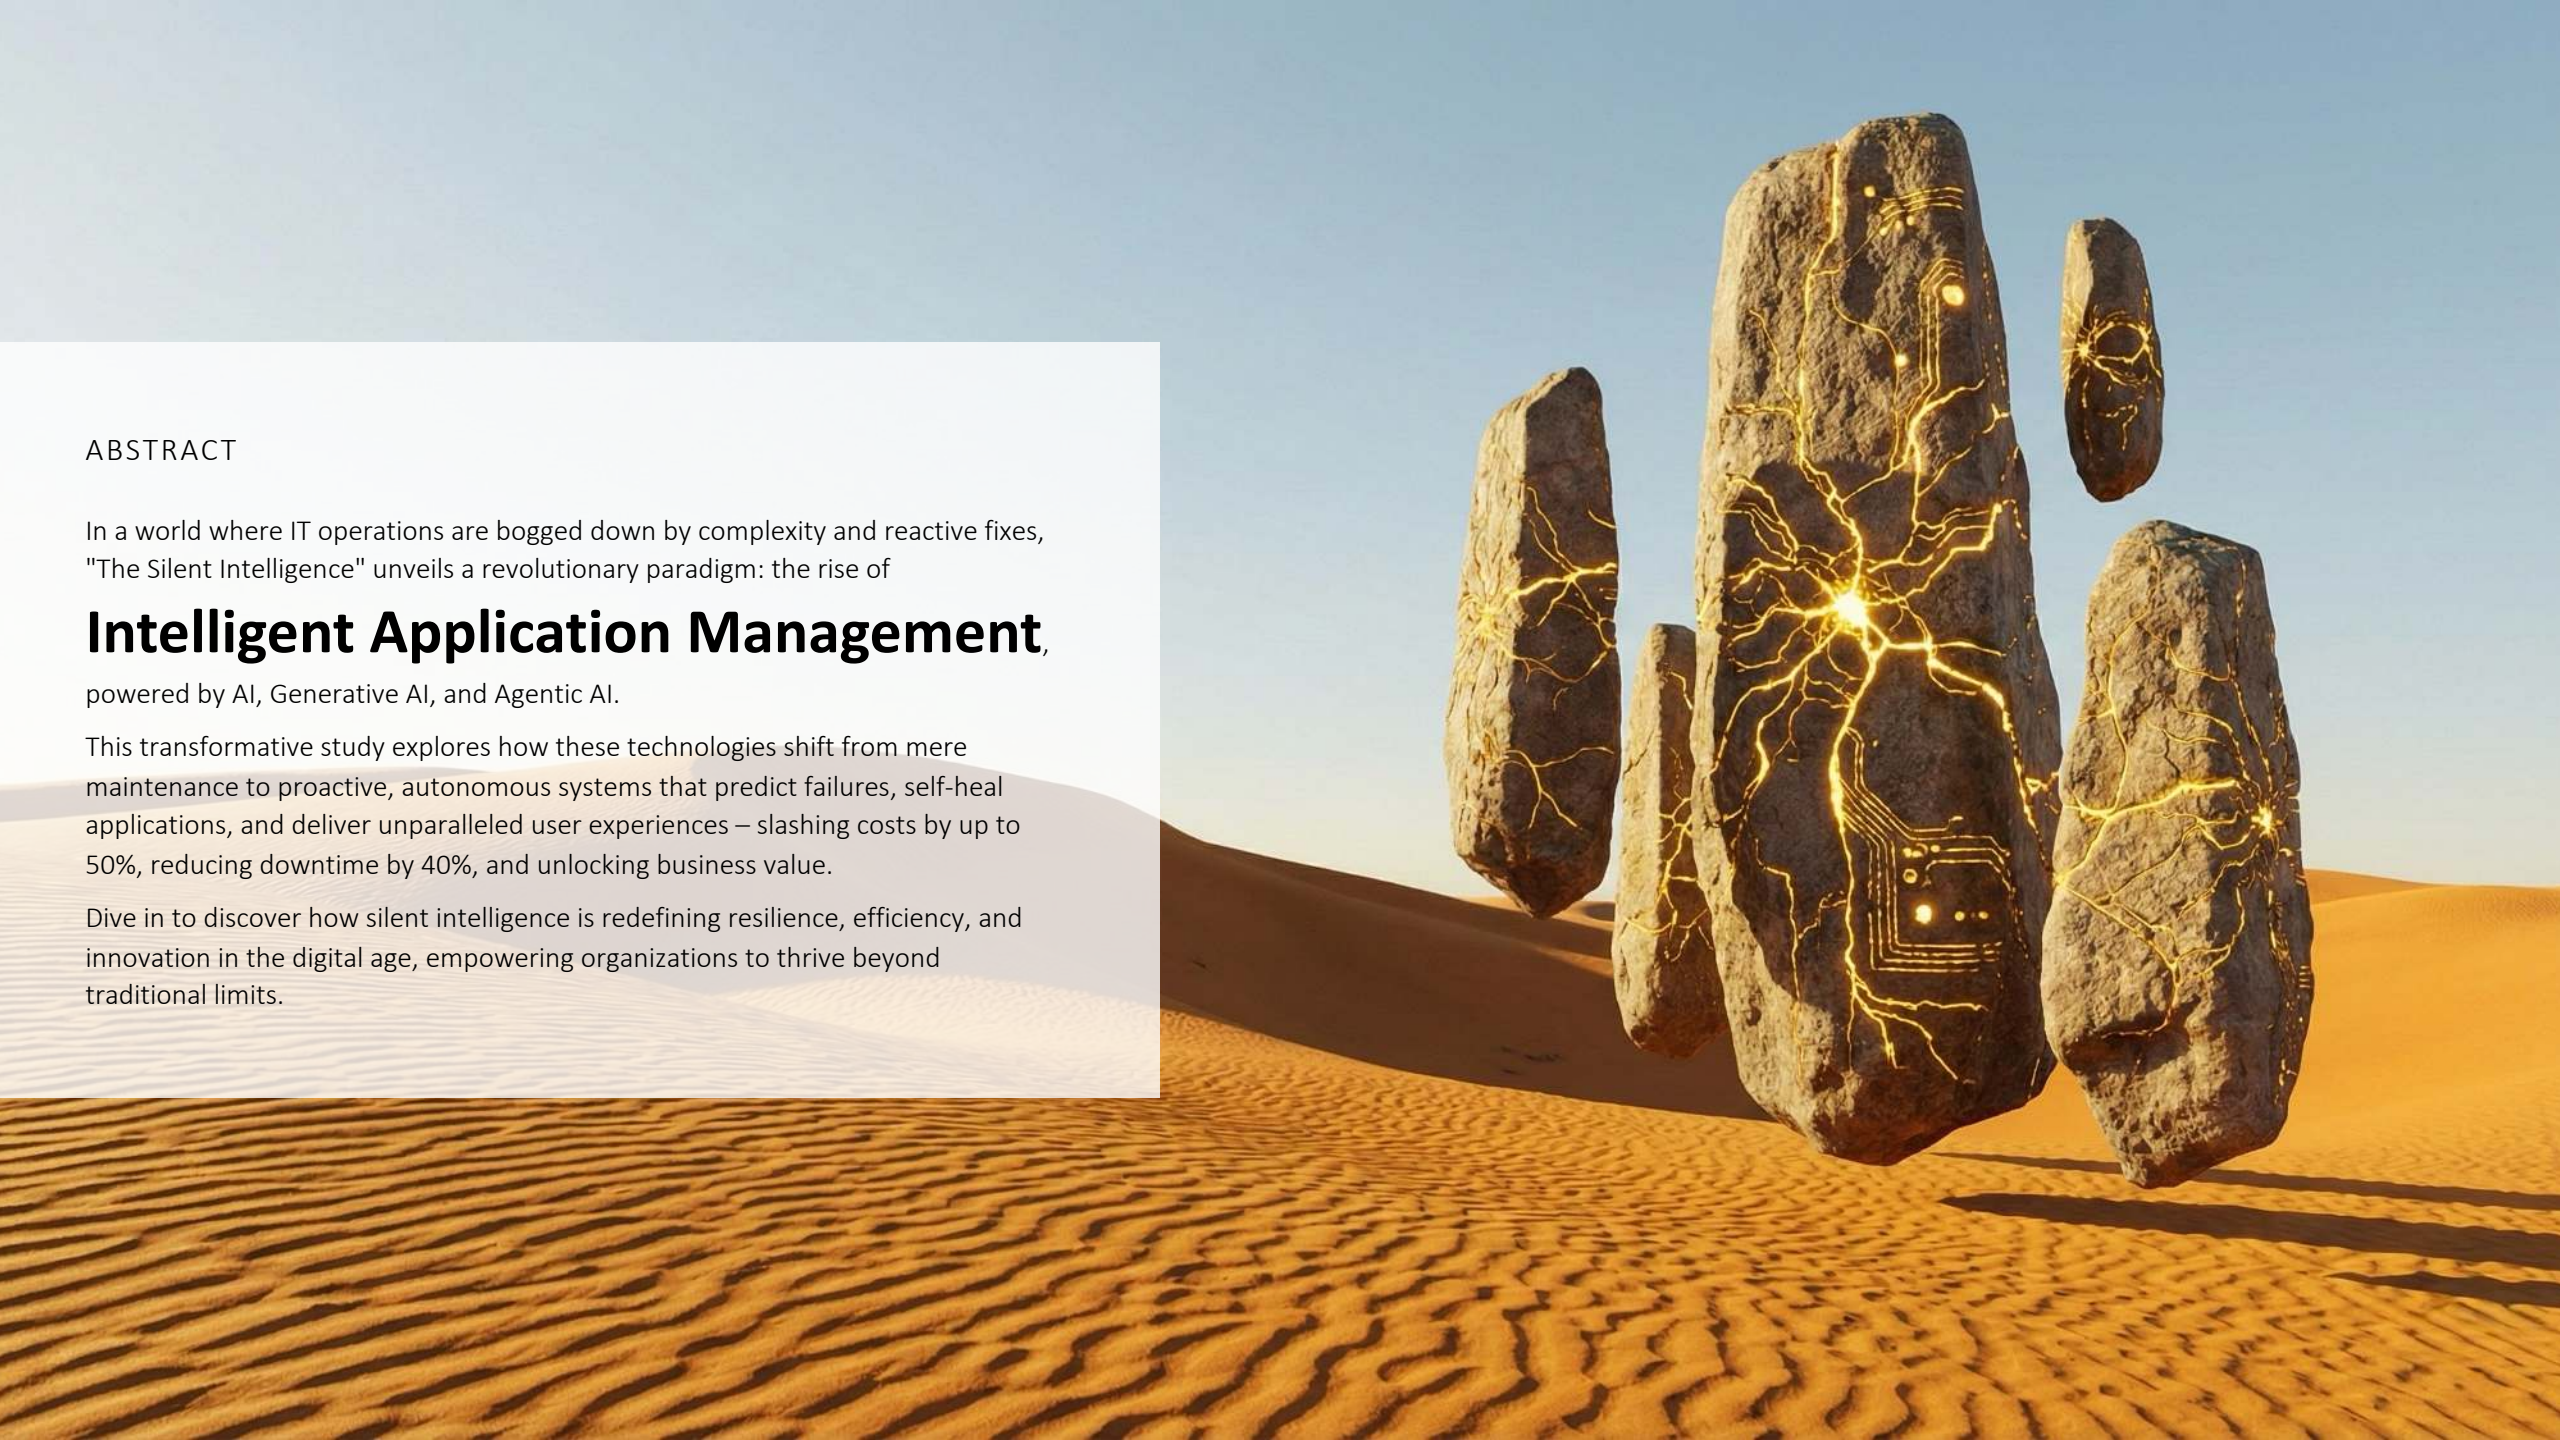

ABSTRACT

In a world where IT operations are bogged down by complexity and reactive fixes, "The Silent Intelligence" unveils a revolutionary paradigm: the rise of

Intelligent Application Management,

powered by AI, Generative AI, and Agentic AI.

This transformative study explores how these technologies shift from mere maintenance to proactive, autonomous systems that predict failures, self-heal applications, and deliver unparalleled user experiences – slashing costs by up to 50%, reducing downtime by 40%, and unlocking business value.

Dive in to discover how silent intelligence is redefining resilience, efficiency, and innovation in the digital age, empowering organizations to thrive beyond traditional limits.

TABLE OF CONTENTS

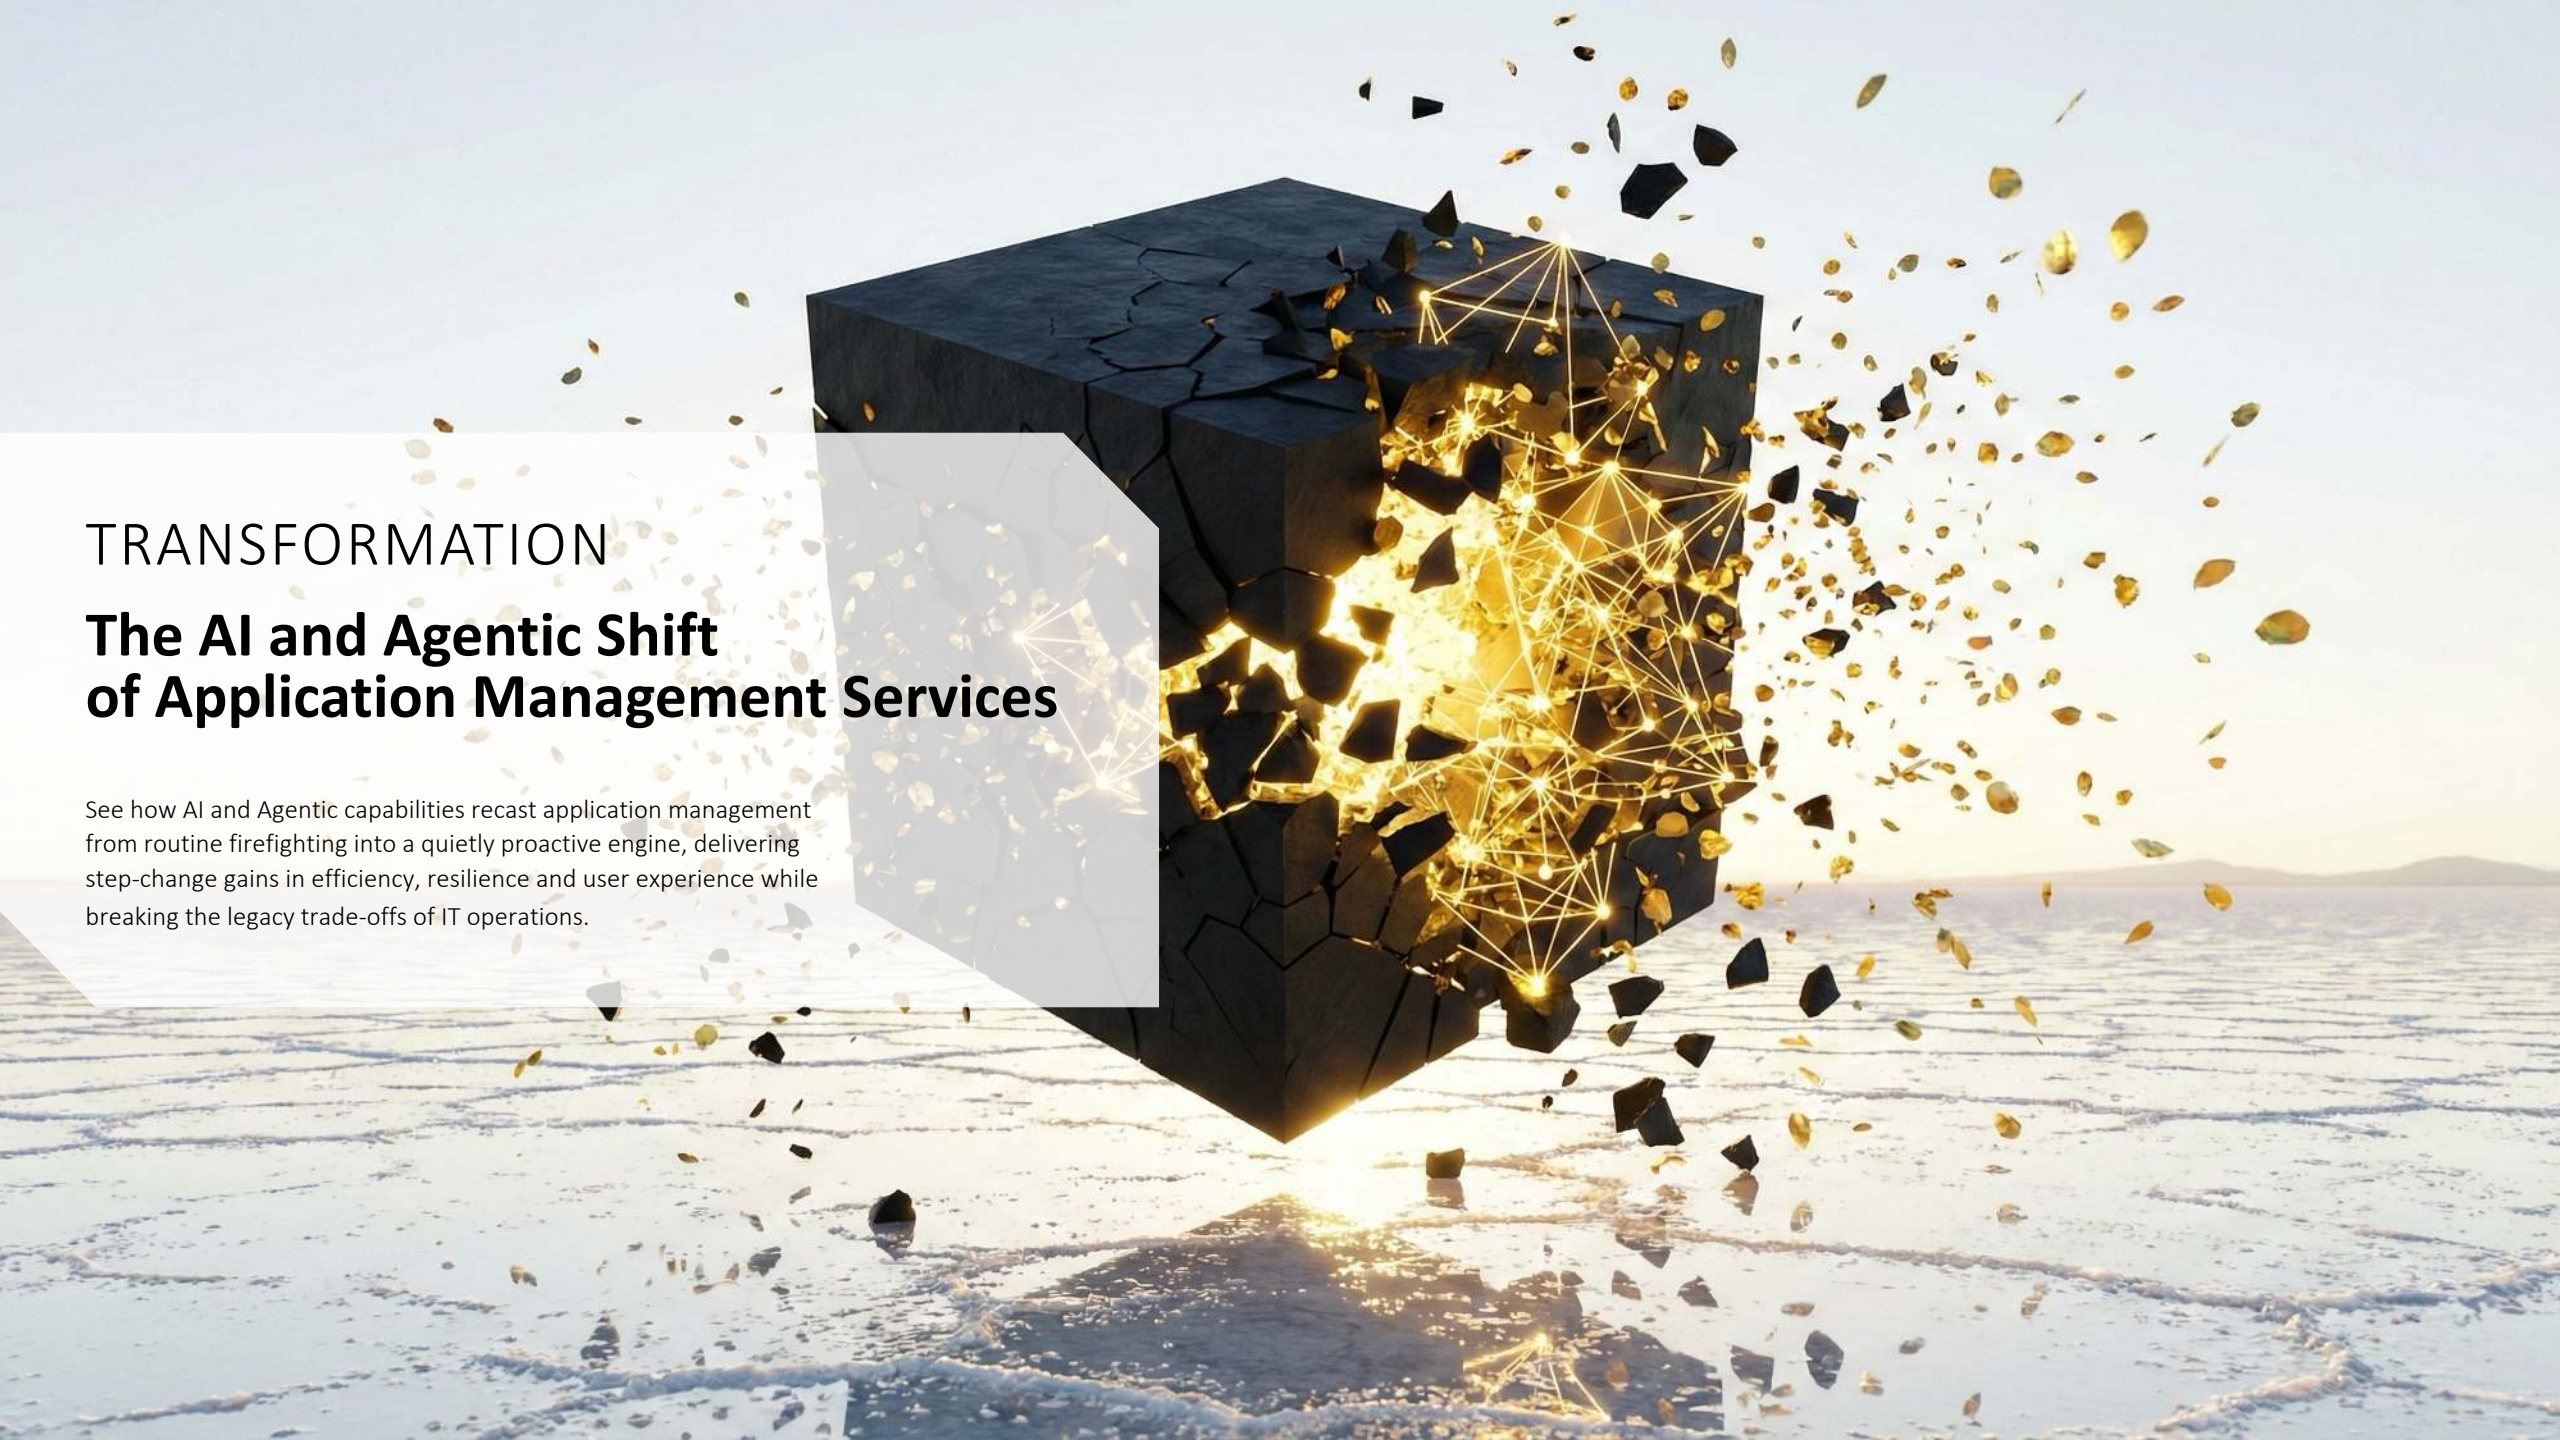

TRANSFORMATION

The AI and Agentic Shift of Application Management Services

See how AI and Agentic capabilities recast application management from routine firefighting into a quietly proactive engine, delivering step-change gains in efficiency, resilience and user experience while breaking the legacy trade-offs of IT operations.

Re-inventing Application Management in the Age of AI

The transformation of Application Management Services (AMS) through the integration of Artificial Intelligence (AI), Generative AI (GenAI), and Agentic AI shifts AMS from a traditional, reactive support function into a proactive, intelligent, and experience-driven operation, often referred to as Intelligent Application Management (iAM).

The need for AI in AMS is driven by increasing application complexity, growing user bases, pressure for rapid issue resolution, and the diminishing returns of conventional, rule-based automation. AI capabilities now play a crucial role in reshaping AMS so it can deliver richer experiences and stronger resilience while still driving cost efficiency.

The industry expects AMS partners to transition to AI-powered service management to achieve significant cost reduction and maximum operational efficiency. This evolution tackles the traditional trade-off triangle in IT operations: value delivered to the business, cost to run the estate, and speed of change. Mature iAM models show that these three dimensions can be improved together instead of being traded off against each other.

Key Shifts toward Intelligent Application Management

Key characteristics of this transformation include:

Shift from Reactive to Proactive: Moving from fixing issues after they occur to predicting and preventing them.

Focus on Experience: Replacing traditional Service Level Agreements (SLAs) with Experience Level Agreements (XLAs) that more precisely track user sentiment and engagement.

Intelligence at the Core: Integrating AI/ML and Generative AI to streamline incident management, automate ticket classification, and enable predictive maintenance.

APPLICATION

Key Use Cases for Intelligent Application Management

Unleash the transformative potential of AI in real-world scenarios, from automating incident handling and predictive maintenance to accelerating root cause analysis and hyperautomation, promising dramatic reductions in downtime and boosts in productivity that redefine what's possible in application support.

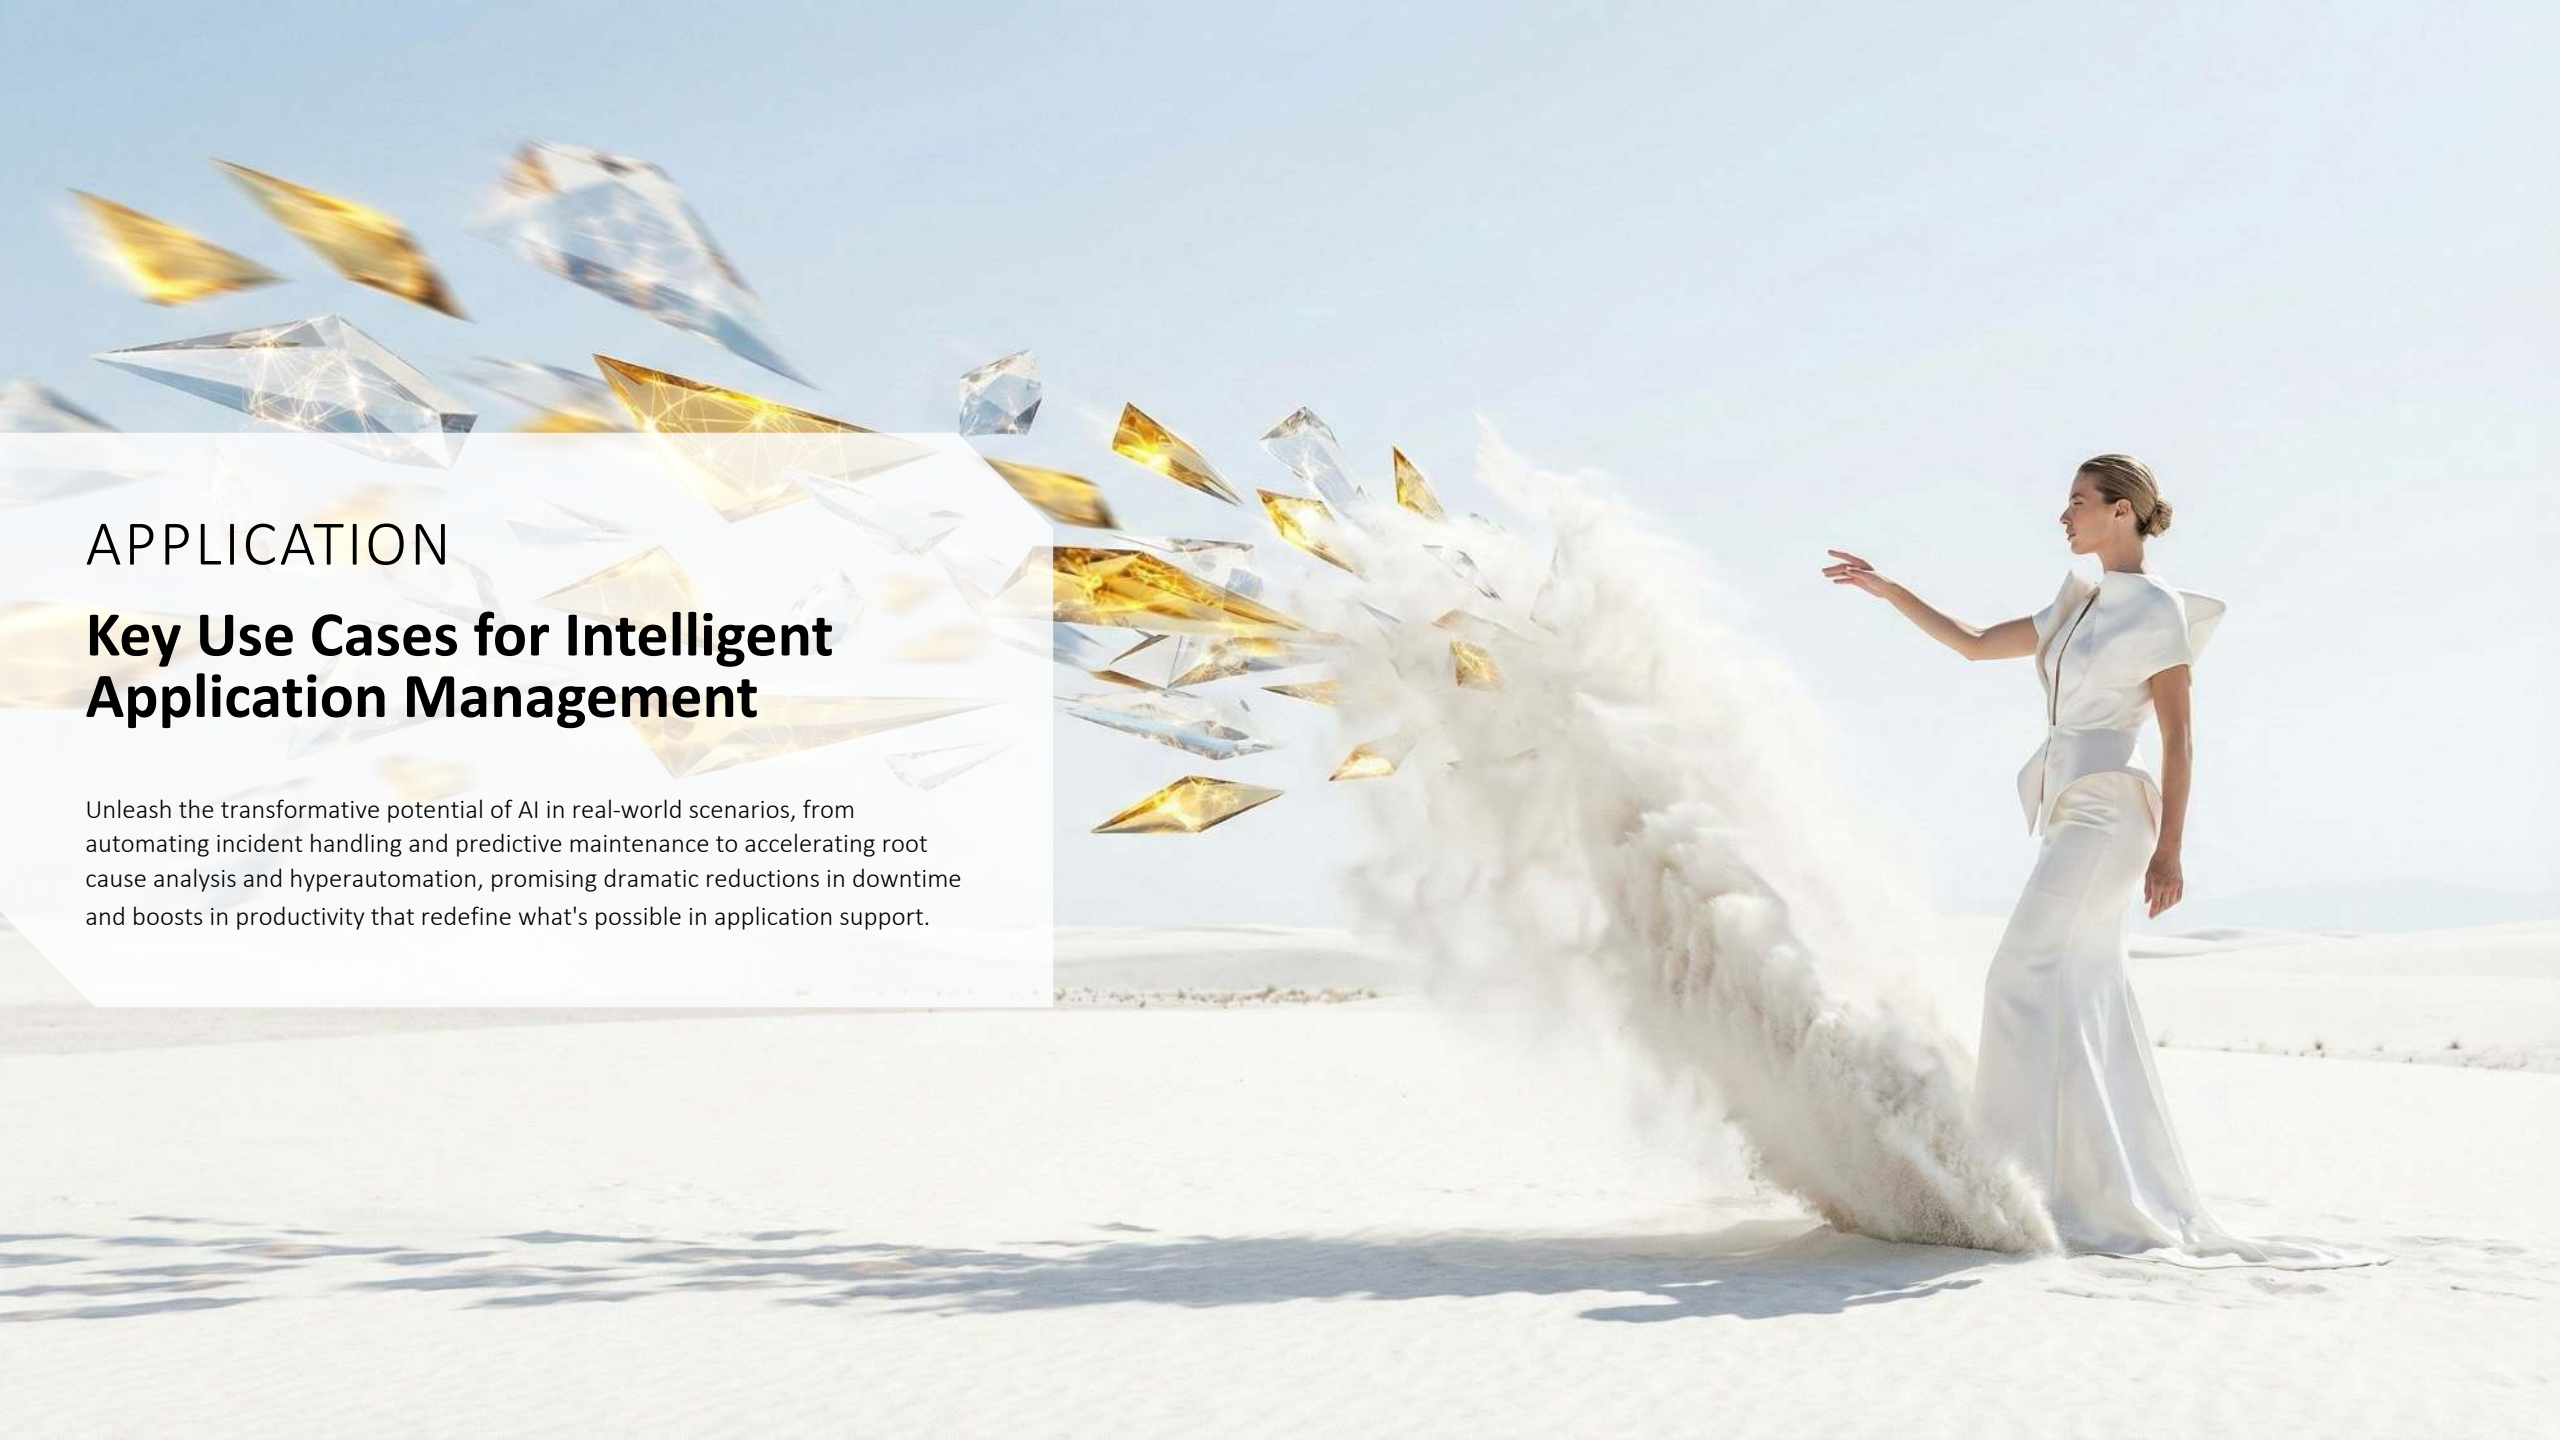

Seven Use Cases & Applications of AI

The core Application Management Services (AMS) transformation into Intelligent Application Management (iAM) lies in several key use cases and applications that drive efficiency, deliver proactive support, and enhance the overall user experience.

This section presents **7 key use cases and applications of AI** in Application Management Services (AMS), reviewing each and highlighting the main benefits that can be expected.

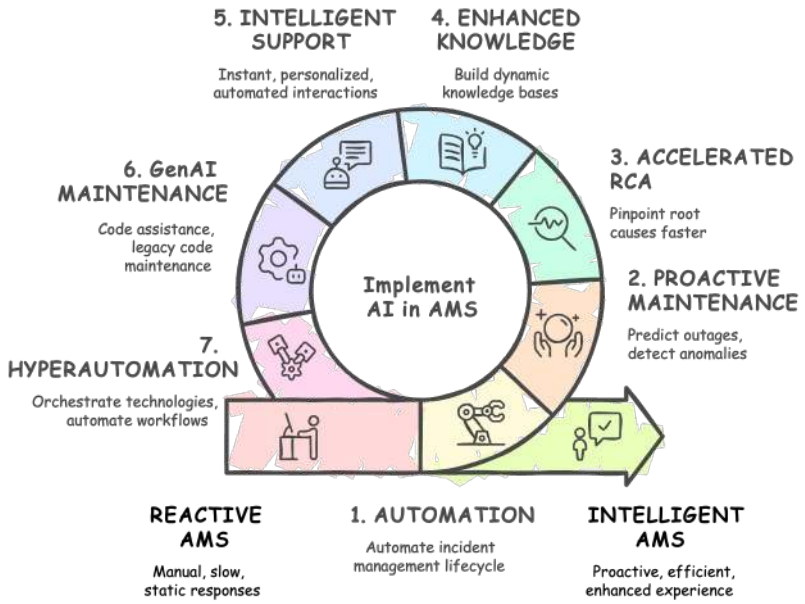

1. AUTOMATION OF INCIDENT MANAGEMENT AND TICKETING

A primary application of AI in AMS is the automation and streamlining of the incident management lifecycle, which is traditionally dependent on manual processes.

Automated Detection and Classification: Instead of relying on manual log reading and ticket review, AI pipelines monitor events, metrics, traces, and user reports in real time. They flag anomalies and new requests, interpret natural language descriptions, and infer the type of issue and its impact. Incidents are automatically enriched with standard categories, priority, and urgency codes before they enter the queue. This cuts repetitive triage work, keeps labels consistent across teams, and reduces the delay between first signal and actionable ticket.

Intelligent Routing and Assignment: Once incidents are classified, AI uses skills, workload, and past ownership to decide who should handle them. Tickets go straight to the right squad or engineer instead of bouncing between queues. This cuts misrouting and reassignments and speeds up response.

Incident Clustering and Correlation: AI/ML models perform cluster analysis to identify patterns and group similar issues, which helps in predicting potential major incidents by spotting fast-growing clusters of related failures. By merging alerts from different monitoring tools into a small number of incident clusters, AIOps cuts alert noise and lets teams focus on real problems.

Case Study Example: ServiceNow utilizes AI-driven virtual agents and ML models within its ITSM platform for automatic categorization, prioritization, and routing of support tickets, which has resulted in customers reporting up to a 30% reduction in manual ticket handling and a 25% improvement in Mean Time To Resolution (MTTR) – SmartDev | Ngoc Nguyen, 2025.

2. PROACTIVE AND PREDICTIVE MAINTENANCE (AIOps/iAM)

AI enables AMS to shift from a reactive, break-fix model to a proactive, predictive service delivery model.

Predictive Analytics: AI models analyze historical incident data, logs, and real-time system metrics to forecast potential application outages, component failures, or performance degradation. This capability allows support teams to take corrective action before users experience service degradation.

Proactive Issue Detection: AI agents continuously monitor metrics and logs, detecting subtle anomalies and deviations from normal behavior that signal an impending issue.

***Example**: An AIOps platform can forecast failures with 85% accuracy and assess an 87% risk of the e-commerce cart becoming unavailable within 24 hours – Quinnox, 2024.*

Autonomous Remediation and Self-Healing: A major application of AIOps is the creation of self-healing systems. AI agents and cognitive bots, powered by AI analytics, can perform required remedial actions without human intervention, such as automatically detecting database performance degradation and triggering DB reindexing, or identifying high CPU usage and initiating auto-scaling to maintain stability. Quinnox (2024) notes that AIOps-based remediation can deliver a 15–45% drop in high-priority incidents and a 20–40% reduction in unplanned downtime.

***Case Study Example**: IBM's Watson AIOps platform leverages AI to predict outages and performance degradation, delivering early warnings that have helped clients achieve up to a 40% reduction in incident volumes and a 35% decrease in downtime – SmartDev | Ngoc Nguyen, 2025.*

3. ACCELERATED ROOT CAUSE ANALYSIS (RCA) AND PROBLEM MANAGEMENT

Traditional RCA is highly laborious and time-consuming, requiring manual log analysis and correlation across multiple data sources. AI dramatically accelerates this process:

Data Correlation and Diagnosis: ML algorithms sift through vast amounts of historical data, logs, metrics, and error patterns to correlate disparate data points and pinpoint probable root causes faster and more accurately than manual methods.

GenAI Summarization: Generative AI solutions are vital here because they can synthesize large volumes of structured and unstructured data (such as logs, manuals, and historical tickets) to generate concise summaries of the root cause, providing guidance for resolution.

Agent Collaboration: Agentic AI is instrumental in identifying the root cause by retrieving topology information and log data from various sources (MELT sources through MCP servers), analyzing this data, and proposing resolution steps.

***Case Study Example**: Splunk integrates AI and ML to automate RCA across complex infrastructures, enabling customers to reduce RCA time by up to 50% and boost operational efficiency by 20% – SmartDev | Ngoc Nguyen, 2025.*

4. ENHANCED KNOWLEDGE MANAGEMENT AND DOCUMENTATION

AI turns AMS knowledge from a static archive into a living system that grows with every incident. Instead of relying only on manual write-ups, AI services analyse resolved tickets, chat transcripts, runbooks, code comments, and monitoring data to keep information aligned with the current state of the applications.

Dynamic capture: AI engines detect recurring symptoms, known errors, and successful fixes, then package them into structured entries with context, diagnostic steps, and remediation actions. Generative models can draft full articles from a single resolved case, while experts focus on review and approval rather than writing from scratch.

Smarter discovery: Semantic search and LLM assistants let engineers query the knowledge base in natural language and retrieve conceptually similar issues, not just keyword matches. They surface relevant patterns, candidate solutions, and runbooks that have worked in comparable situations, reducing time spent hunting for information.

Operational impact: When these capabilities are embedded into AMS workflows, organizations see faster resolution, less duplication of effort, and smoother transitions between teams and providers. Knowledge becomes a shared, continuously refreshed asset rather than a collection of static documents.

5. INTELLIGENT USER SUPPORT AND CONVERSATIONAL AI

AI turns the service desk into an always-on, conversational interface. Virtual agents powered by NLP become the front door for users: they understand free-form requests, detect intents, and guide people through tasks such as password resets, access requests, status checks, or simple troubleshooting. Available 24/7 across web, chat, and mobile, they resolve many routine questions without waiting for a human agent.

GenAI-assisted conversations: By combining LLMs with enterprise knowledge and context about the user, GenAI assistants deliver answers that reflect the user's environment instead of generic FAQs. They can ask clarifying questions, propose next steps, and escalate to a human with a concise summary when needed, cutting the time agents spend on repetitive exchanges.

Service performance: Organizations deploying these capabilities see tangible gains: a large share of level-1 demand can be absorbed by automation, response times fall from hours to minutes, and first-contact resolution rates increase. Zendesk, for instance, reports that its Answer Bot can handle up to 70% of routine tickets and contribute to a 25% uplift in first-contact resolution (SmartDev | Ngoc Nguyen, 2025). Over time, analysis of language, tone, and feedback across conversations also reveals sentiment trends and recurring pain points, helping AMS teams refine knowledge and move towards experience-based XLAs.

6. GENAI IN MAINTENANCE AND MODERNIZATION (CODE AND LEGACY SYSTEMS)

GenAI extends AMS capabilities into the application codebase, helping manage technical debt and supporting modernization efforts.

Code Assistance for Engineers: GenAI agents (Code Assistant) improve the productivity of support engineers through capabilities like code auto-completion, code translation, and code summarization.

Legacy Code Maintenance: GenAI models analyze legacy codebases to detect security problems, suggest security patches, and even rewrite old modules in modern languages. It helps maintain stability by generating comprehensive test cases and technical documentation for outdated applications.

Automated Patch and Update Management: In patching, AI replaces manual lists and rigid calendars with risk-based decisions. Models combine vulnerability feeds and asset inventories to flag the most exposed components and the updates that matter most. Instead of treating every patch as equal, platforms propose prioritized waves and trigger rollouts through DevOps pipelines. Microsoft Azure DevOps uses predictive analytics to orchestrate critical updates and help reduce incidents tied to unpatched systems – SmartDev | Ngoc Nguyen, 2025.

7. HYPERAUTOMATION AND ORCHESTRATION

Hyperautomation refers to the intelligent orchestration of multiple technologies – including RPA, AI, and iPaaS – to automate entire, complex workflows.

Cognitive Workflows: AI-first automation utilizes cognitive bots powered by AI analytics and recommendation engines. Automation agents can ingest tickets, infer which standard steps are required for resolution, and complete the workflow from end to end autonomously.

***Example**: For a retail client, a hyperautomation layer orchestrated workflows between SAP, ServiceNow, and Slack, using RPA bots to auto-restart failed jobs, which resulted in a 35% reduction in ticket volume within three months and 50% faster resolution time for L1 tickets – Accenture | Devansh Sinha, 2025.*

Agentic Collaboration: Agentic AI coordinates phased AMS transitions and orchestrates the actions of specialized AI agents for functions such as observability, reliability engineering, and service management, driving end-to-end decision automation.

OUTCOME

Business Benefits: from Efficiency to Value through Resilience

Explore the game-changing business impacts of AI-driven application management, where cost savings soar up to 50%, downtime plummets by 40%, and user satisfaction skyrockets, turning IT from a cost center into a strategic engine for innovation and growth.

The primary goal of integrating AI, Generative AI (GenAI), and Agentic AI into Application Management Services (AMS) – a shift often referred to as Intelligent Application Management (iAM) – is to move beyond mere operational maintenance and transactional efficiency to deliver measurable business value and strategic outcomes. This transformation enables organizations to maximize value, optimize costs, and accelerate speed simultaneously.

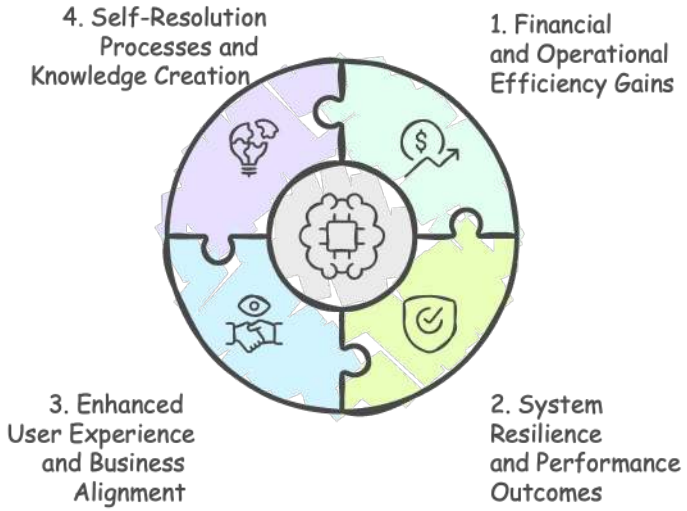

Delivering Measurable Business Impact with AI in AMS

The business value of AI in AMS is quantified through key performance indicators (KPIs) related to cost reduction, improved agility, enhanced system resilience, and superior user experience.

1. FINANCIAL AND OPERATIONAL EFFICIENCY GAINS

A core outcome of AI adoption is the significant improvement in productivity and efficiency, which directly translates into substantial cost savings.

TOTAL COST OF OWNERSHIP (TCO) AND COST REDUCTION

AI-powered AMS significantly reduces operational expenses by minimizing manual labor, optimizing resource allocation, and preventing costly downtime.

TCO Reduction: Accenture cites cases where redesigning AMS with hyperautomation, cloud optimization and integrated tooling delivers up to 50% lower TCO, 2025. Wipro's NextGen AMS, powered by the SPRINT framework, aims for a 35–40% TCO reduction – Wipro, 2025.

Incident Volume Reduction: Automation of routine tasks and proactive incident prevention lead to a high reduction in ticket volume. For instance, clients have reported a 35% reduction in ticket volume within three months through hyperautomation – Accenture | Devansh Sinha, 2025 – a 30–35% reduction in tickets through touchless operations – Wipro, 2025 – and up to a 40% reduction in incident volumes using IBM Watson AIOps – SmartDev | Ngoc Nguyen, 2025.

Operational Productivity: Organizations using AI in customer operations have reported productivity gains, including a 14% increase in issue resolution per hour and a 9% reduction in handling time per issue – SmartDev | Ngoc Nguyen, 2025. AIOps can boost productivity by 35–45% by automating routine work – Quinnox, 2024.

Resource Optimization: AI enables dynamic scaling of resources based on workload demands, ensuring optimal performance while minimizing costs.

ACCELERATED PROBLEM RESOLUTION

AI dramatically shortens the time spent diagnosing and fixing issues, enhancing service delivery efficiency.

Reduced Mean Time To Resolution (MTTR): AI-driven root cause analysis (RCA) correlates disparate logs and error patterns, minimizing MTTR. ServiceNow customers reported a 25% improvement in MTTR. Splunk's AI-powered RCA reduced incident analysis time by up to 50% – SmartDev | Ngoc Nguyen, 2025.

Faster Response Times: Conversational AI chatbots reduce response times from hours to mere minutes. Faster response times improve user satisfaction and trust.

Higher First Contact Resolution (FCR): AI-powered solutions, such as Zendesk's Answer Bot, have handled up to 70% of routine inquiries autonomously, resulting in a 25% increase in first-contact resolution rates – SmartDev | Ngoc Nguyen, 2025.

2. SYSTEM RESILIENCE AND PERFORMANCE OUTCOMES

AI-driven AMS shifts operations from reactive firefighting to proactive prevention, resulting in highly resilient applications and reduced downtime.

Improved Availability and Resilience: Predictive maintenance and self-healing systems drastically reduce unplanned downtime. Quinnox (2024) notes that organizations running AIOps platforms have seen unplanned downtime fall by 20–40%, while IBM cites around a 35% downtime reduction for clients using Watson AIOps – SmartDev | Ngoc Nguyen, 2025. Our ambition is to reach touchless operations with near-zero downtime.

Outage Reduction: Accenture reports that scaling automation and observability cuts outages by around 60% while sustaining 99.9% availability, 2025. Predictive AMS helped one financial institution cut downtime by 30% – YASH | Veera Venkata Naga Ravi Krishna Desaraju, 2025.

Proactive Issue Prevention: AIOps platforms, which enable predictive analytics, allow teams to forecast potential issues and act proactively. Quinnox (2024) finds that proactive, predictive incident practices can cut high-priority incidents by 15–45%.

Noise Reduction: AI-driven event correlation aggregates alerts to eliminate noise and false alerts. Quinnox (2024) finds that automated correlation can lower alert noise by up to 95%, while Wipro (2025) cites a 75% drop in alert volume.

3. ENHANCED USER EXPERIENCE AND BUSINESS ALIGNMENT

The strategic value of AI is moving AMS governance beyond technical SLAs to focus on measurable user and business outcomes.

SHIFT TO EXPERIENCE LEVEL AGREEMENTS (XLAs)

AI enables a focus shift from measuring technical SLA compliance to measuring Experience Level Agreements (XLAs).

User Satisfaction: AI-powered systems improve service quality and foster stronger customer loyalty. GenAI-enabled support has shown measurable improvements in customer satisfaction (CSAT) scores, with one case reporting a CSAT improvement from 4.0 to 4.5 – Xoriant | Sandeep Raje, 2024 – and a 15% enhancement in satisfaction scores – Cognizant | Stefano Montanari, 2025. Infosys anticipates a 25% NPS improvement through AI-first experiences – Infosys | Praveen Pejaver Rao & Vineeth Av, 2024.

Personalization: AI analyzes client usage trends to offer personalized recommendations and tailored services, enhancing the client experience.

Sentiment Analysis: AI tools automatically gauge user sentiment from chat interactions and ticket comments, providing measurable insights for service improvement.

BUSINESS ALIGNMENT AND STRATEGIC OUTCOMES

AI refocuses IT operations on business priorities, turning technology services from a budget line into a source of strategic value.

Business Observability: AI adds a business lens to operations by turning metrics, logs, and support interactions into decision-ready views tied to outcomes and KPIs. This directs effort toward the most critical services and limits customer-visible impact.

Accelerated Growth and Innovation: AI frees up AMS experts from routine tasks, allowing them to focus on complex problem-solving and strategic improvements like cloud migration and continuous modernization. This transformation supports strategic planning and continuous service improvement. GenAI shortens innovation experimentation cycle time by 50% – Infosys | Praveen Pejaver Rao & Vineeth Av, 2024.

Measurable Business Impact: Data-driven AMS solutions are expected to improve IT-business alignment by 30% – Cognizant | Stefano Montanari, 2025. In specific scenarios, AI-first insights helped prevent regulatory penalties amounting to millions of dollars for a health insurance payer, and reduced billing delays by 54% for a utilities company – Infosys | Praveen Pejaver Rao & Vineeth Av, 2024.

4. SELF-RESOLUTION PROCESSES AND KNOWLEDGE CREATION

GenAI and Agentic AI introduce novel benefits, particularly concerning knowledge creation and self-resolution capabilities:

Agentic Self-Resolution: Agentic AI-enabled frameworks aim to resolve over 50% of issues autonomously – ISG & HCLTech | Akhila H & Oliver Nickels, 2025. Wipro targets 85% of tickets resolved through self-heal/self-service – Wipro, 2025.

Knowledge Transfer Acceleration: GenAI automates documentation and knowledge capture. This capability streamlines troubleshooting and is reported to reduce time spent searching for solutions by up to 30% – SmartDev | Ngoc Nguyen, 2025.

Legacy System Maintenance: GenAI reduces the effort required to maintain outdated systems. One provider achieved a 40% reduction in average defect resolution time for a procurement software provider using GenAI-enabled sustenance services – Xoriant | Sandeep Raje, 2024.

Security and Compliance: AI-driven vulnerability management has led to a 40% reduction in security incidents related to unpatched vulnerabilities and a 35% improvement in update deployment speed – SmartDev | Ngoc Nguyen, 2025.

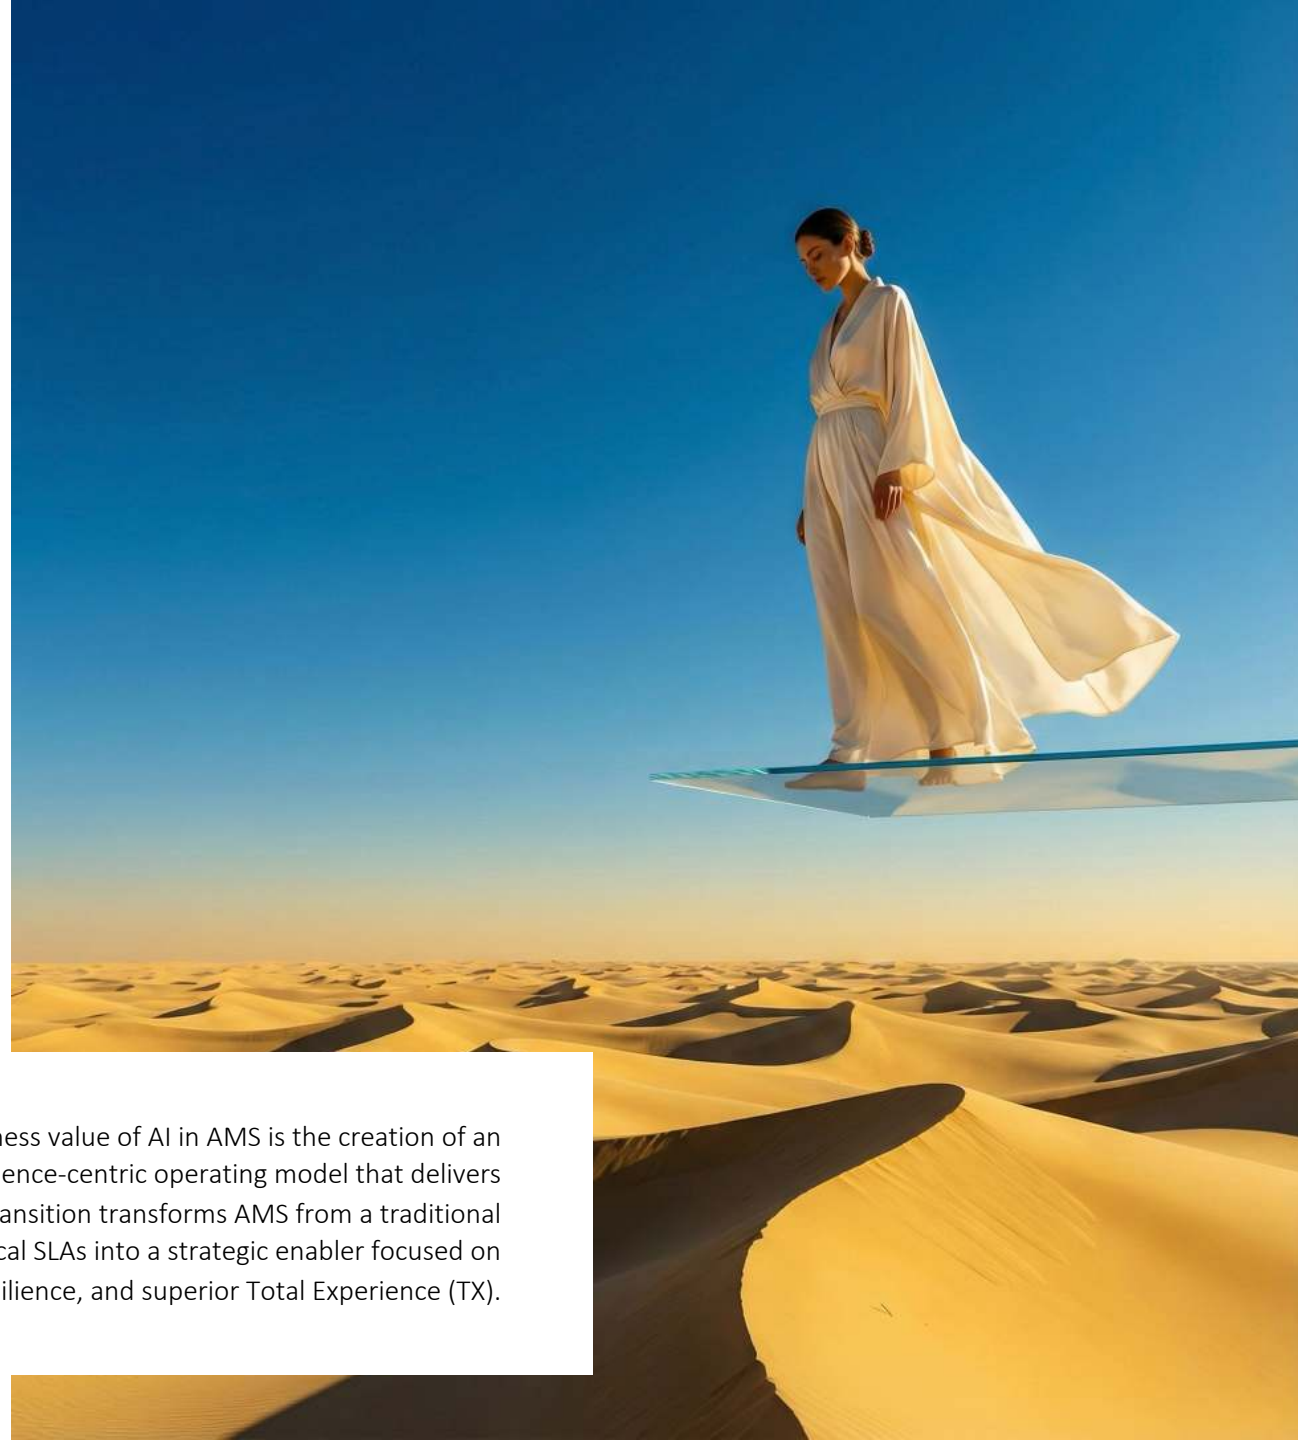

To summarize, the ultimate business value of AI in AMS is the creation of an intelligent, autonomous, and experience-centric operating model that delivers maximum business value. This transition transforms AMS from a traditional support function obsessed with technical SLAs into a strategic enabler focused on driving revenue, resilience, and superior Total Experience (TX).

FOUNDATION

Core Concepts and Technologies of Application Management Powered by AI

Dive into the foundational AI technologies – like machine learning, generative AI, and AIOps – that empower intelligent application management, enabling self-healing systems and autonomous agents to create a seamless, intelligent ecosystem that anticipates needs and drives superior outcomes.

Foundational and Emerging AI Shaping the Future of AMS

The integration of Artificial Intelligence (AI) and its associated technologies is fundamentally reshaping Application Management Services (AMS), moving the practice from a traditional, reactive, and manual model to an intelligent, proactive, and experience-driven framework.

The core concepts and technologies driving this transformation within AMS include foundational AI components, advanced platforms like AIOps and Observability, and the emerging capabilities of Generative and Agentic AI.

AI Technologies in Application Management Services

Characteristic	Core Function	AMS Application	Key Benefit
Machine Learning (ML)	Analyzes data, predicts outcomes	Predicts failures, classifies tickets	Faster root cause analysis
Natural Language Processing (NLP)	Understands user queries, interprets text	Automates ticket classification	Reduced manual effort
Generative AI (GenAI)	Generates new content, code, documentation	Automates knowledge creation	Increased efficiency
Agentic AI	Autonomous agents handle incidents	End-to-end decision automation	Autonomous issue resolution
AIOps	Applies AI to IT operations	Predictive insights, anomaly detection	Proactive problem-solving
Observability	Provides visibility across applications	Links technical performance to business KPIs	Business-driven prioritization
Hyperautomation	Automates workflows across systems	Automates incident triage and resolution	Reduced operational costs

1. FOUNDATIONAL AI TECHNOLOGIES AND CONCEPTS

Artificial Intelligence (AI) refers to computer systems capable of performing tasks that typically require human intelligence, such as learning, reasoning, and decision-making. These capabilities are built upon several key technologies:

Machine Learning (ML): ML enables systems to analyze vast datasets, identify recurring patterns, and predict outcomes without explicit programming. In AMS, ML models are trained on historical incident data and real-time system metrics. ML is crucial for predicting system failures, classifying tickets automatically, and accelerating Root Cause Analysis (RCA) by sifting through incident data to pinpoint probable causes faster than manual methods.

Natural Language Processing (NLP) and Conversational AI: NLP allows AI tools to understand user queries, interpret unstructured text data from tickets or emails, and determine user intent. This is essential for:

_Automated Ticket Classification: NLP-powered tools can interpret user-submitted tickets, classify issues accurately, and route them to the appropriate teams without human intervention.

_Intelligent Chatbots and Virtual Assistants: Conversational AI, enhanced by advancements in NLP, provides instant, 24/7 support by understanding user intent and responding with contextually relevant answers. These tools automate routine user queries, such as password resets and ticket status updates.

_Large Language Models (LLMs): LLMs form the basis of Generative AI (GenAI) solutions. They leverage deep learning techniques to analyze and process large volumes of data, which is crucial for achieving high efficiency and superior user experience in IT application management.

2. GENERATIVE AI (GenAI)

GenAI can generate brand-new, unique artifacts such as text, code, documentation, synthetic data, and workflows. It leverages Large Language Models (LLMs) which are essential for achieving high efficiency and superior user experience in IT application management.

GenAI's role in AMS includes:

Automated Knowledge Creation: Generating new knowledge documents, troubleshooting guides, or technical documentation based on solved problem tickets, user interactions, or code summaries. This creates "Living Documentation" that is continuously updated.

Code Assistance: Improving support engineer productivity through smart code hints, cross-language refactoring, plus brief function summaries. It helps manage technical debt by generating code summaries and documentation for legacy codebases.

Accelerating RCA: Generating concise summaries of the root cause by synthesizing large volumes of structured and unstructured data across logs, manuals, and historical tickets.

Conversational Interfaces: Powering intelligent chatbots that provide human-like responses, guidance on next best actions, and instant self-service.

3. AGENTIC AI: INTELLIGENT AND SPECIALIZED AI AGENTS

Agentic AI refers to intelligent, autonomous agents embedded across support workflows that proactively handle incidents, automate analysis, assist in resolutions, and continuously learn. The concept of Intelligent Application Management (iAM) is built upon a spectrum of agents:

Deterministic Agents: These are the rule-based workhorses that operate on predefined rules and workflows. They excel at predictable, repetitive tasks, such as monitoring against static thresholds or restarting failed services. Their primary role is to provide stability in well-defined scenarios.

Predictive AI Agents (Adaptive Intelligence): These agents leverage machine learning models to analyze data and make predictions, bringing context-aware intelligence. They predict system failures based on historical trends, allowing for proactive management, but they depend heavily on high-quality data.

Generative AI Agents (Agentic AI): These are the creative problem-solvers powered by LLMs. Agentic AI refers to intelligent, autonomous agents embedded across support workflows that proactively handle incidents, automate root-cause analysis, and continuously learn from operations. They don't just analyze data; they generate new content or solutions, such as generating automation scripts, curating resolutions from knowledge bases, or drafting solution proposals. This capability is driving the transformation towards autonomous IT operations.

In AMS, Agentic AI moves beyond simple automation to end-to-end decision automation, autonomously handling escalation, remediation, and system changes. This includes self-service, self-healing, and intelligent coordination to resolve more than 50% of issues autonomously – ISG & HCLTech | Akhila H & Oliver Nickels, 2025.

4. ADVANCED AI FRAMEWORKS

The convergence of these foundational technologies leads to comprehensive systems designed to manage complex IT environments:

AIOps AND OBSERVABILITY

The shift in AMS relies heavily on the capabilities of AIOps and Observability, particularly to enable proactive management:

AIOps (Artificial Intelligence for IT Operations): AIOps is defined as the use of AI/ML and data analytics within event management layers to boost, speed up, and automate manual tasks in IT operations processes. AIOps acts as the intelligence layer, turning observability data into actionable insights. Key capabilities include predictive insights, anomaly detection, noise reduction, event correlation, and automated remediation. The objective of AIOps is to create self-healing systems and enable predictive maintenance, moving application support from reactive problem-solving to proactive service delivery.

Observability: Provides continuous and complete visibility across applications, infrastructure, and services by collecting real-time operational data – specifically metrics, logs, and traces. Observability platforms must be continuously aware of the business context of the systems they monitor, linking technical performance to business metrics and KPIs (Key Performance Indicators). Cloud-Native Observability ensures a continuous stream of data to feed AIOps intelligence.

HYPERAUTOMATION

Hyperautomation blends RPA, AI, iPaaS and other automation tools into a unified fabric that spans applications, platforms and workflows, coordinating tasks, decisions and handoffs from end to end. In AMS, this fabric links ERP landscapes, ITSM platforms such as ServiceNow, and custom-built applications, so that incidents can be detected, triaged and resolved automatically, with steep drops in ticket volume and handling time.

5. APPLICATION AND USE CASES WITHIN AMS

These technologies are applied across the AMS value chain to achieve productivity, efficiency, and resilience goals. The ultimate goal of applying these core concepts is to enable support teams to focus on complex problem-solving and strategic improvements rather than routine diagnostics. The future of AMS is seen as a collaboration where AI agents and human expertise work in tandem to achieve unprecedented efficiency, enhancing productivity and service quality.

The transition to this AI-first model, however, faces challenges such as ensuring high data quality, managing integration complexities with legacy systems, addressing skill gaps, and mitigating risks associated with AI model bias, accuracy, and compliance.

In essence, AI technologies act as the digital nervous system for AMS. Predictive AI is the diagnostic tool anticipating illness, Agentic AI is the skilled surgeon performing automated repair, and Observability provides the continuous real-time vital signs of the application environment.

GENERATIVE AI & AGENTIC AI

Automated RCA: Generates concise root cause summaries from logs, knowledge, and historical tickets.

Code Assistant: Improves support engineer productivity through code auto-completion, translation, and summarization.

Living Documentation: Automates knowledge capture, creating dynamic, searchable documentation from tickets and conversations.

Legacy Code Maintenance: Generates code summaries, technical documentation, and test cases to reduce technical debt.

AIOps & PREDICTIVE AI

Proactive Issue Detection: Analyzes historical and real-time data to forecast system failures or bottlenecks, enabling teams to act before users are impacted (predictive maintenance).

Incident Clustering: Identifies patterns and groups similar issues to predict major incidents.

Self-Healing: Automatically diagnoses issues (like high CPU usage) and triggers remedial actions (like auto-scaling or DB reindexing) autonomously.

NLP & MACHINE LEARNING

Automated Ticketing: Streamlines incident management by classifying tickets, determining urgency, and routing them to the correct teams.

Service Desk Automation: Chatbots handle routine queries, reducing manual effort and workload on human agents.

ANALYTICS & METRICS

Experience-Driven Operations (XLA): Shifts governance focus from pure SLA compliance to measuring user sentiment and engagement, leveraging AI tools for sentiment analysis from user interactions.

Business Alignment: Correlates IT operation metrics with business KPIs (like revenue, cost, or customer churn) to prioritize issues based on commercial impact.

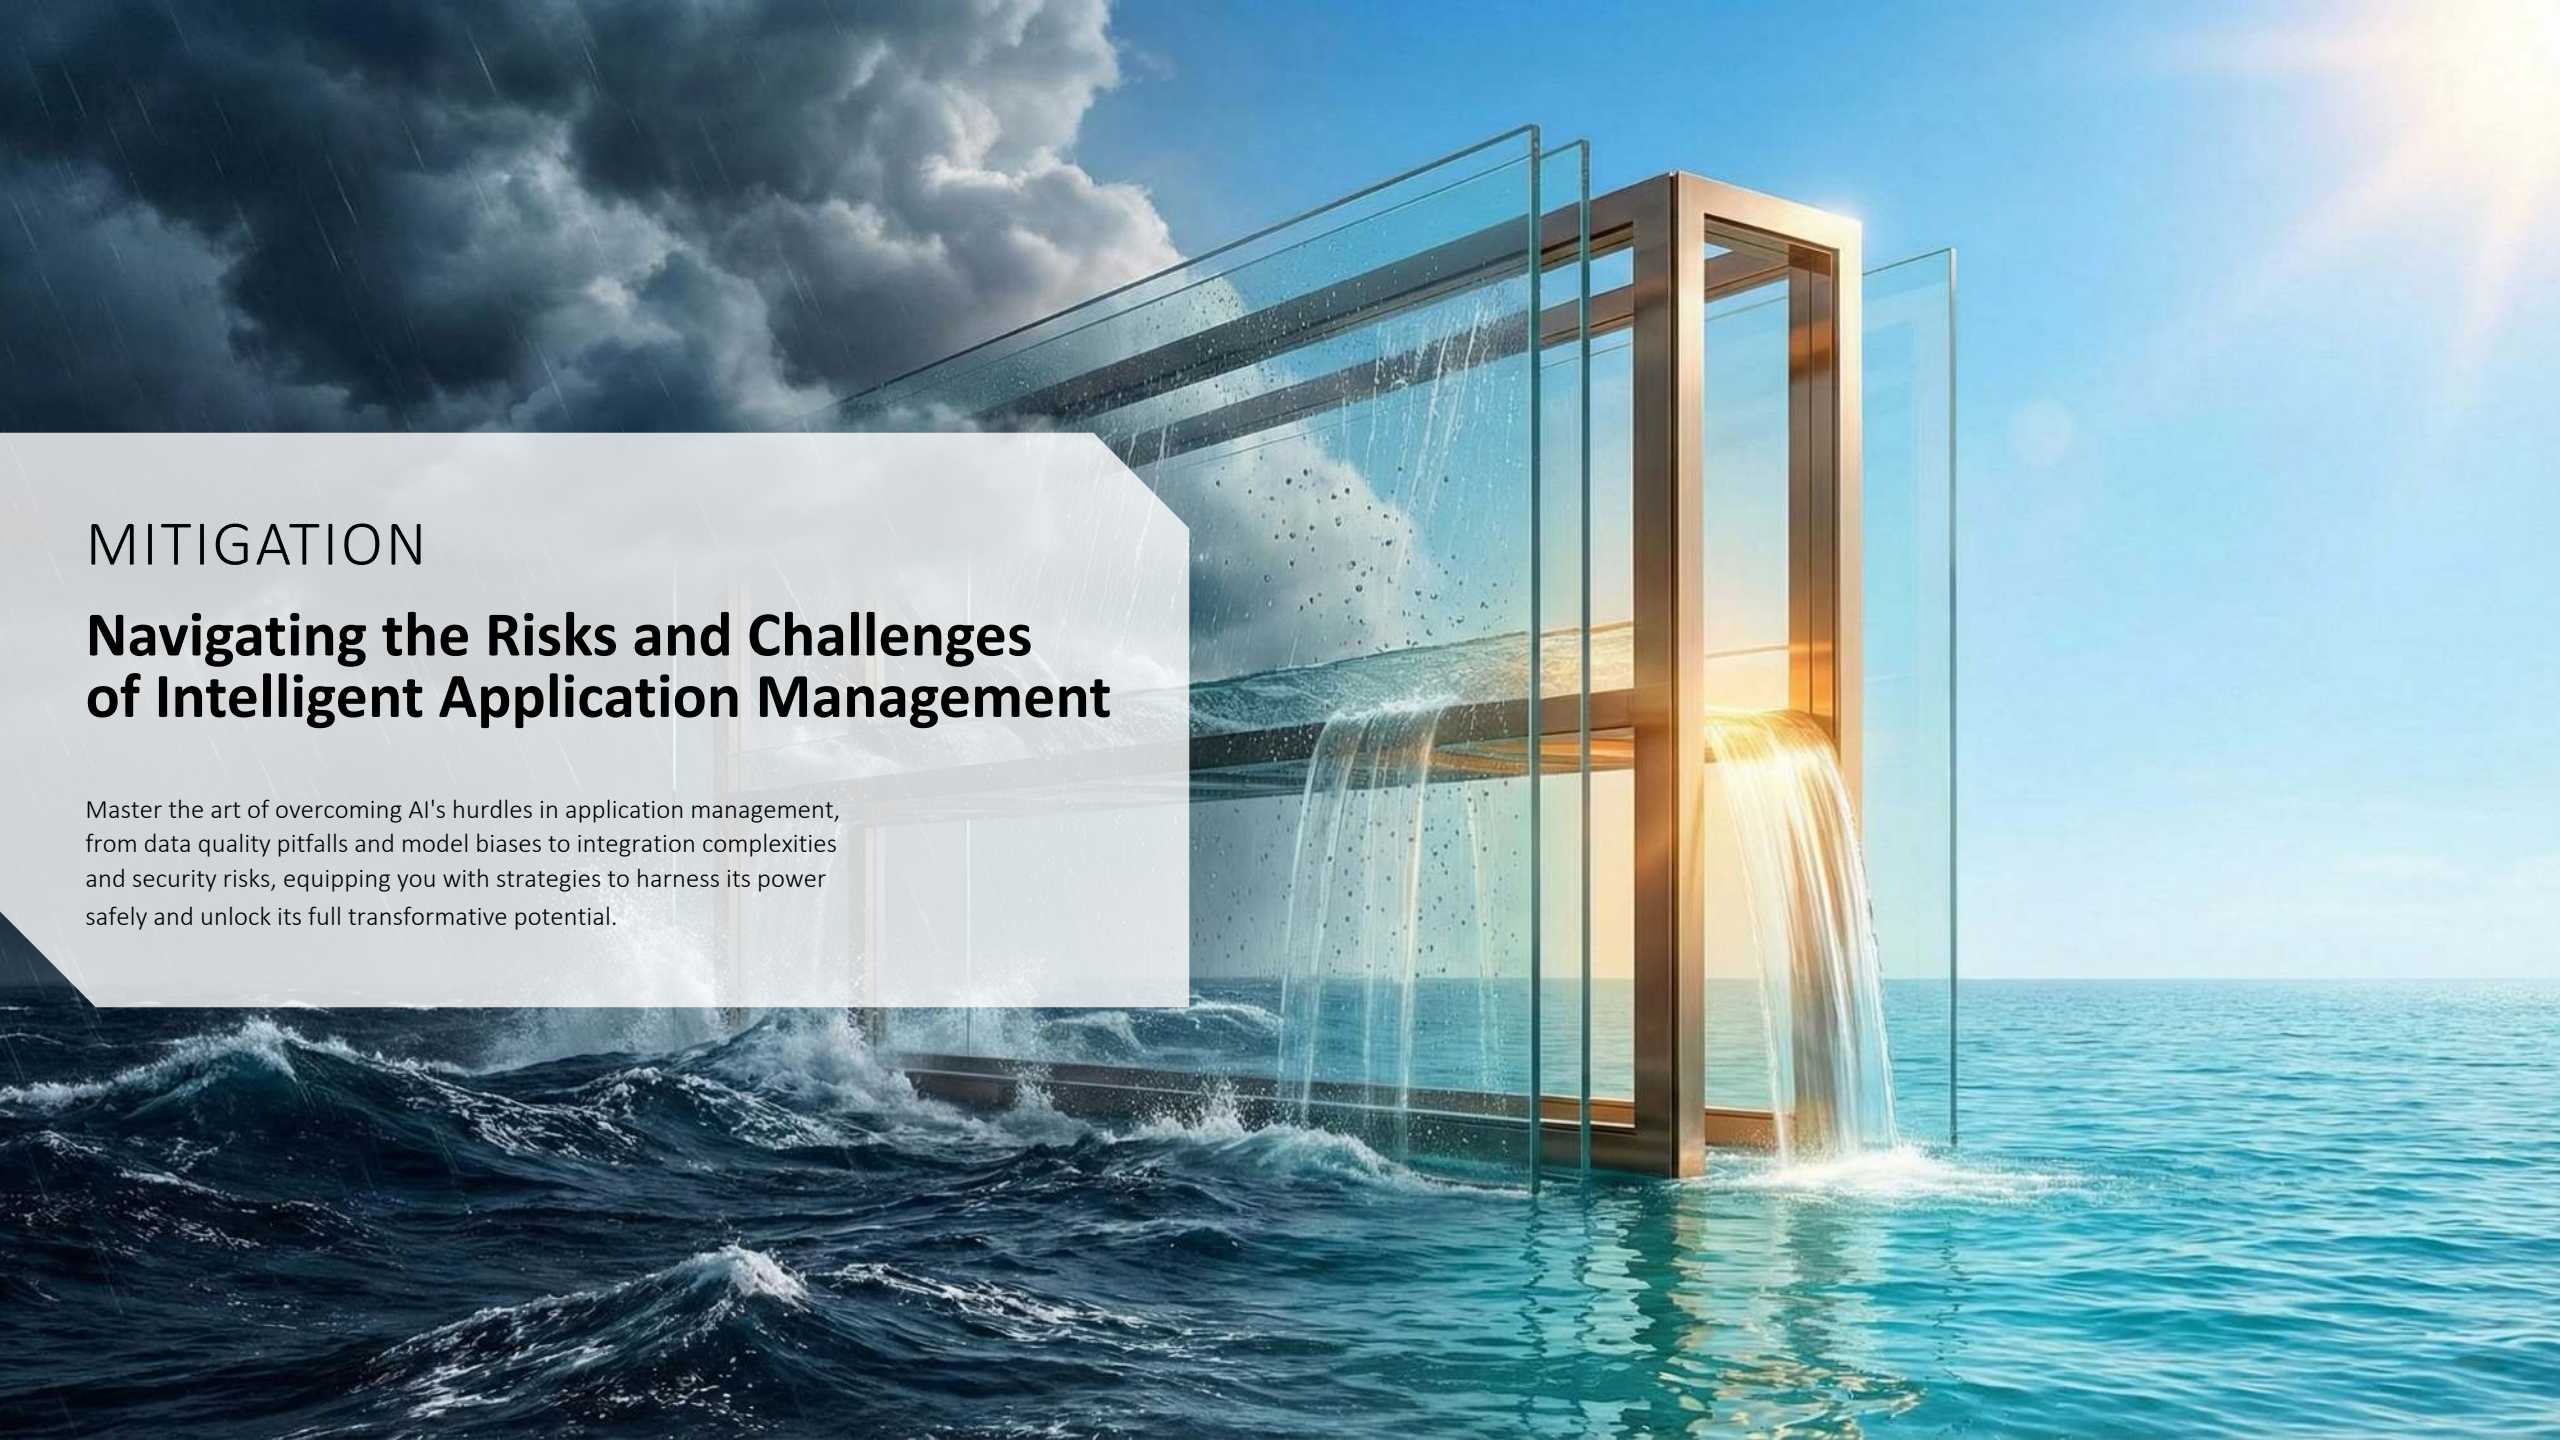

MITIGATION

Navigating the Risks and Challenges of Intelligent Application Management

Master the art of overcoming AI's hurdles in application management, from data quality pitfalls and model biases to integration complexities and security risks, equipping you with strategies to harness its power safely and unlock its full transformative potential.

Overcoming the Hidden Pitfalls of the iAM Transformation

The transition from Application Management Services (AMS) to Intelligent Application Management (iAM) is accompanied by substantial challenges and risks that must be carefully managed. These challenges span technical hurdles, operational complexities, security concerns, and human factors, ultimately threatening the realization of predicted benefits if not addressed proactively.

We classify these challenges into five major categories, including data quality, model design, application integration, change management and security.

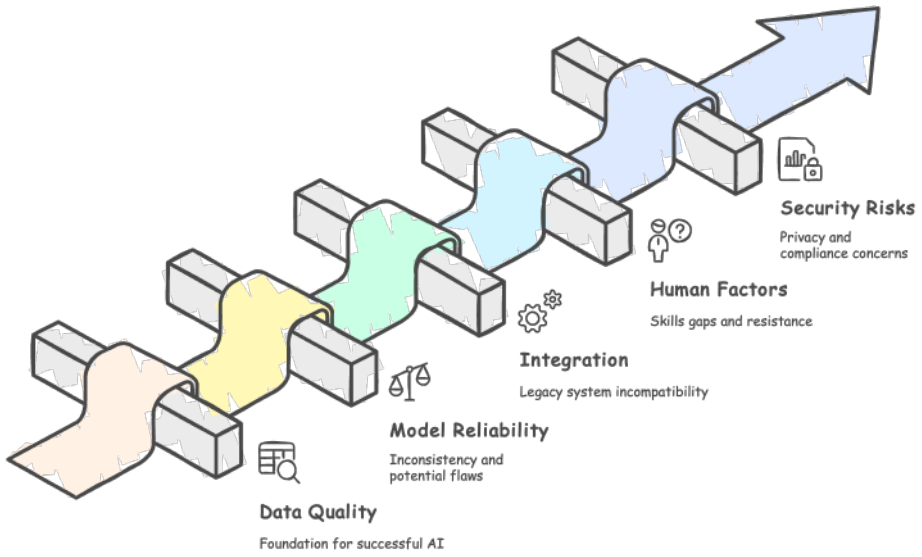

1. DATA QUALITY AND FOUNDATION COMPLEXITIES

Advanced AI agents are only as good as the data that feeds them, and AMS data is rarely in ideal shape. Operational information that matters for application support – tickets, logs, metrics, traces, CMDB records – is spread across tools and often incomplete, inconsistent, or poorly aligned.

Fragmentation: Different teams may use their own taxonomies, naming conventions, and severity scales, which makes it hard for AI to build a coherent picture of the environment. If this fragmented data is ingested as-is, models may miss early warning signs, misjudge impact, or draw misleading correlations between events.

Data foundations: Moving toward iAM therefore requires deliberate work on data foundations: rationalizing schemas and reference data across tools, cleaning and enriching key fields, consolidating telemetry into well-governed observability platforms, and setting up clear data ownership. Without this groundwork, predictive agents and GenAI copilots remain limited to narrow use cases or kept in “read-only” modes because stakeholders do not trust the underlying data.

Inference economics and cost volatility: Beyond technical hurdles, the shift to Agentic AI introduces a critical financial variable: the cost of inference. Unlike traditional fixed-cost software, LLMs and autonomous agents consume tokens for every reasoning step and data retrieval. Poor data quality or ambiguous instructions can lead agents into expensive "reasoning loops” – where the AI consumes vast amounts of tokens trying to make sense of fragmented data without reaching a resolution. To ensure the economic viability of iAM, organizations must implement strict "AI FinOps" mechanisms, monitoring token consumption and setting budget caps to prevent efficiency gains from being eroded by exorbitant operational costs.

2. MODEL RELIABILITY, ACCURACY, AND BIAS RISKS

Even with good data, AI models introduce a new set of risks around the quality and stability of their outputs. Recommendations are probabilistic and can be wrong in subtle ways: a misclassified incident, an overconfident root-cause suggestion, or an automation rule triggered on weak evidence.

Inherited bias: Models learn from past incidents, routing choices, and resolution behaviors. If the history reflects skewed priorities, mis-tagged tickets, or gaps in documentation, the model will inherit these patterns and amplify them. Certain incident types may be systematically under-prioritized, while noisy but low-impact alerts are escalated too often, diluting attention from the signals that really matter.

GenAI behaviors: GenAI capabilities add specific behaviors that need to be controlled. LLMs can generate responses that look plausible while being factually wrong, omit relevant constraints, or contradict previous guidance. Outputs can vary from one interaction to another, making it harder to reproduce decisions and explain why a certain recommendation was made.

Lifecycle management: To manage these risks, organizations must treat models as living assets. They require continuous evaluation, drift monitoring, and regular retraining. Explainability features, testing in shadow mode, and human-in-the-loop patterns help engineers understand when to follow or override AI recommendations and keep accuracy and fairness at an acceptable level.

3. INTEGRATION AND LEGACY SYSTEM CHALLENGES

Introducing iAM capabilities into existing landscapes is rarely a clean greenfield exercise. Enterprise portfolios contain a mix of cloud-native services, packaged applications, and heavily customized legacy systems. Many of these platforms were never designed to expose the detailed telemetry or APIs that AI tools expect.

Bridging the gap: Connecting AI to this reality often requires adapters, middleware, and custom data pipelines. Log formats must be normalized, events re-keyed, and relationships between components reconstructed before models can learn meaningful patterns. Where observability is limited, AI is effectively blind and can only operate on partial information.

Controlled rollout: Integration work also carries a risk of disruption. Changes to alerting, routing, and automation can impact on-call practices, escalation paths, and compliance controls. Without a staged approach – using shadow modes, pilots on limited scopes, and clear fallback options – organizations may struggle to scale iAM beyond a few isolated experiments.

4. HUMAN FACTORS, SKILLS GAPS, AND CHANGE MANAGEMENT

The shift from traditional AMS to iAM changes the day-to-day reality of support teams. Engineers are asked to work with AI copilots and agents, trust recommendations from systems they do not fully understand, and adjust long-standing practices. Without explicit attention to people aspects, resistance and skepticism are inevitable.

New skills: Teams need capabilities such as understanding how models behave, reading confidence signals, crafting effective prompts, and interpreting dashboards that present probabilistic insights rather than deterministic rules. Many AMS professionals are experts in specific technologies but have had limited exposure to data science or AI, which can create confidence gaps and slow adoption.

Change dynamics: Communication and involvement are therefore critical. Teams must see AI as a way to remove toil and elevate their work, not as a threat. Co-designing new workflows with practitioners, providing hands-on experimentation time, and recognizing new roles such as “AI-enabled SRE” help anchor the change, while clear guardrails avoid over-reliance on automation for complex or high-impact decisions.

5. SECURITY, PRIVACY, AND COMPLIANCE RISKS

Because iAM relies on access to operational and sometimes user-level data, it sits at the crossroads of security, privacy, and regulatory constraints. Logs, tickets, and monitoring events can expose configuration details, business transactions, and personal information that must be handled with care.

Regulatory constraints: When this data is ingested into AI platforms, organizations need to ensure that collection, storage, and processing comply with regulations such as GDPR, HIPAA, or CCPA. This includes strict access controls, encryption, clear retention rules, and auditability around who can see what and how data is used for training or inference. Multi-tenant or cloud-hosted AI services raise additional questions about data residency and cross-border transfers.

Agent exposure: Agentic and GenAI capabilities add further exposure. Agents that can read logs, generate scripts, or trigger actions in production environments must operate within tightly defined guardrails. Excessive permissions, weak segregation of duties, or poor prompt hygiene can be exploited to exfiltrate confidential information or run harmful commands. Security teams therefore need a precise view of how AI components connect to the rest of the stack.

Responsible use: Beyond classical security, organizations must manage copyright and IP risks linked to training data, avoid generating misleading content that could be perceived as authoritative, and document how AI is used in critical processes. Embedding security, privacy, and compliance requirements into the design and governance of iAM from day one is essential to balance innovation with acceptable risk.

IMPLEMENTATION

Foundational Pillars of AI-Driven Application Management

Embark on a strategic journey to build AI-first application management through essential pillars like robust data foundations, AIOps observability, and Agentic automation, paving the way for autonomous operations that deliver measurable business value and a superior total experience.

Designing the Path from Reactive AMS to Intelligent Operations

The shift from Application Management Services (AMS) to Intelligent Application Management (iAM) or AI-first application management requires a thoughtful, phased, and highly strategic approach. We outline detailed roadmaps, critical foundational steps, and strategic frameworks necessary to transition from traditional, reactive AMS to a proactive, intelligent service delivery model.

In a nutshell, the strategy for AI in AMS requires meticulous execution across four layers: establishing a clean data foundation, implementing AIOps and Observability, enabling intelligent self-service, and governing the entire process through experience-led outcomes (XLAs), all while managing the human and integration complexities inherent in dealing with legacy systems.

AI Implementation in Application Management

1	**Strategic Vision and Mindset Shift** Transition from technical compliance to business outcomes
2	**Foundational Implementation Pillars** Establish core pillars for AI-driven AMS
2.1	**Data Foundation Layer** Create a unified and high-quality data ecosystem
2.2	**AIOps & Cloud-Native Observability Layer** Integrate AIOps for proactive management
2.3	**Self-Service Knowledge Management Layer** Enable faster resolution and user self-service
3	**Operational Strategy: Agentic AI and Hyperautomation** Embed intelligent automation across AMS
4	**Organizational and Talent Strategy** Develop new skills and manage change

1. STRATEGIC VISION AND MINDSET SHIFT

The strategy for adopting AI in AMS begins with a fundamental change in mindset, moving away from being focused purely on technical compliance to delivering measurable business outcomes and superior user experiences.

THE "AI-FIRST" AND "EXPERIENCE-LED" MANDATE

AI integration is no longer optional but a critical necessity for organizations seeking to excel in complex application environments. This paper aligns with the AI-first and experience-led positioning seen in recent AMS frameworks from Infosys and HCLTech.

Reinventing the Operating Model: Enterprises must move from function-based, reactive operations to Product-Aligned Operating Models (PAOM), as described for example in HCLTech's research on AI-led operating models. In this structure, AMS teams are organized around digital products and value streams, linking day-to-day support to business results, key journeys, experience goals, and cost-to-serve indicators.

Shifting Governance (SLA to XLA): The strategy must pivot governance from traditional Service Level Agreements (SLAs) to Experience Level Agreements (XLAs). XLAs measure outcomes in terms of user sentiment and engagement, ensuring service managers focus on making the user experience seamless and productive, rather than just closing tickets on time.

Goal of Autonomous Operations: The ultimate strategic goal is to establish self-healing systems, achieve zero-touch automation, and target autonomous operations capable of settling more than 50% of tickets via AI-driven engagement, self-service agents, and predictive operations – a target widely referenced in provider frameworks (e.g. Infosys, HCLTech, Wipro) for mature AI-first AMS models.

A PHASED, CONSULTING-LED APPROACH

Successful AI deployment requires a structured, step-by-step strategy to ensure smooth adoption and maximum benefits. Infosys, Cognizant, and other providers utilize multi-phased frameworks (like Cognizant's ADAPT methodology – Cognizant | Stefano Montanari, 2025) to manage transition and transformation.

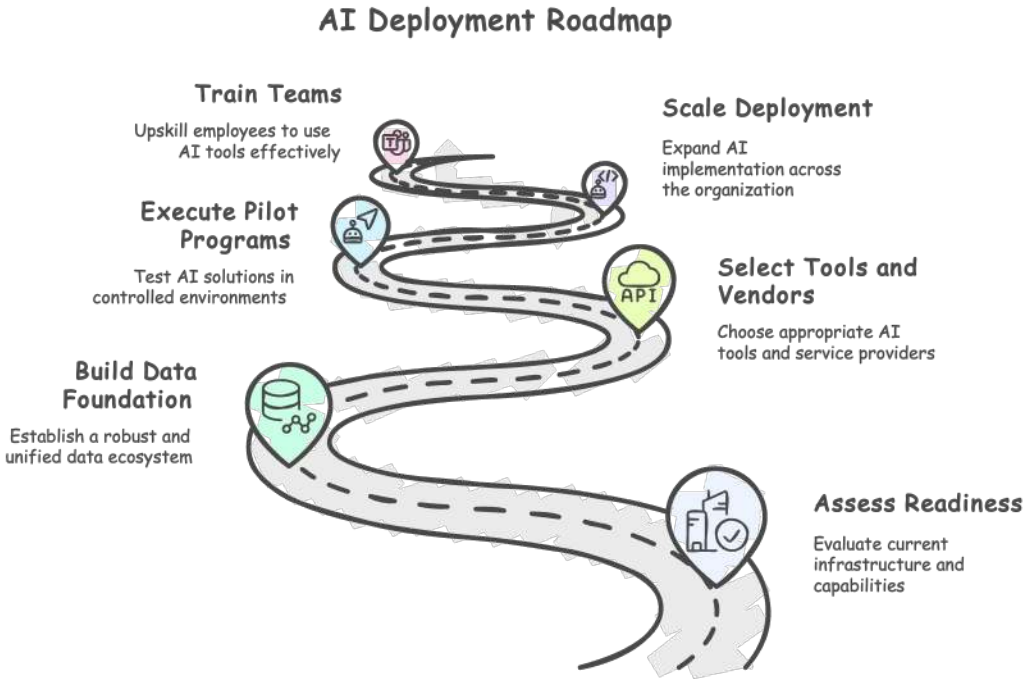

PILLAR 1
DATA FOUNDATION LAYER

Data is the cornerstone of any successful AI implementation. Without a robust and high-quality data foundation, AI algorithms will produce inaccurate predictions or flawed recommendations, leading to risk amplification.

Unified Ecosystem: Strategy requires creating a unified, high-quality data ecosystem that acts as the single source of truth for all analytics, AI, and governance activities. This involves consolidating all tickets and data streams (from end-users, system logs, and collaboration channels) into a single ITSM system like ServiceNow.

Data Collection and Standardization: Organizations must systematically collect diverse datasets, including error logs, performance metrics, and historical ticket data. Processes must be established to clean and standardize this data, removing inconsistencies that could skew AI model training and outputs.

Addressing Data Challenges: This phase is complex and time-consuming, often requiring significant investment in data integration platforms and data governance policies to ensure data integrity, privacy, and security.

PILLAR 2
AIOps & CLOUD-NATIVE OBSERVABILITY LAYER

This layer provides the continuous visibility and intelligence needed for proactive management.

Observability: Implementation must focus on Cloud-Native Observability (Lever 1), which involves continuously collecting real-time operational data: metrics, logs, and traces. This proactive monitoring feeds the AIOps engine.

AIOps Integration: AIOps (Lever 2) acts as the intelligence layer, using machine learning to detect anomalies, correlate related alerts (reducing alert noise by up to 95% – Quinnox, 2024), and predict potential issues. The strategy must include integrating AIOps platforms, such as IBM Watson AIOps, with existing monitoring and ITSM tools.

Self-Healing Capabilities: The deployment of AIOps must aim for capabilities like automatically detecting database degradation and triggering DB reindexing or identifying high CPU usage and initiating auto-scaling.

PILLAR 3
SELF-SERVICE KNOWLEDGE MANAGEMENT LAYER

The goal is to make knowledge dynamic and context-aware, enabling faster resolution and user self-service.

AI-Generated Knowledge: Strategy should utilize GenAI to automatically curate, update, and recommend relevant knowledge articles, creating "living documentation" from resolved incidents and problem records.

Conversational AI: Deploying virtual agents or chatbots (e.g., built using Microsoft Copilot Studio or Zendesk's Answer Bot) to provide guided troubleshooting steps and instant answers is key.

2. FOUNDATIONAL IMPLEMENTATION PILLARS FOR BUILDING AN AI-DRIVEN AMS FRAMEWORK

3. OPERATIONAL STRATEGY: AGENTIC AI AND HYPERAUTOMATION

The execution strategy involves embedding intelligent automation across the entire AMS value chain, moving from basic automation to cognitive workflows.

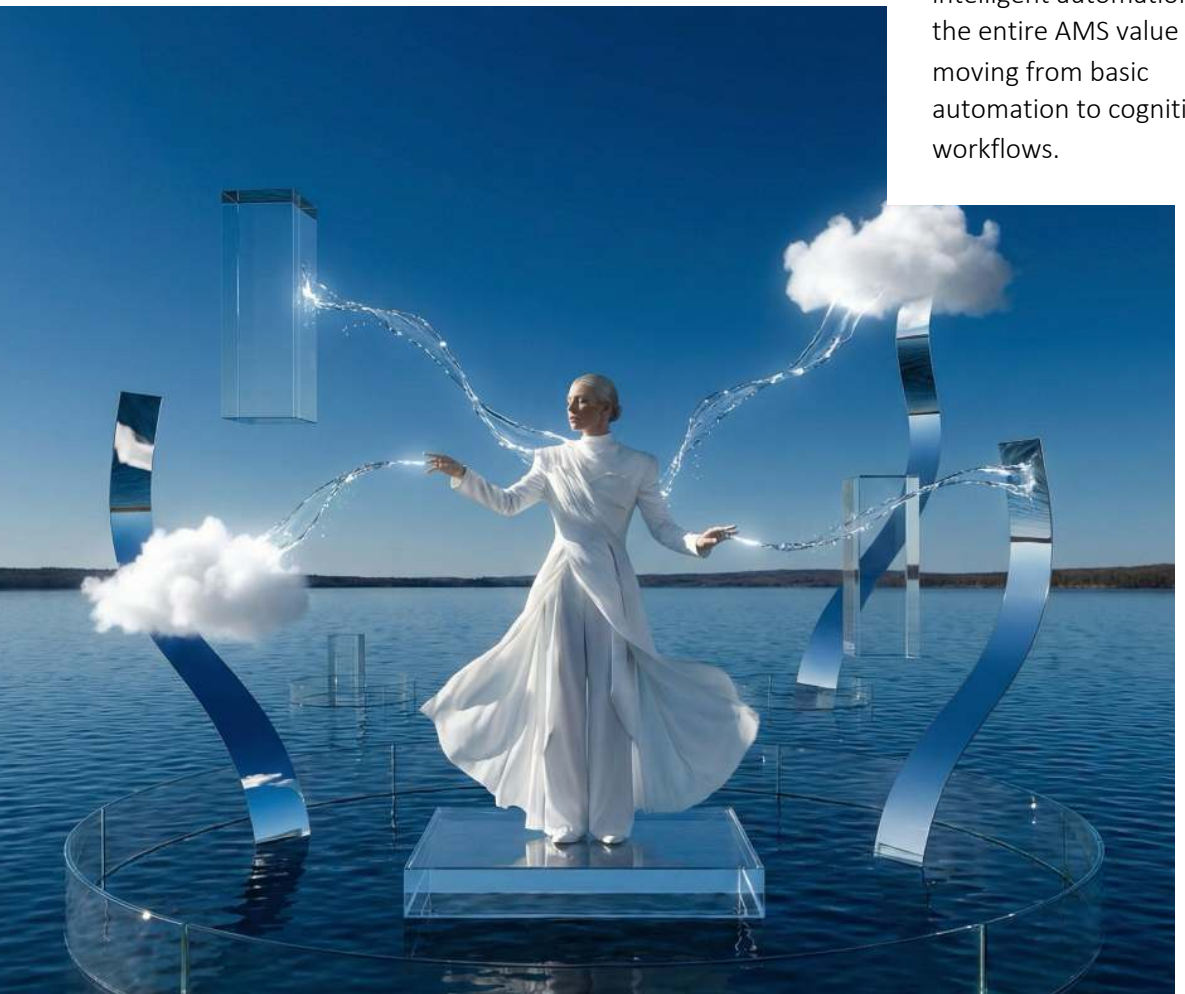

Hyperautomation Orchestration: This strategy involves leveraging hyperautomation, which intelligently automates workflows by combining RPA, AI, and iPaaS orchestration. This layer connects tools like ERP (e.g., SAP), ITSM (e.g., ServiceNow), and custom applications to automate incident triage and resolution.

Agentic AI Deployment: Following the taxonomy used by Quinnox and ISG/HCL for iAM, this paper discerns three families of agents across the AMS value chain:

_Deterministic agents, implementing codified rules and repeatable workflows for stable, well-understood tasks (threshold-based alerts, standard operating procedures, scheduled jobs).

_Predictive AI agents, leveraging machine-learning models to detect anomalies, anticipate capacity or performance issues and recommend preventive actions.

_Generative AI agents, using LLMs to summarize tickets, propose remediation steps, generate scripts or playbooks and support conversational interactions.

End-to-End Value Stream Instrumentation: HCLTech suggests mapping the end-to-end AMS value stream – from ticket prevention to resolution and feedback loop – and inserting AI at each point to maximize measurable leverage.

_Ticket Prevention: Using full-stack observability and AI to spot patterns before users are affected.

_Triage and Assignment: Using Natural Language Understanding (NLU) to classify intent, route to the right resolver group, and suggest playbooks.

_Resolution and Execution: Agentic workflows propose and execute fixes in bounded cases.

Risk Mitigation via Guardrails: Given the risks associated with GenAI (hallucinations, bias, security), strategic deployment requires AI guardrails. Infosys mandates the use of a Responsible by Design AI framework (Infosys, 2024) to systematically evaluate solutions against ethical, regulatory, legal, and internal policies and governance rules before they are put into production. Human-in-the-loop systems must validate critical decisions, ensuring that agents operate with explicit, confidence-based autonomy tiers.

4. ORGANIZATIONAL AND TALENT STRATEGY

The organizational strategy must focus on enabling the workforce to partner effectively with AI.

Addressing the Skill Gap: The transition demands that support teams develop new skills in data literacy, AI tools, and interpreting machine-generated insights. Investment in training and upskilling is necessary to ensure employees understand AI's capabilities and how to leverage them effectively.

Change Management: Organizations must proactively manage resistance and fear of job displacement by emphasizing that AI augments human effort, handling routine tasks while freeing human talent for complex problem-solving, strategic planning, and innovation.

Partnerships: Forging the right partnerships with industry-leading solution providers, hyperscalers (like Azure OpenAI, Google Vertex AI), and AI infrastructure OEMs is crucial to access cutting-edge technologies and accelerate time to value. Companies like Infosys (with their LEAP platform – Infosys | Praveen Pejaver Rao & Vineeth Av, 2024) and HCLTech (with AI Force and Agentic Studio – ISG & HCLTech | Akhila H & Oliver Nickels, 2025) rely on industrialized frameworks and platforms to structure these transformations.

Phased Rollout: Beginning with pilot programs on specific, high-impact tasks (like automated password resets or categorizing incoming tickets) in a controlled environment is critical for validating the tools and mitigating risk before gradually scaling adoption across all operations.

CONCLUSION

From Managed Services to Silent Intelligence

Application Management Services are at an inflection point. For years, AMS has been organized around tickets, SLAs, and cost containment. This paper has shown how **AI, Generative AI and Agentic AI** shift this paradigm toward **Intelligent Application Management (iAM)**: an operating model where applications are continuously observed, risks are anticipated, and remediation is increasingly autonomous and invisible to end users. Value no longer comes from adding operational "noise", but from cultivating a layer of silent intelligence embedded in the IT estate.

The promise of iAM is **already tangible**. The combination of **AIOps, observability, GenAI assistants and Agentic automation** delivers reduced total cost of ownership, fewer high-priority incidents, less downtime, and higher productivity and satisfaction. More importantly, it **changes the role of operations: from "keeping the lights on" to quietly safeguarding revenue, protecting brand, and enabling digital growth** – without demanding constant attention from users or business stakeholders.

Four shifts characterize this transformation in modern AMS operating models:

From reactive firefighting to predictive resilience. AIOps and predictive agents continuously monitor the digital nervous system, correlate weak signals, and prevent failures before they hit the business. Self-healing turns incident response into zero-touch services, where the most critical incidents are those that never surface because they are resolved silently.

From SLA compliance to experience-driven operations. XLAs and sentiment analytics refocus governance on what users feel and achieve, not just on uptime and ticket metrics. The best service is often the one that users barely notice, because disruptions are rare and friction is low.

From static documentation to living knowledge and self-resolution. GenAI and Agentic AI transform interactions and logs into structured, reusable knowledge. Copilots assist users and engineers, while autonomous resolution grows. Knowledge works in the background, shortening paths to resolution and reducing the need for visible escalation.

From operational cost to growth platform. By eliminating toil and automating routine tasks, iAM frees experts to focus on modernization, cloud migration, and optimization of critical journeys. The operational landscape becomes less cluttered, making room for focused, high-impact change.

This vision is **demanding**. Fragmented data, probabilistic model behavior, legacy integration, skills gaps, and security, privacy and compliance constraints can all derail the journey. Intelligent Application Management is therefore as much about data engineering, architecture, governance and risk management as it is about AI. **Silent intelligence must not be opaque intelligence**: transparency, control and human-in-the-loop design remain essential.

The way forward requires an **AI-first, experience-led ambition for AMS**, a solid backbone of **data, ITSM and observability**, a **GenAI-powered knowledge layer**, and a **pragmatic roadmap for Agentic automation and hyperautomation**. It also requires investment in people and operating models, so that automation can run in the background under clear and trusted guardrails.

The most advanced iAM models will be judged less on algorithmic sophistication than on the **simplicity of the experience** they create. **Silent intelligence is this simplicity made operational**: users no longer see the complexity of the stack or the orchestration of agents that keep it healthy. They simply experience fast, reliable, intuitive applications that evolve at the speed of the business. For CIOs and AMS leaders, the challenge is no longer **whether to embrace this shift**, but **how quickly they can industrialize it, responsibly and at scale** – before a noisy, fragile estate becomes the loudest signal that silent intelligence is missing.

GLOSSARY

A Guided Tour of Silent Intelligence in Application Management

ADAPTIVE INTELLIGENCE (PREDICTIVE AI AGENTS)

Predictive AI agents learn from historical incidents, logs, and metrics to anticipate failures, performance degradation, or demand peaks. In iAM, they continuously scan the “nervous system” of applications to flag weak signals and propose preventive actions before users notice any impact. They sit between static automation and fully autonomous agents, providing context-aware recommendations that operations teams can trust.

AGENTIC AI

Agentic AI designates autonomous, goal-driven software agents embedded across AMS workflows. In iAM, these agents not only analyze data but decide on and execute remediation steps end-to-end: gathering evidence, correlating symptoms, applying fixes, and validating outcomes. Their ambition is to resolve a majority of incidents without human intervention while still operating under explicit guardrails and human oversight.

AIOps

AIOps (Artificial Intelligence for IT Operations) is the intelligence layer that ingests massive volumes of observability data – metrics, logs, traces – and turns them into actionable insights for operations teams. It automates anomaly detection, alert de-duplication, event correlation, and root cause hypotheses, enabling a shift from reactive firefighting to proactive, self-healing operations. Within iAM, AIOps is the analytical core on which predictive and Agentic behaviors depend.

APPLICATION MANAGEMENT SERVICES (AMS)

AMS refers to the traditional managed services responsible for maintaining, supporting, and evolving enterprise applications. Historically, AMS has been ticket-driven, reactive, and cost-focused. Intelligent Application Management re-imagines AMS as a value-centric, AI-powered capability that maximizes resilience, speed, and user experience – not just cost savings.

AUTONOMOUS RESOLUTION RATE

The autonomous resolution rate measures the proportion of issues resolved without human intervention, through self-service, self-healing, or fully agent-driven workflows. It is a signature KPI of iAM, indicating how far an organization has progressed from manual ticket handling to Agentic operations. Ambitious iAM strategies often target a majority of incidents to be resolved autonomously while reserving experts for high-impact and complex cases.

BUSINESS OBSERVABILITY

Business observability links technical signals (alerts, performance metrics, incident trends) to business outcomes such as revenue, conversion, or regulatory risk. Instead of optimizing infrastructure in isolation, iAM uses business observability to prioritize work on what truly matters: the applications and user journeys that carry the most business value. It is the key to aligning AMS with executive-level objectives.

CODE ASSISTANT

Code Assistant denotes a GenAI-powered assistant for application engineers and L2/L3 support teams. It can read and summarize unfamiliar code, propose fixes, translate legacy languages to more modern ones, and generate targeted test cases. Within iAM, Code Assistant significantly shortens the feedback loop between production incidents and code-level remediation, effectively shrinking technical debt over time.

CONVERSATIONAL AI

Conversational AI (Virtual Agents) provides natural-language interfaces – chatbots or voicebots – that understand user intent, answer questions, troubleshoot, and create or update tickets. In iAM, these agents serve as the “front door” of support, resolving a large share of routine requests instantly and collecting high-quality context for the remaining ones. They are also a primary source for experience and sentiment data used in XLAs.

DIGITAL NERVOUS SYSTEM

The digital nervous system is a metaphor for the continuous flow of data – metrics, logs, traces, events, tickets – connecting every component of the application landscape. Observability and AIOps together form this nervous system for iAM: predictive agents act as the diagnostic layer, Agentic AI plays the role of the "surgeon," and governance defines which reactions are allowed in which contexts.

DETERMINISTIC AGENTS

Deterministic agents are rule-based bots that follow predefined runbooks and SOPs. They excel at stable, repetitive tasks such as restarting services, cleaning queues, or executing standard maintenance steps. In an iAM architecture, deterministic agents form the reliable base layer on top of which predictive and generative agents orchestrate more complex, adaptive workflows.

EXPERIENCE-DRIVEN OPERATIONS

Experience-driven operations place user perception and business impact at the center of AMS governance. Instead of optimizing solely around uptime or ticket volumes, iAM measures how users actually feel and perform when using applications – through sentiment analysis, behavior analytics, and survey signals – and uses these insights to drive prioritization, automation, and investment decisions.

EXPERIENCE LEVEL AGREEMENT (XLA)

An Experience Level Agreement is a contract or internal target focused on user experience – speed, ease of use, emotional response – rather than only technical metrics like response time or incident counts. In an iAM model, XLAs complement or partially replace SLAs, capturing satisfaction, adoption, and productivity as primary measures of success. They are fed by AI-derived sentiment and behavior analytics.

GENERATIVE AI (GenAI)

Generative AI creates new artifacts – text, code, documentation, synthetic data – based on patterns learned from large datasets. In Intelligent Application Management, GenAI is used to draft incident summaries, propose remediation steps, generate or refactor code, and produce "living" technical documentation that evolves with every solved case and conversation. It is the creative engine behind many Agentic capabilities.

GUARDRAILS

AI guardrails are the explicit constraints and policies that limit what AI agents are allowed to do in production. They cover which systems can be modified, which actions require human approval, how changes are logged, and how risky operations are simulated before being executed. In iAM, robust guardrails are essential to manage hallucinations, bias, and stochastic behaviors of GenAI while preserving trust in autonomous operations.

HYPERAUTOMATION

Hyperautomation is the orchestration of multiple automation technologies – RPA, AIOps, GenAI, integration platforms – across end-to-end workflows. Instead of automating isolated steps, iAM uses hyperautomation to handle entire incident lifecycles: detection, triage, diagnosis, remediation, validation, and communication. The result is significant reductions in ticket volume, cost, and time-to-resolution.

INCIDENT CLUSTERING AND CORRELATION

Incident clustering groups similar events or tickets based on patterns in their symptoms, contexts, or impacted components. Correlation then links these clusters back to shared underlying causes. In iAM, this capability is central to AIOps: it compresses thousands of noisy alerts into a handful of meaningful "situations," allowing agents and humans to focus on what truly matters and to predict emerging major incidents early.

INTELLIGENT APPLICATION MANAGEMENT (iAM)

Intelligent Application Management is the evolution of classic AMS into an AI-powered, agent-orchestrated, experience-centric operating model. It combines AIOps, observability, GenAI, Agentic automation, and hyperautomation to move from “keep the lights on” support to proactive, autonomous, and business-aligned operations. Its ambition is to relax the usual tension between business value, cost efficiency and delivery speed in IT operations, so that improvements on one dimension do not degrade the others.

LIVING DOCUMENTATION

Living documentation refers to knowledge assets – KB articles, troubleshooting guides, architecture notes – that are continuously regenerated and enriched by AI as new incidents are resolved and new patterns discovered. In an iAM context, GenAI agents assemble this documentation automatically from tickets, logs, chat transcripts, and code repositories, ensuring that operations teams always have up-to-date, searchable knowledge without separate manual effort.

OBSERVABILITY

Observability is the practice of instrumenting systems so that their internal state can be inferred from external outputs such as metrics, logs, and traces. For iAM, observability is the raw sensory layer: it feeds AIOps and agents with the continuous data needed for anomaly detection, RCA, and predictive maintenance. Mature observability explicitly ties technical signals to services, user journeys, and business KPIs.

PREDICTIVE MAINTENANCE

Predictive maintenance uses ML models and AIOps to forecast failures and performance degradation before they impact users. In Intelligent Application Management, this means identifying slow-growing memory leaks, capacity bottlenecks, or unusual transaction patterns early and triggering preventive actions – from scaling and configuration changes to code fixes – thus reducing downtime and high-priority incidents.

ROOT CAUSE ANALYSIS (RCA)

Root Cause Analysis seeks to identify the underlying technical or process cause behind incidents or clusters of issues. iAM leverages AI to automate much of this work: correlating logs and metrics, reconstructing timelines, and generating concise RCA narratives for humans to validate. Agentic RCA dramatically shortens MTTR and feeds higher-quality insights back into automation runbooks and living documentation.

SELF-HEALING

Self-healing describes the ability of systems and agents to detect anomalies and autonomously trigger corrective actions, restoring normal operation without human intervention. In iAM, self-healing may cover everything from restarting failed services and applying known patches to executing complex remediation playbooks orchestrated by Agentic AI. The objective is a “silent” reduction in outages and user-visible incidents.

SILENT INTELLIGENCE

Silent intelligence is the principle that the most advanced AI in operations should be largely invisible to end users. They experience faster, more stable, and more intuitive applications, but the algorithms, models, and agents remain in the background, amplifying human expertise rather than competing with it.

TOTAL EXPERIENCE (TX)

Total Experience shifts attention from isolated end-users to the combined journeys of everyone involved: customers, support staff, business leaders, and ecosystem partners. Under iAM, TX becomes the ultimate outcome: smoother workflows for operations teams, resilient and performant applications for business stakeholders, and effortless digital journeys for customers and employees. AI and agents are judged not only on efficiency gains but on their contribution to this shared experience.

INSPIRING SOURCES

Accenture | Devansh Sinha. (2025, September 29). Hyperautomation in AMS: The New Frontier of Efficiency.
https://www.linkedin.com/pulse/hyperautomation-ams-new-frontier-efficiency-devansh-sinha-ixqie

Accenture | Devansh Sinha. (2025, August 7). Rethinking AMS Transitions: The Next-Gen Model Powered by GenAI and Agentic AI.
https://www.linkedin.com/pulse/rethinking-ams-transitions-next-gen-model-powered-genai-devansh-sinha-qocgf

Accenture | Devansh Sinha. (2025, May 30). Reimagine AMS: Leveraging Intelligent Agents for Proactive and Efficient Application Management.
https://www.linkedin.com/pulse/reimagine-ams-leveraging-intelligent-agents-proactive-devansh-sinha-1fjqc

Accenture. (2025). Run to new: Outpace change with value from gen AI and cloud.
https://www.accenture.com/us-en/services/cloud/application-managed-services

Ahilan Ponnusamy. (2024, December 12). Enhancing Stability and Maintainability of Legacy Apps With Gen AI.
https://medium.com/@ahilanp/enhancing-stability-and-maintainability-of-legacy-apps-with-gen-ai-4e35bc382671

Capgemini | Maxime Girardeau & Sharp Delusion. (2025). L'IA Invisible.
https://www.capgemini.com/fr-fr/perspectives/publications/ia-invisible

CGI | Chandra Guntury. (2025, March 26). The Future of AMS – Key Transformational Aspects.
https://www.linkedin.com/pulse/future-ams-key-transformational-aspects-chandra-guntury-t0rxc

Coforge | Gaurav Gupta. (2025, November 20). Reimagining AMS: How AI is Transforming Application Managed Services into Experience Driven Operations.
https://www.linkedin.com/pulse/reimagining-ams-how-ai-transforming-application-managed-gaurav-gupta-cbk6c

Cognizant | Stefano Montanari. (2025, January 21). Next-Gen Application Management Services Transforming RCG Industry.
https://www.cognizant.com/dk/en/insights/blog/articles/next-gen-application-management-services-transforming-rcg-industry

EMA | Torsten Volk. (2024, Q1). AI-Boosted Observability.
https://community.ibm.com/community/user/blogs/trent-shupe/2024/08/13/ema-analyst-paper

First Line | Roman Borovsky. (2025). How GenAI Improves Application Maintenance and Support Metrics.
https://firstlinesoftware.com/blog/how-genai-improves-application-maintenance-and-support-metrics

FlairsTech | Hagar Hadad. (2025, May 21). AI in Software Maintenance Support: What is New?
https://flairstech.com/blog/ai-software-maintenance-support

HCLTech | Kandasamy "Jam" Ramanujam. (2025, November 7). The impact of GenAI and Agentic AI on application management.
https://www.hcltech.com/trends-and-insights/the-impact-of-genai-and-agentic-ai-in-app-management

IBM. (2025). Smart, scalable and AI-driven: The next era of application management.
https://www.ibm.com/think/insights/next-era-application-management

Infosys | Praveen Pejaver Rao & Vineeth Av. (2024). AI-First Application Management Services with Generative AI.
https://www.infosys.com/services/adm/documents/application-management.pdf

ISG & HCLTech | Akhila H & Oliver Nickels. (2025, September). HCLTech's AMS in the Age of Agentic AI.
https://www.hcltech.com/sites/default/files/documents/resources/whitepaper/files/2025/10/03/ISG-HCLTech-Agentic-AI-AMS-Framework.pdf

INSPIRING SOURCES

Quinnox | Arun C.R. (2024, December 19). How AI Agents are Revolutionizing Application Management: Moving Beyond the AI Hype.
https://www.quinnox.com/blogs/how-ai-agents-are-revolutionizing-application-management

Quinnox. (2024, September 4). The State of AIOps 2024 and the Emergence of iAM (Intelligent Application Management).
https://www.quinnox.com/lens/the-state-of-aiops-and-emergence-of-intelligent-application-management-iam

Phenom^{e}Cloud. (2025, March 11). How AI is Revolutionizing Application Management Services.
https://blog.phenomecloud.com/featured-blogs/how-ai-is-revolutionizing-application-management-services

SmartDev | Ngoc Nguyen. (2025, July 2). AI in Application Support: Top Use Cases You Need To Know.
https://smartdev.com/ai-use-cases-in-application-support

TCS | Raka Sengupta & Swati Gupta. (2025). Next-gen application management services get an AI makeover.
https://www.tcs.com/what-we-do/services/enterprise-solutions/white-paper/next-gen-ams-operations-enterprise-applications

Wipro | Ashish Narayan. (2025, June 27). Agentic AI in Application Maintenance and Support Projects.
https://wiprotechblogs.medium.com/agentic-ai-in-application-maintenance-and-support-projects-043984d32a83

Wipro. (2025). Next Gen AMS Powered By SPRINT.
https://www.wipro.com/applications/services/next-gen-ams-powered-by-sprint

Xoriant | Sandeep Raje. (2024, November 5). Revolutionizing Product Sustenance: Leveraging GenAI for Smart Support and Engineering.
https://www.xoriant.com/thought-leadership/article/revolutionizing-product-sustenance-leveraging-genai-for-smart-support

YASH | Veera Venkata Naga Ravi Krishna Desaraju. (2025, November 18). Redefining AMS Talent: How AI Elevates Human Roles in Managed Services.
https://www.yash.com/blog/how-ai-elevates-human-roles-in-managed-services

YASH | Veera Venkata Naga Ravi Krishna Desaraju. (2024, August 7). Generative AI in the context of AMS.
https://www.yash.com/blog/generative-ai-in-the-context-of-ams

Jérôme Fenyo

AUTHOR

Jérôme Fenyo is a consultant and interim executive specialized in digital, data and AI-driven transformation, with over twenty-five years of experience at Capgemini, Atos and Klee Group. He has led technology business lines in data and AI and enterprise content management as well as sales engineering organizations in application services. Through his own brand, Goyave Consulting, he advises large organizations, IT service providers, start-ups and scale-ups, and investors, designing intelligent application strategies and operating models that combine innovation, industrialization and measurable business impact.