

The 2026 **AI playbook** **for enterprise IT**

5 AI and analytics capabilities every
CIO must adopt today



Table of contents

■	Reimagining enterprise ITOps for the modern era	3
■	Accurately predicting IT demands with dynamic foresight	4
■	Automating critical ITOps initiatives with agentic AI	8
■	Future-proofing security for volatile landscapes	11
■	Optimizing cloud costs and utilization: FinOps reimagined	13
■	Building an always-on enterprise with unparalleled reliability through data-driven SRE	18
■	Reliable, AI-powered ITOps is a strategic imperative for 2026	21
■	About ManageEngine Analytics Plus	22

Reimagining enterprise ITOps for the modern era

Enterprise IT is undergoing its most significant operational overhaul since virtualization and the rise of hybrid infrastructure. By 2026, achieving seamless ITOps resilience and efficiency will no longer be an aspiration or a competitive advantage; it will be a fundamental prerequisite for sustained enterprise growth.

Yet most IT teams are still constrained by manually governed, siloed, and tool-centric operations. These legacy approaches, characterized by scaling teams, adding more tools, tightening SLAs, and performing reactive fixes, can't keep pace with the velocity, volatility, and volume of modern digital environments.

This makes AIOps an essential priority for IT leaders in 2026.

To achieve truly efficient and resilient ITOps, organizations must shift from traditional manual oversight to an AI-powered operations model—one built on adaptive intelligence, automated decision-making, and minimal guesswork. In this model, the ITOps ecosystem continuously senses patterns, recognizes anomalies, predicts disruptions, self-optimizes, and evolves into a strategic growth engine.

GenAI and agentic AI sit at the heart of enterprise ITOps transformation. GenAI turns raw data into foresight, while agentic AI transforms that foresight into autonomous action. Together, they create an IT environment that is predictive, self-driving, cost-efficient, secure, and consistently reliable.

This e-book outlines five AI and analytics capabilities that serve as a practical blueprint for building predictive, autonomous, always-on enterprise ITOps, setting the stage for how IT teams must operate by 2026 and beyond.

01

Accurately predicting IT demands with dynamic foresight

Timely, accurate, and data-backed predictions are the reigning currency of control in IT. For years, predictive IT management has remained a top strategic priority for CIOs, and its importance will intensify in 2026 and beyond.

Traditionally, most enterprises rely on historical trends and apply standard univariate forecasting to anticipate future IT events, incidents, capacity shifts, and security risks. While this approach sounds foolproof on paper, real-world IT environments rarely operate under such stable or predictable conditions.

Effective forecasting must account for two critical aspects of IT operational data:

- The multitude of dependent factors or variables in data that collectively impact prediction outcomes.
- How an unanticipated change or variation in just one or two of these variables can drastically influence predicted outcomes.

Forward-looking IT teams are now adopting two advanced, AI-powered analytics techniques to ensure predictions are highly accurate and precisely reflect real-world scenarios:

a) ML-driven multivariate forecasting: By unifying operational data across ITSM, ITAM, IAM, cloud monitors, service desks, business systems, and HR tools, ML-based multivariate forecasting incorporates multiple dependent variables, not just historical patterns. This allows organizations to streamline forecasting, accurately capturing demand patterns, seasonal trends, risk indicators, and impending shortages, long before they have a chance to disrupt the business.

b) Scenario based what-if analysis: Layering scenario analysis on top of forecasting enables IT teams to simulate how various unanticipated future events and changes might impact the initial predicted outcomes.

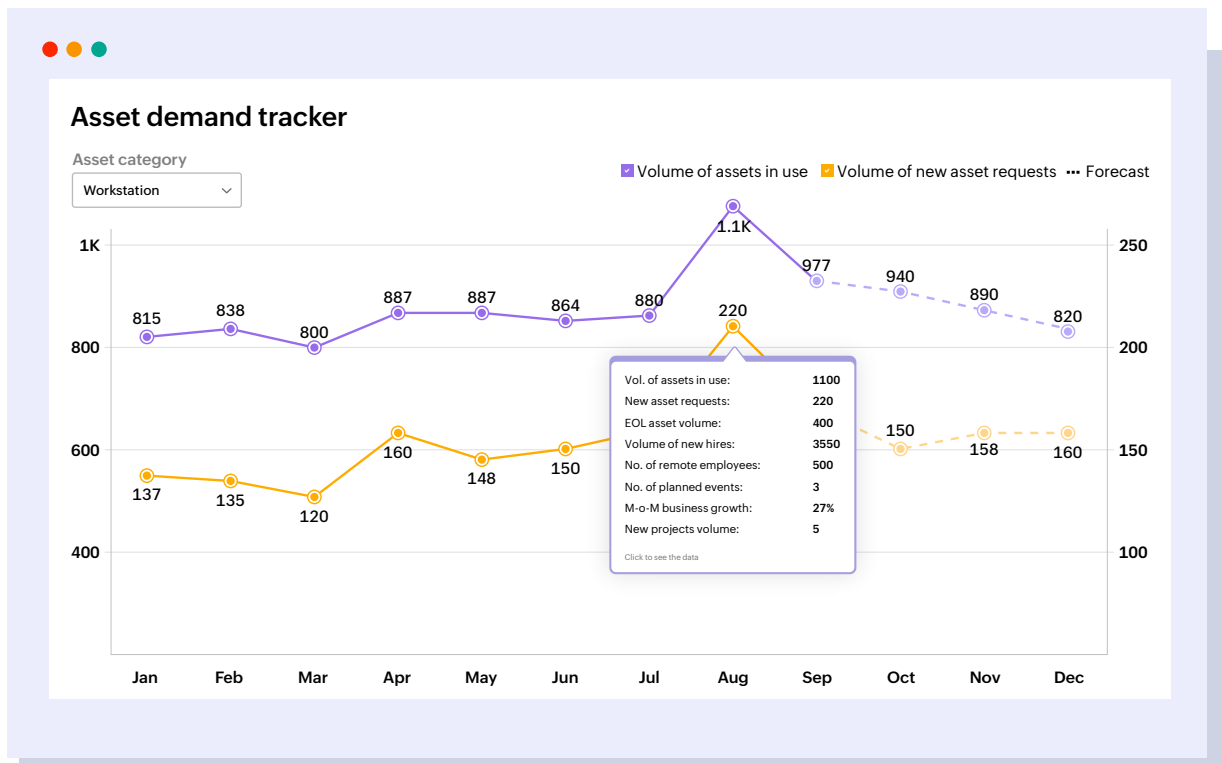
These two AI-driven capabilities are transformative for domains historically plagued by guesswork, last-minute procurement, sudden shortages, and budget overruns, such as workforce planning, capacity allocation, asset procurement, and budget management.

Consider the example of IT asset management. The following text explores how CIOs can eradicate reactive, cost-draining asset practices through strategic, precision forecasting using multivariate ML algorithms.

The procurement of end-user assets (like workstations and mobile devices) and software licenses is often based on linear projections of past consumption. Such a flawed approach frequently results in:

- **Over-procurement:** Leading to significant CapEx waste and unused shelf-ware.
- **Under-procurement:** Causing lengthy provisioning delays that hinder employee productivity and stall project starts.

Predictive asset demand intelligence applies multivariate ML models to forecast the complete life cycle, usage, demand, renewal cycles, and depreciation patterns of all IT assets—including workstations, laptops, VMs, network equipment, software licenses, and cloud subscriptions.

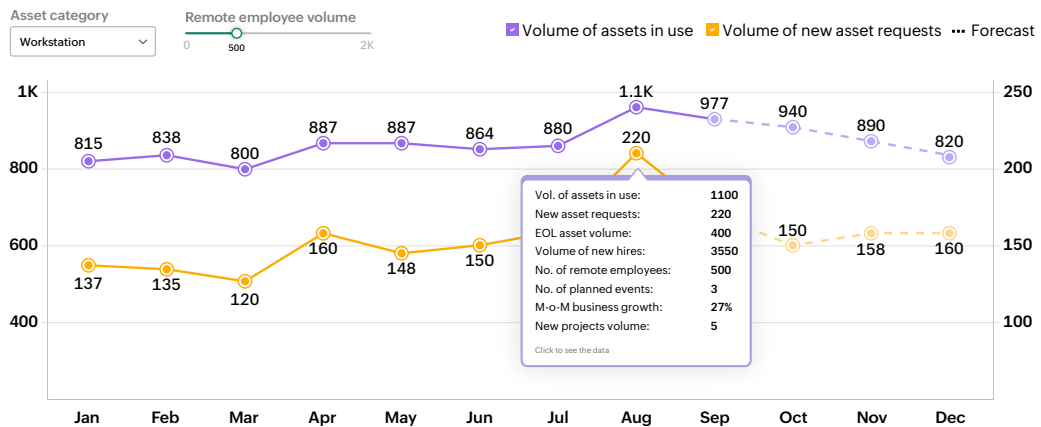


In the above analysis, Analytics Plus' multivariate forecasting models analyze how dozens of technical and business variables interact to shape workstation consumption and refresh demands. This can accurately predict workstation needs over upcoming months as seen in the analysis.

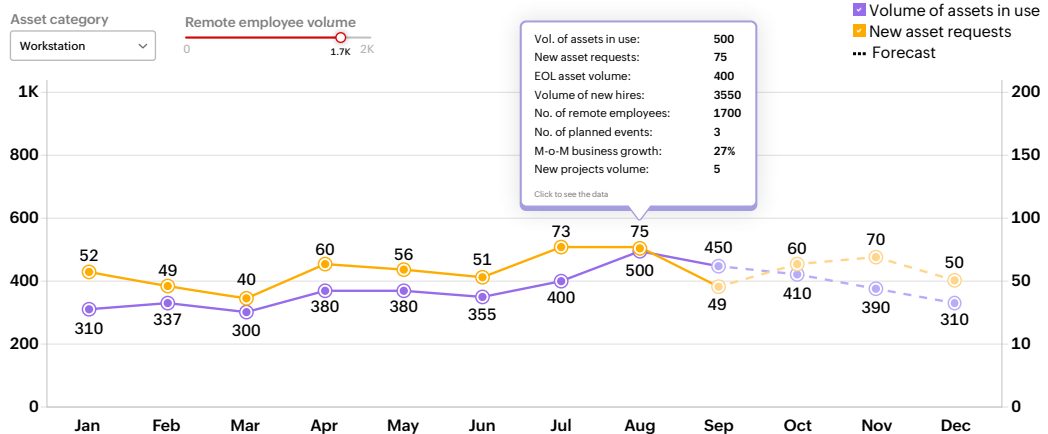
The forecasting can be further advanced using scenario analysis. Scenario analysis helps understand exactly how consumption patterns and the overall prediction would vary in the event of an unexpected change in a single dependent factor, such as a sudden increase in the remote workforce.

Impact of remote workforce on workstation demand

Workstation demand vs. No. of remote employees (Low)



Workstation demand vs. No. of remote employees (High)



The above comparative visualization dynamically forecasts that an increase in the proportion of remote workers (who often use mobile devices like laptops) will lead to a decrease in the need for and usage of traditional fixed workstations.

This strategic forecasting capability allows IT to dynamically identify when assets will fail, expire, reach peak usage, require upgrades, or hit renewal windows. By providing employees with the necessary tools immediately, it ensures operational agility and minimizes productivity loss. Furthermore, preventing wasted CapEx on shelf-ware ensures financial resilience, transforming asset management from simple reactive inventory tracking into strategic resource planning.

Automating critical ITOps initiatives with agentic AI

Traditional, human-centric IT operations, even those augmented with basic automation, are no longer viable. The sheer complexity of distributed tech stacks, hybrid environments, and the volume of events overwhelm technicians. This fuels chronic alert fatigue, slow MTTR, and persistent human error.

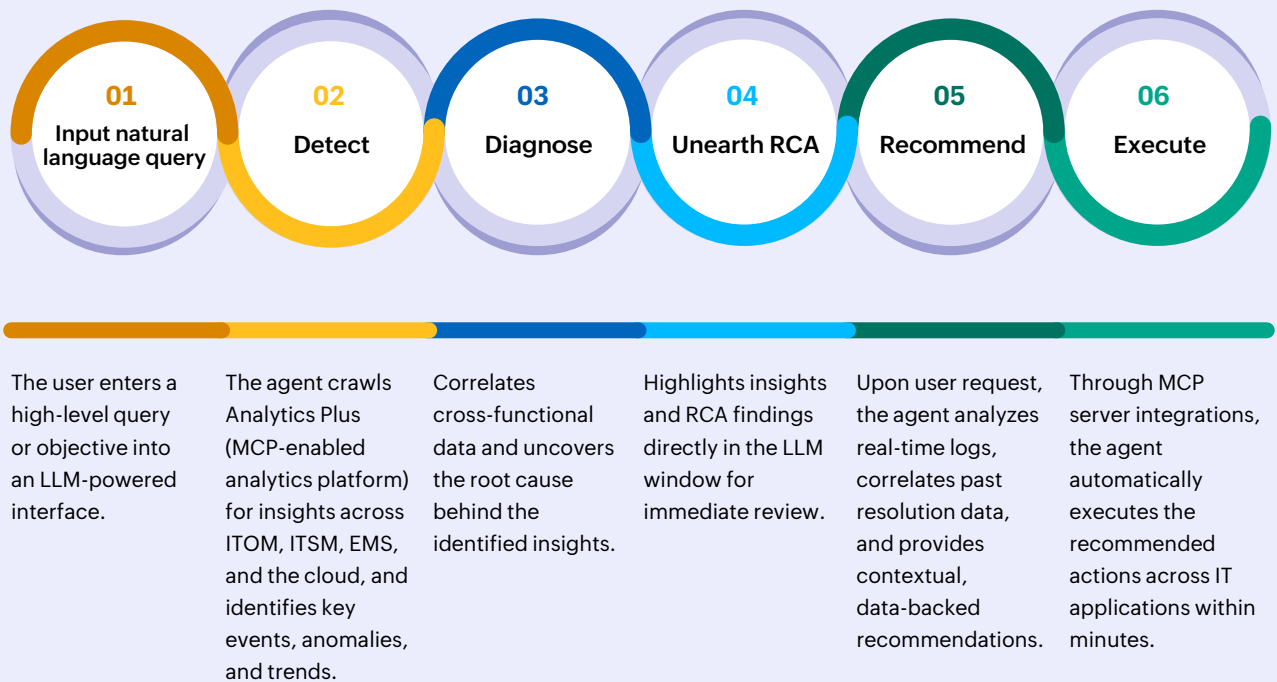
To ensure resilience, enterprises must move beyond this reactive approach toward self-healing, self-optimizing, closed-loop ITOps.

Agentic AI will be the backbone of this autonomous ITOps in 2026. Unlike traditional chatbots, agentic AI acts as a context-aware decision-making system capable of:

- Self-diagnosing issues.
- Planning multi-step remediation strategies.
- Executing actions across disparate systems.
- Optimizing outcomes continuously.

All with minimal to zero human intervention.

A future-proof, self-healing IT environment powered by agentic AI



AI-powered IT analytics platforms like Analytics Plus make it possible to build enterprise-grade agentic AI workflows with ease:

- It unifies real-time and predictive insights across the entire IT landscape.
- It enables deeper reasoning and automated root cause analysis (RCA).
- It generates GenAI-powered contextual recommendations for any IT problem.
- Its built-in Model Context Protocol (MCP) server support lets external applications and third-party LLMs (Claude, GPT, etc.) execute actions directly.

The visualization below demonstrates how a custom agentic AI workflow, powered by a centralized intelligence platform, identifies, assesses, and mitigates application performance issues using simple natural language queries.

Claude



List the average response time and daily outage hours for the top 5 high usage applications in my IT infrastructure



Response time and outage performance of 5 high usage applications:

Application name: HRMS

- Avg. response time (ms): 110
- Avg. daily downtime (hours): 0.5

Application name: CRM

- Avg. response time (ms): 280
- Avg. daily downtime (hours): 1.5

Application name: Email

- Avg. response time (ms): 230
- Avg. daily downtime (hours): 0.3

Application name: Cloud drive

- Avg. response time (ms): 700
- Avg. daily downtime (hours): 5

Application name: Support desk

- Avg. response time (ms): 350
- Avg. daily downtime (hours): 2.5



What's causing consistent high outage and slow response time for the Cloud drive?



Parameter	Downtime contribution	Impact on cloud drive downtime hours
Server downtime	2.8 hrs	35%
CPU saturation	1.9 hrs	24%
Planned downtime	1.7 hrs	21%
New DB patch deployments	1.6 hrs	20%



Give me recommendations for improving Cloud drive availability



Recommendations to improve Cloud drive availability and performance:

1. Replace problematic application server with a cloud instance or more stable on-premises compute.
2. Upgrade capacity allocation at HQ to accommodate the 40% increase in user volume.
3. Shift mandatory patch deployments to non-business hours to reduce productivity loss during planned downtime.
4. Streamline pre-deployment patch testing for DB patches to avoid patch-related compliance issues, critical vulnerabilities, and application breakages.



Looks good. Let's implement these recommendations and analyze immediately



Raising a change request and PO in the service desk for server replacement...
Creating an AWS EC2 RI purchase request via network manager application...
Updating patch deployment schedule in the endpoint manager...
Validating DB test automation workflows...
Analyzing new test automation workflow for DB server patches...

📄 📌 🔄 Retry

Claude can make mistakes. Please double-check responses.

Reply to Claude...

+ 🗨️ 🔍 Research

Claude Sonnet 4 📶

In this analysis:

- The agent first assesses application performance logs to understand application response times and outages, comparing them against established baselines.
- Once an application is identified as slow and outage-prone, the agent uses follow-up queries to find the RCA.
- The agent crawls the unified IT insights, uncovering all supporting components (application server, DB server, middleware etc.) and their respective metrics.
- Leveraging automated anomaly detection, the agent pinpoints metrics that deviated heavily from observed norms.
- Using advanced reasoning capabilities, the agent lists strategic, data-backed, and actionable recommendations to improve performance and prevent downtime.
- Finally, any tangible actions can be immediately triggered via an MCP server-enabled application directly from the LLM interface.

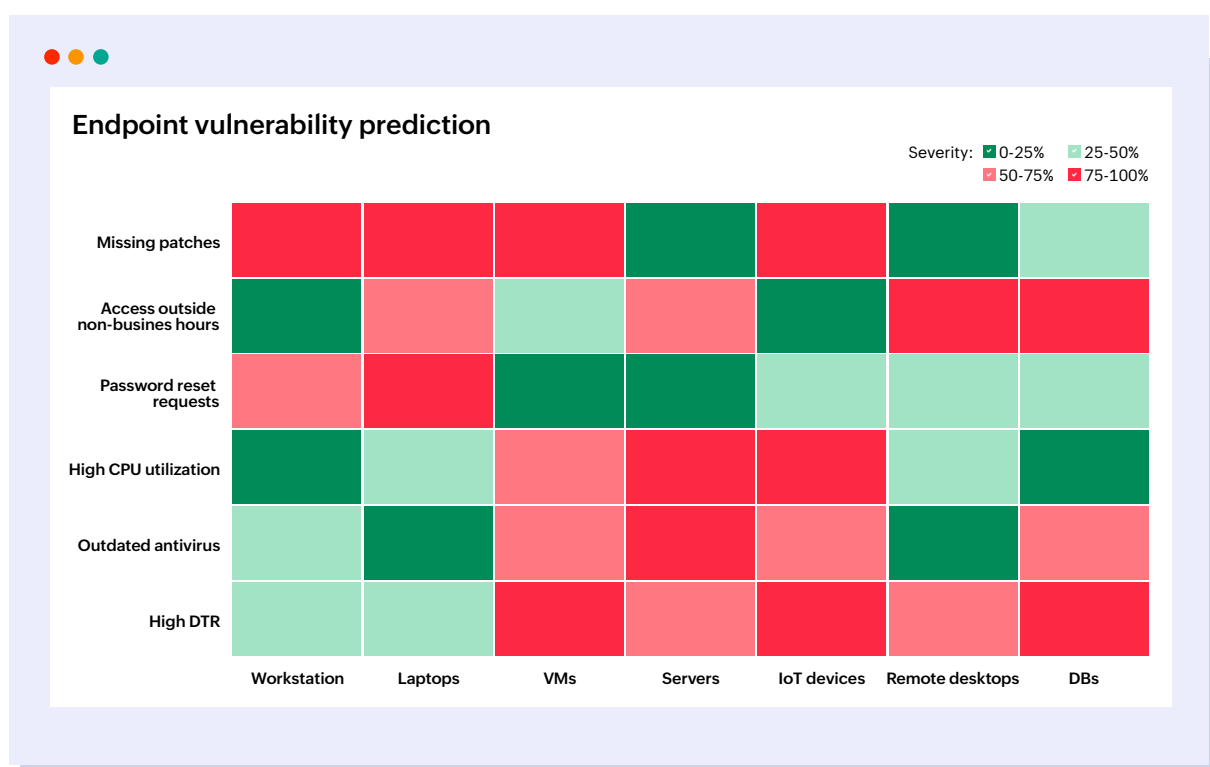
By minimizing human intervention and reducing operational costs, agentic AI ensures a significantly higher degree of continuous IT service availability, reliability, and robust operational resilience.

03 | Future-proofing security for volatile landscapes

The enterprise security landscape is more volatile than ever. Traditional, signature-based defenses can no longer keep up with modern sophisticated threats and AI-generated attacks. As threats grow in volume and complexity, security teams face widening detection gaps and slower response times. IT leaders need a faster, smarter way to uncover hidden risks and strengthen proactive threat detection.

One such challenge lies in critical, high-usage endpoints that harbor vulnerabilities often missed by standard security tools and monitoring applications. These blind spots, ranging from missing patches to abnormal access behaviors, frequently escalate into serious breaches.

In such scenarios, no-code ML becomes transformative for enterprise IT security management. Using AutoML capabilities in Analytics Plus, security teams can quickly build custom vulnerability prediction models without complex coding or data science expertise. These models analyze unified endpoint and security data from UEM and SIEM solutions, service desks, and other sources to identify which assets are most exposed to high-risk threat vectors that have resulted in critical breaches in the past.



Armed with this intelligence, security teams can take precise, proactive action, such as patching, tightening access controls, or removing risky shadow IT software, before threats materialize.

ML-driven vulnerability predictions are tailored to an enterprise's exposure to each specific vulnerability parameter, effectively reducing the attack surface and improving overall readiness against sophisticated threats.

With MCP server integrations, these actions can be triggered automatically, as with autonomous application performance workflows, to quickly contain and neutralize critical threats as soon as they are predicted.

04 | Optimizing cloud costs and utilization: FinOps reimagined

Cloud FinOps has become a strategic priority for modern enterprise IT teams. While cloud infrastructure offers agility, elasticity, and near-infinite scalability, it also introduces cost volatility and usage unpredictability. Without proactive, rigorous monitoring and management, enterprises quickly fall into patterns of cloud misuse, including:

- Overutilization of expensive services.
- Underutilization of provisioned resources.
- Cloud sprawl.
- Wasted or forgotten instances.
- Last-minute procurement at premium pricing.

This results in significant budget leakage from:

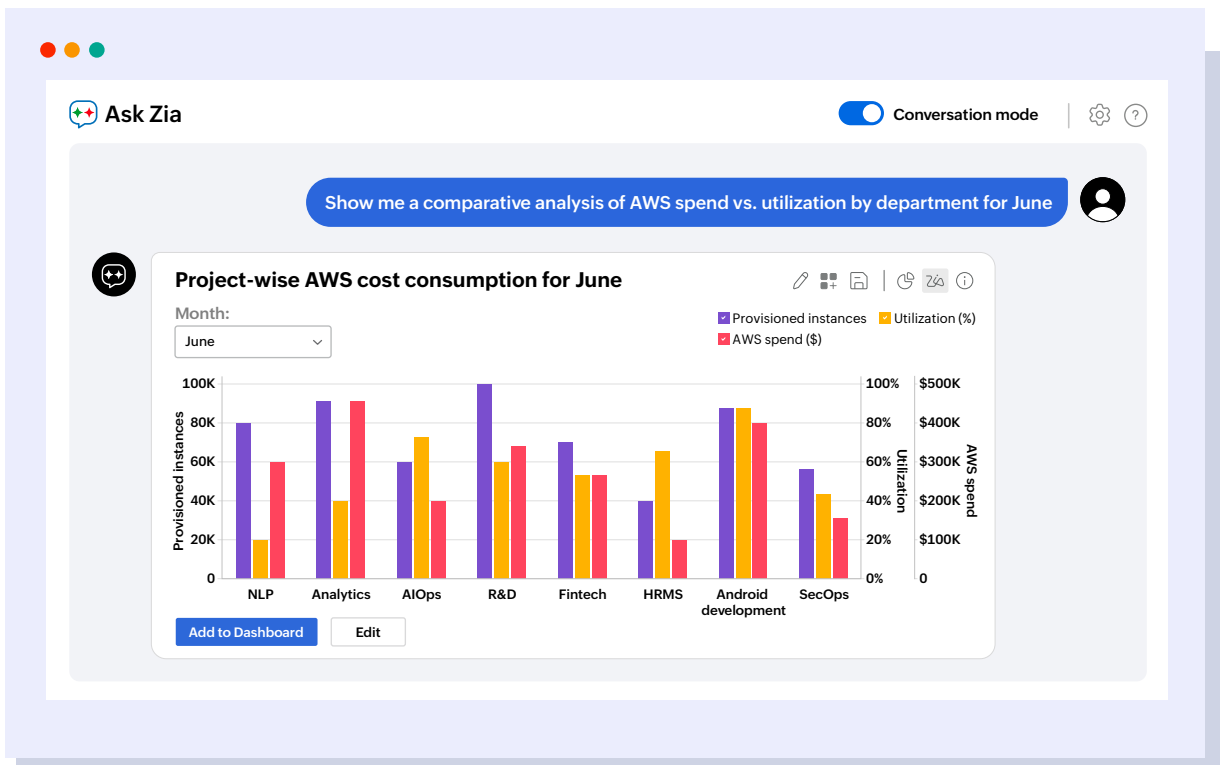
- Over-provisioned resources.
- Forgotten or wasted idle instances.
- Reliance on suboptimal pricing models for last-minute, urgent procurement.

GenAI now emerges as an essential and intelligent "cloud police" that empowers NOC teams to maintain optimum infrastructure performance while simultaneously controlling and reducing costs.

With GenAI-powered assistants like Ask Zia, cloud FinOps evolves from basic cost reporting into a continuous optimization engine, driven entirely through natural language queries rather than complex manual analysis.

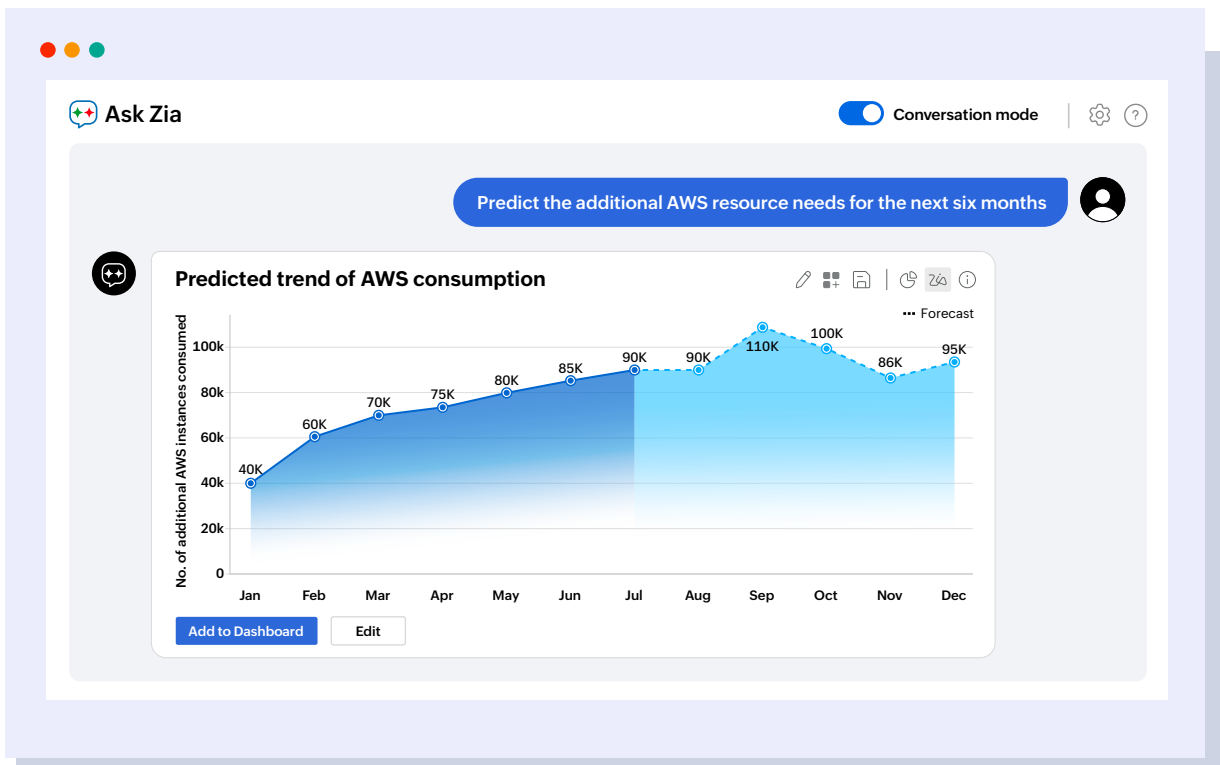
● Key GenAI-driven FinOps capabilities

a) Automated right-sizing and deprovisioning: GenAI assistants can continuously assess cloud resource utilization against provisioned capacity to immediately identify pockets of wastage. If a department consistently underutilizes resources, Zia flags and even executes right-sizing or reallocation actions to shift unused instances to units facing shortages.



This powerful analysis allows IT teams to quickly pinpoint major cloud wastage (departments with high spend where resources were idle or unused). Conversely, it also identifies mission-critical workloads that are frequently throttled and urgently require an upgrade. The GenAI analytics assistant helps teams triage cloud resource misuse and highlight opportunities for right-sizing and cost optimization within minutes, all without requiring analysts to dig through endless reports or input complex queries.

b) Intelligent capacity planning and procurement: Beyond utilization visibility, Zia helps forecast resource needs for upcoming months.



Once demand is predicted, Zia assists in evaluating the cost implications of different procurement options to arrive at an optimal procurement strategy.



This analysis identifies the right blend of additional instance procurement options:

- **Reserved instances (RIs):** Ideal for stable, predictable workloads, offering guaranteed capacity at significant discounts when purchased in advance.
- **Spot instances (SIs):** These are cheaper than RIs but can be interrupted in short notice. SIs are suitable for fault-tolerant, flexible tasks that can handle unplanned interruptions.

In addition to the above two scenarios, GenAI's prowess extends to other critical aspects of proactive cloud resource management, including:

- Anomaly detection for flagging unexpected cost spikes.
- Automated, optimal workload distribution across regions, projects, and multiple service providers.
- Quickly identifying and terminating idle or abandoned resources.

AI-driven cloud FinOps ensures cost resilience by ruthlessly cutting down wasted resources and intelligently optimizing cloud spend, thus freeing up crucial budget for innovation. Simultaneously, it enhances operational resilience through proactive right-sizing, guaranteeing that critical services always possess the necessary resources without excessive over-provisioning, preventing both performance degradation and budget overruns.

Building an always-on enterprise with unparalleled reliability through data-driven SRE

Any critical IT service downtime carries severe financial, strategic and operational consequences. Every minute of interruption translates directly into lost revenue, irreparable reputational damage, and rapid customer dissatisfaction.

Site reliability engineering (SRE) is a transformative discipline that applies core software engineering principles to IT operations (ITOps). SRE empowers organizations to fundamentally enhance reliability, minimize disruptive downtime, and scale efficiently by emphasizing proactive monitoring and data-driven decision-making.

An AI-powered SRE dashboard serves as the central nervous system for this effort. It expertly aggregates and correlates critical operational data from diverse sources, including:

- ITOps logs, metrics, and traces.
- Incident management.
- User behavior patterns.

All of this information is delivered in real time within a unified, intelligent view of live ITOps health.

360-degree SRE dashboard



Command center for always-on operations

The unified SRE dashboard above acts as a unified command center for achieving the goal of the always-on enterprise. It expertly bridges data boundaries, dismantles technological silos, and harmonizes cross-functional perspectives.

This unparalleled, comprehensive visibility empowers IT teams to gain a clear, immediate understanding of their IT health, facilitating rapid issue identification and swift resolution by:

- Standardizing reliability across distributed systems.
- Highlighting SRE breaches and compliance gaps.
- Enabling faster releases and infrastructure upgrades without risking stability.
- Forecasting potential pitfalls to ITOps reliability and efficiency.
- Connecting critical incidents and events to root causes.

When implemented effectively, such data-rich dashboards provide instant clarity across the entire IT infrastructure. This enables IT teams to instinctively prioritize critical issues and strategically address key gaps in SRE management practices—a pivotal step toward achieving an always-on and always-available status.

By incorporating robust SRE practices and leveraging AI-driven, unified observability, enterprises can establish measurable and enforceable strategies for elevating ITOps reliability.

Reliable, AI-powered ITOps is a strategic imperative for 2026

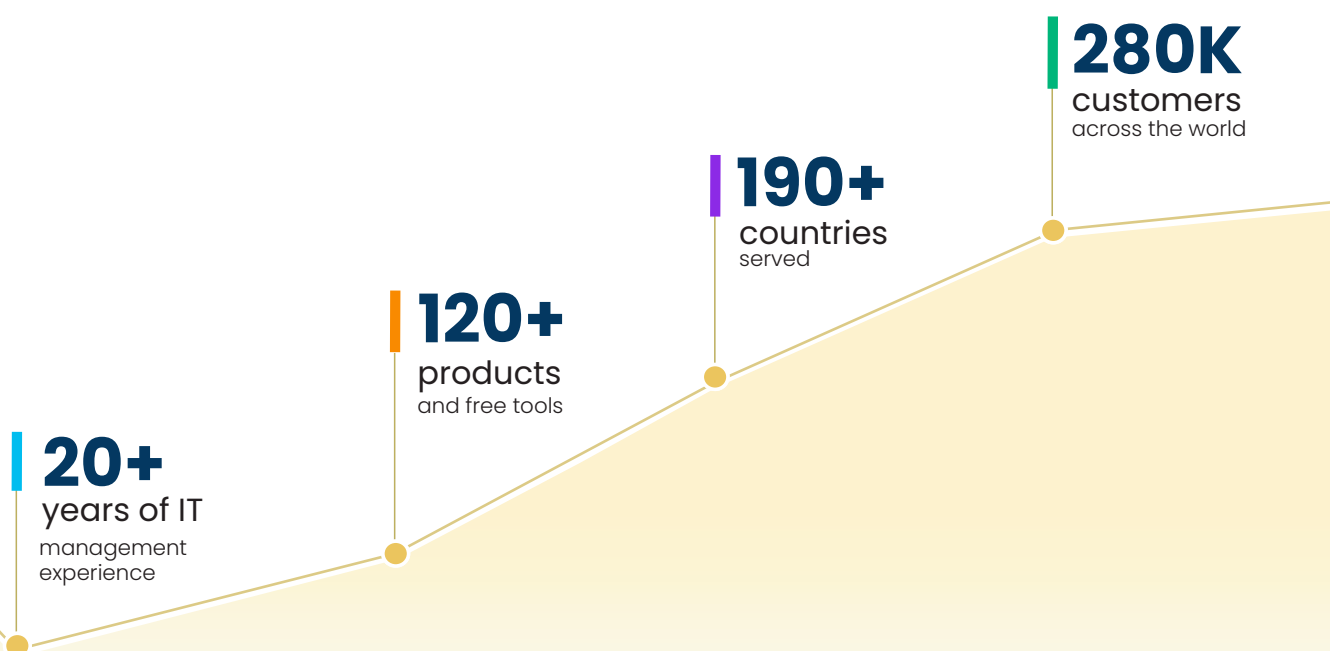
Building a resilient enterprise in 2026 is not a routine operational task; it is a strategic necessity. As ITOps becomes more complex and distributed, manual-centric models can't keep up. CIOs must embrace GenAI-powered analytics and agentic AI to strengthen resilience, control costs, and transform ITOps from a reactive upkeep function into a strategic driver of efficiency and growth. Enterprises that adopt the five analytical capabilities outlined in this guide will not only future-proof their operations but also unlock new levels of agility, performance, and competitive advantage. In 2026 and beyond, successful IT leadership will be defined by how intelligent, adaptive, and autonomously their ITOps functions, and AI-driven intelligence will be the essential engine that makes it possible.

About

ManageEngine Analytics Plus is an IT analytics and decision intelligence solution designed to provide organizations with a unified view of their IT operations, correlate interdependencies and derive meaningful insights. It breaks down data silos by consolidating both on-premises and cloud infrastructure KPIs. Analytics Plus measures the efficiency of network operations, tracks the responsiveness and availability of business applications, evaluates technician performance, assesses the progress of processes and flags security anomalies. This comprehensive analysis is achieved by connecting to all IT software that forms the backbone of an IT infrastructure. These consolidated insights enable organizations to make data-driven decisions that enhance operational efficiency and drive business success.

For more information about Analytics Plus,

visit: www.manageengine.com/analytics-plus/





© ManageEngine, a division of Zoho Corporation