

Revolutionize enterprise service desks with agentic AI

✦ Your roadmap to the future of IT service management.

Table of contents

■	Introduction	3
■	Resolution of complex asset management bottlenecks	4
■	Proactive SLA breach prevention	9
■	Optimizing technician performance	15
■	Predictive incident management	18
■	Conclusion	23
■	About ManageEngine Analytics Plus	24

Introduction

Enterprise service desks are under more pressure than ever. The volume of user requests is rising, the complexity of issues increasing, and end-user expectations are no longer satisfied by responses and eventual resolutions.

Traditional ITSM practices—ticket triage, back-log burning, SLA catch-ups—are no longer sufficient. What organizations need is a service-desk that not only responds to tickets but predicts, reasons, recommends, and in some cases, acts.

This is where generative and agentic AI come into play. GenAI-powered analytics transforms raw data into smart, narrative insights, supports natural-language questions, and surfaces hidden patterns. Agentic AI takes it further—it enables the system to prioritize action, automate workflows and optimize the loop from “analyze” to “execute”.

This e-book dives into the four most persistent service desk challenges and shows how agentic AI solves them for immediate, measurable results.

The e-book explores actionable frameworks, realistic outcomes and strategic insights to redefine your service desk for the era of autonomous operations that delivers tangible business value: fewer tickets, faster resolution, smarter teams and optimized assets.

Resolution of complex asset management bottlenecks

In today's sprawling hybrid IT environments, managing thousands of servers, endpoints, and network devices across cloud and on-premises setups has become one of IT's most complex balancing acts. IT asset management is inherently challenging because it demands orchestration across numerous disparate systems: the CMDB, AD, licensing platforms, endpoint management tools, and finance systems. Even a fundamental process like asset onboarding can involve dozens of manual, siloed steps, inevitably leading to significant bottlenecks, data inaccuracies, and security risks.

While IT teams tend to rely on monitoring dashboards and periodic maintenance schedules, these methods are inherently reactive. They offer a snapshot of past utilization or tell you which asset has already failed. They fail to answer critical forward-looking questions: Which asset will fail next? What is the optimal maintenance plan? Which devices are silently draining the budget? As the enterprise asset footprint and life cycle costs surge, this lag in insight directly translates into higher operational expenses, mismanaged maintenance, erratic utilization, and lost productivity.

The core challenges in traditional asset management

IT teams primarily wrestle with two critical, costly asset management challenges:


- **Unoptimized asset life cycle:** Assets are often purchased, deployed, and replaced based on arbitrary, time-based schedules rather than actual utilization, performance, or cost-efficiency indicators. This poor decision-making leads to premature purchases and replacements (underutilization) or, conversely, the continued use of inefficient, aging assets long after they should have been retired (overutilization).
- **Unexpected asset failures:** Despite routine maintenance, assets fail unpredictably due to usage or operational anomalies, aging hardware, or overlooked configuration changes. These failures not only halt IT operations but also trigger downstream incident surges and SLA breaches.


The necessary transformation to smarter, outcome-driven asset management is powered by agentic AI. This serves as a powerful orchestrator, capable of executing complex, cross-functional asset management workflows with minimal human intervention.

Crucial IT infrastructure health data—including the volume of failure alarms, performance issue requests, mean time to repair (MTTR), mean time between failures (MTBF), and error frequency—holds a wealth of early-warning signals that are routinely overlooked. By leveraging the Model Context Protocol (MCP) server, LLMs can link up with IT analytics platforms like Analytics Plus, pulling in such comprehensive data from the service desk, endpoint management and security, asset management, and other monitoring tools.

This creates a powerful agentic AI setup that utilizes the analytics platform's no-code machine learning (ML) models to analyze various cross-functional parameters. The models automatically learn the correlations between current asset performance metrics and historical failure events to identify critical assets highly likely to fail in the upcoming weeks. The LLM then returns quick, actionable insights in natural language on the top problematic assets in any IT landscape.

Claude

 Which assets are likely to fail in the next 30 days based on critical health parameters such as average MTTR, MTBF, recent error rates, service request history, and asset age?



I'll analyze asset reliability based on key health indicators such as average MTTR and MTBF, current error rates, total service requests in the past 6 months, and asset age.

Let me query existing data.

Reading Analytics Plus data...




Evaluating MTBF, MTTR, error rate...

Calculating the failure probability...

Here are the top 10 assets with the highest probability of failure in the next 30 days:




Assets with high failure probability


Asset name	Asset type	Failure probability in next 30 days	Requests in last 6 months	Avg. MTTR (hrs)	Avg. MTBF (days)	Error rate (%)	Asset age (Years)
Ser1023	Application server	92%	24	18	15	18	2.5
RT1956	Core router	85%	18	9	22	15	4
DBS1409	Database Server	72%	20	12	29	13	5
APF1029	Access Point	63%	12	7	31	11	2
FWP1702	Firewall	58%	16	8	32	9	3.5
MAC2043	Workstation	55%	10	5	38	10	4
SVM 5220	Virtual Machine	44%	9	6	52	6	1.5
ERT2201	Edge Router	49%	9	9	45	5	1
NSW2009	Network Switch	39%	8	5	48	3.5	3
HPP8512	Printer	34%	7	4	55	3	2


 Retry

Claude can make mistakes. Please double-check responses.

Reply to Claude...

 Research

Claude Sonnet 4 

 ManageEngine
Analytics Plus

6

The visualization above identifies troublesome IT assets with a high probability of failure, which would directly impede organizational productivity and drive up operational and repair costs.

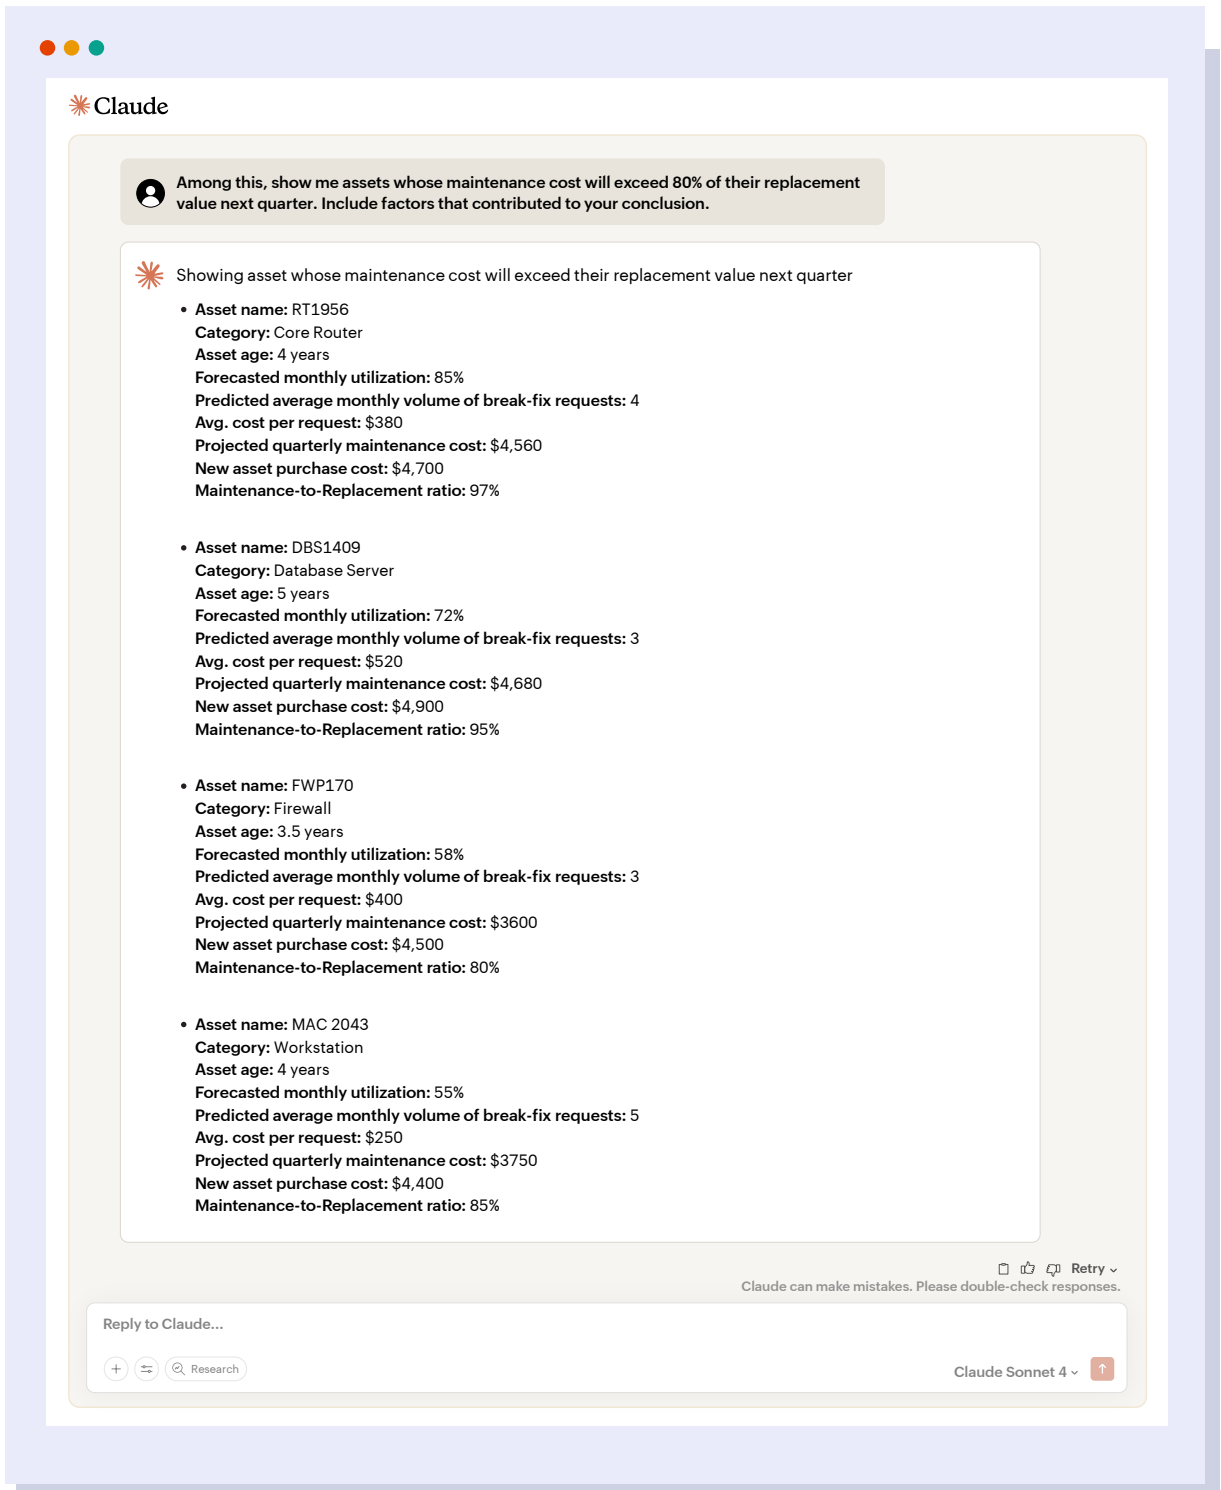
By anticipating asset failures weeks in advance, this in-depth representation enables smarter and more agile maintenance scheduling, leading to reduced MTTR, less downtime, and a higher MTBF. Organizations can then make data-backed decisions: retain, repair, or replace.

Agentic AI also plays a critical role in solving the unoptimized life cycle challenge by forecasting when an asset's cost of upkeep outweighs its operational value.

By seamlessly blending utilization trends with financial data (maintenance costs, downtime losses, residual value, and even new asset purchase costs), agentic AI, powered by tailored ML models, pinpoints each asset's optimal replacement window.

With the built-in AutoML capability of Analytics Plus, IT managers can train, validate, and deploy custom ML models effortlessly, eliminating the need for an expert data scientist and without writing a single line of code.

For instance, when prompted through the LLM interface, the agent can discover that certain aging assets exhibit rising maintenance, and support costs, coupled with decreasing utilization. This signals that the total cost of ownership will soon exceed the cost of purchasing a more efficient alternative.



This GenAI-driven analytics transforms the asset purchase, maintenance, and replacement decisions from intuition-driven to data-driven, saving significant operational costs while ensuring optimal uptime.

These data-driven insights form the bedrock of smarter ITAM. Provided the necessary applications and platforms are MCP-enabled, IT teams can even execute required corrective actions such as replacing costly, aging assets or proactively updating maintenance—directly from the LLM screen.

This creates a 360-degree, closed-loop automation of insights, recommendations, and remediation, without ever switching applications or performing time-consuming analysis.

A holistic, agentic AI-powered approach transforms asset management from a necessary cost center into a strategic function that actively mitigates failure risk, avoids downtime, improves utilization, and ensures optimal financial planning aligned with overall business goals.

02

Proactive SLA breach prevention

SSLAs are more than just targets; they are the universal standard for your IT service desk's performance and trust. While technological advancements like automation and digital self-service have successfully streamlined ticket resolution and reduced common, mundane issues, a persistent challenge of SLA violations still remains. In fact, with the evolution of new age technologies like AI and NLP, the complexity and urgency of tickets have increased—pushing SLA compliance risks further.

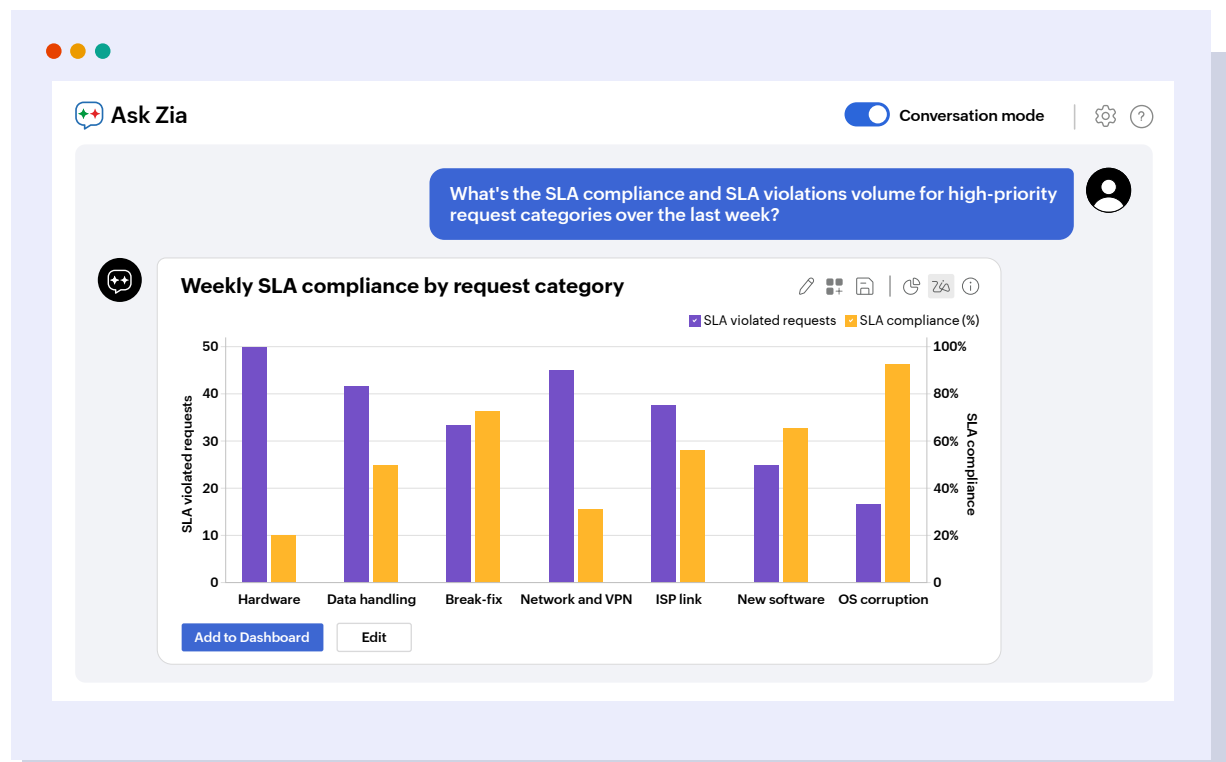
A high volume of cascading SLA violation carries significant consequences. It directly erodes user trust, damages IT productivity, and can lead to notable business and financial issues. The paramount objective for today's ITSM leaders is clear: to drastically lower the SLA breach rate by establishing a system for proactive prevention. This is essential for meeting ever-rising user expectations and contributing to overarching business goals.

The era of manually sifting through mountains of data is over. That cumbersome, time-consuming process of pulling SLA reports, cross-referencing request details, and manually identifying the root cause of violations is now obsolete.

Today, the entire process—monitoring, identification, RCA, and remediation—can be executed from a single screen with unparalleled simplicity and speed, all thanks to conversational analytics. This powerful approach removes the need for complex SQL queries and tedious number-crunching.

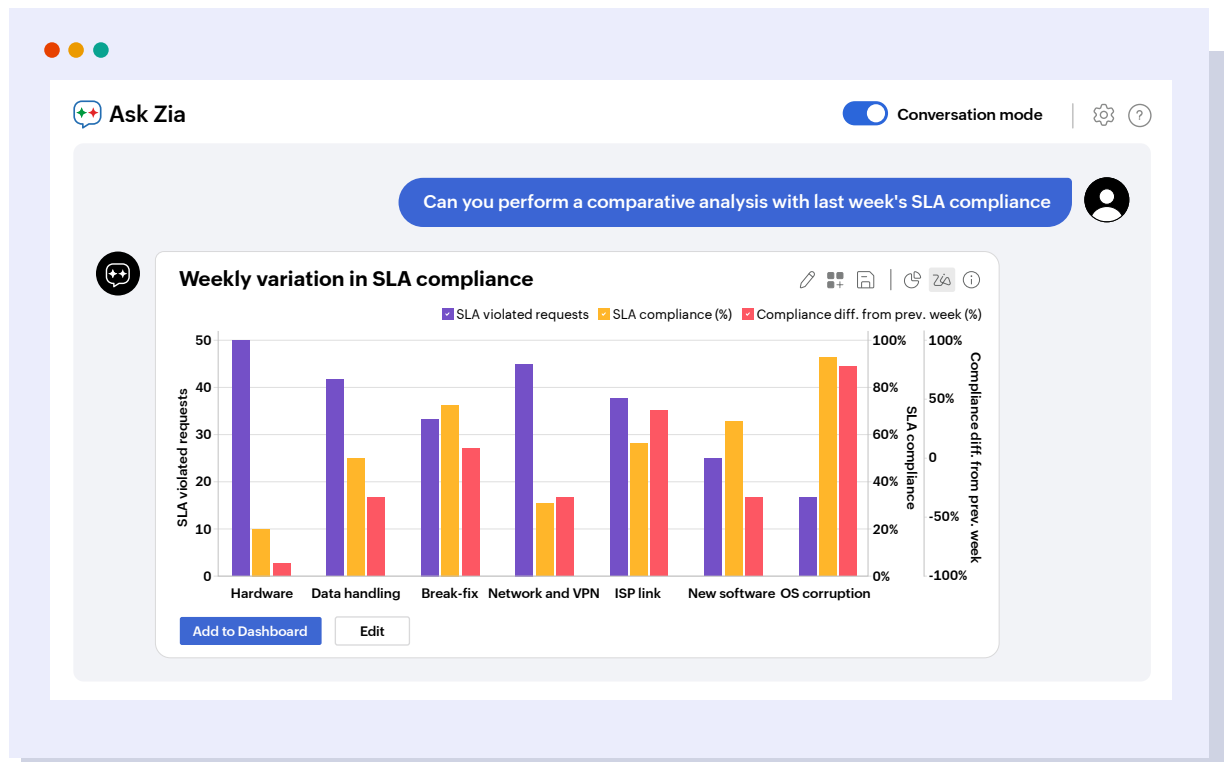
This is where a tool like Zia, the GenAI-powered conversational assistant supporting Analytics Plus, revolutionizes IT operations.

Service desk managers often face a flood of high-priority SLA escalations. Instead of spending endless hours manually digging through service desk logs and raw data, they can simply ask Zia to determine the volume of SLA compliance and violations for high-priority request categories.

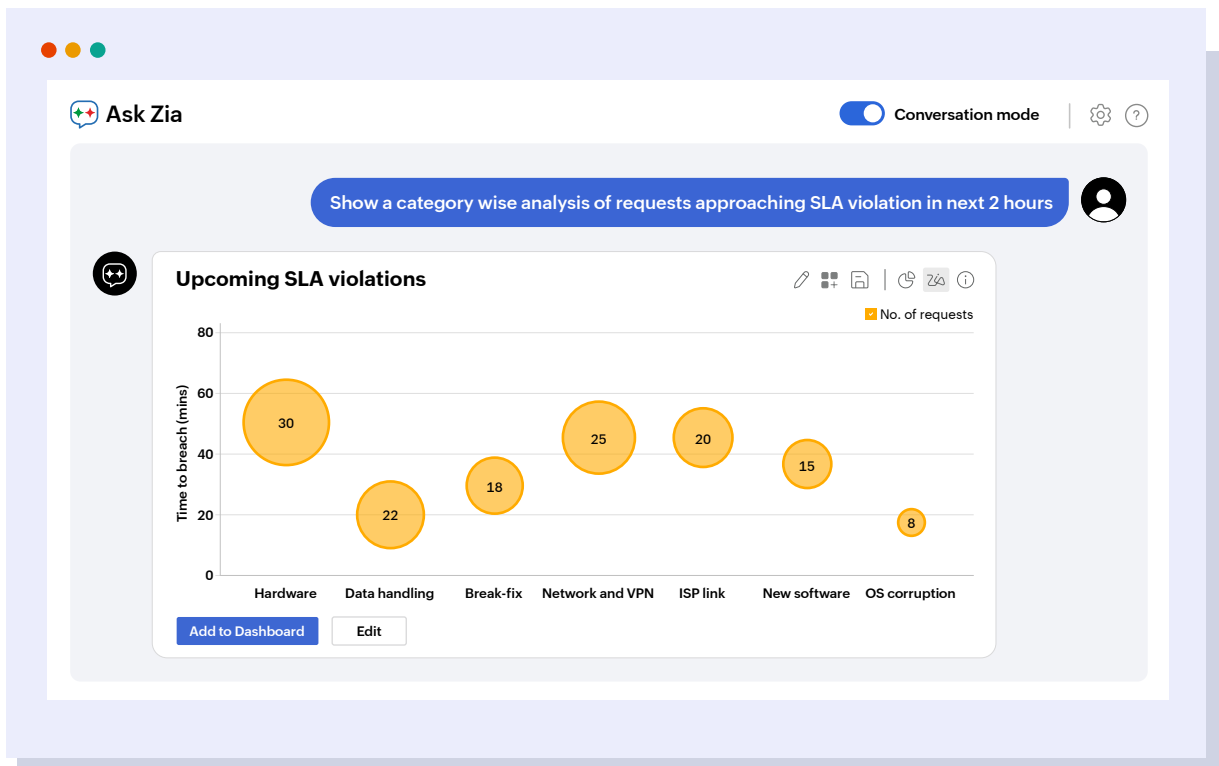


The analysis here reveals that specific request categories have a significantly lower compliance percentage and a higher volume of violations compared to others.

However, the compliance number alone is just a signal, and does not tell the full story. To gain further insights and a clearer picture, ITSM leaders need to see the deviation from the previous week's SLA performance as shown below.



GenAI-powered analytics also enables service desk managers to gain further visibility into potential SLA breach risks days in advance.



With these insights, instead of operating in a perpetual state of reactive firefighting, IT teams can course-correct proactively. This transition from "What just happened?" to "What's about to happen?" is the essence of modern ITSM.

Identifying the critical SLA violation drivers

Quick remediation is not possible without quicker RCA. IT teams can use Zia's deep reasoning and cross-functional correlation to identify the critical factors that led to an increase in SLA violations.

Consider a simple query: "What's driving a high weekly growth in SLA violations for hardware requests?"

In a matter of seconds, the analytics assistant analyzes the entire IT landscape, correlating SLA-linked metrics across diverse sources, including:

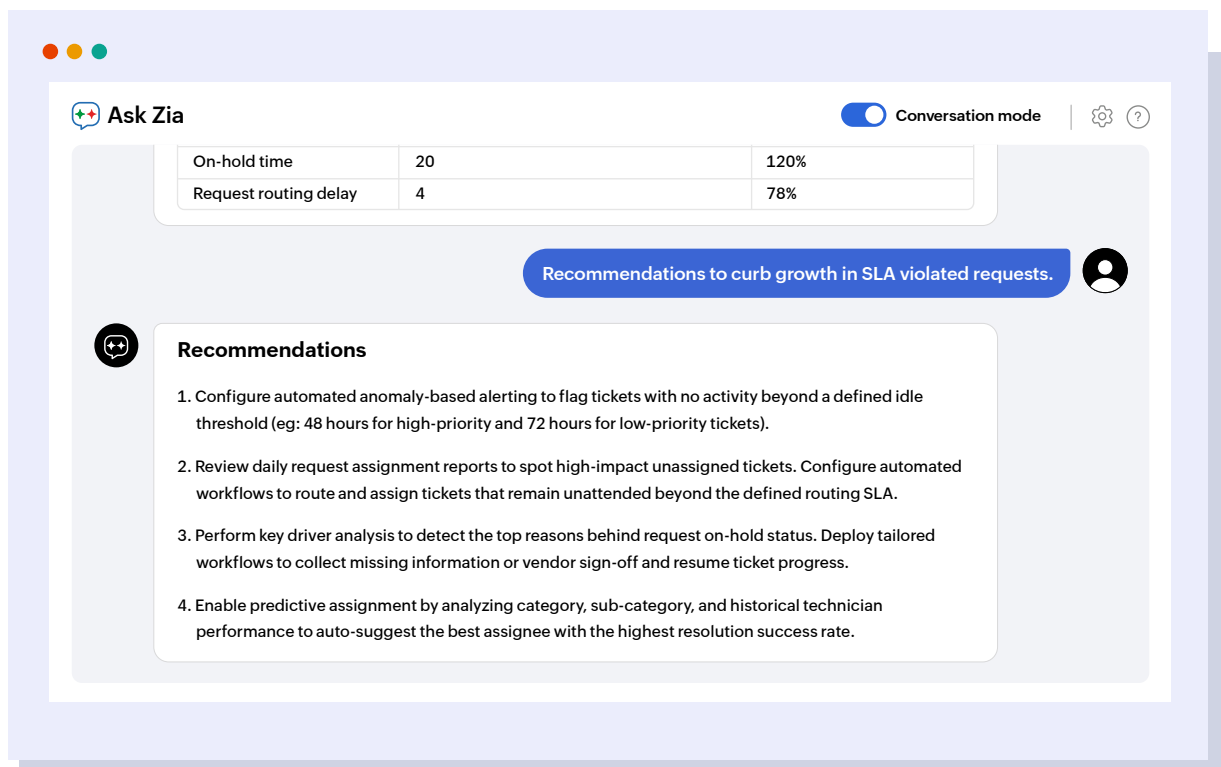
- Request priority
- Workload distribution
- Technician queues
- Ticket routing

It unearths the top critical factors responsible for the growth in breaches across specific request categories.

The screenshot displays the 'Ask Zia' interface within a window. At the top, there's a header with the 'Ask Zia' logo, a 'Conversation mode' toggle, and settings/help icons. Below the header, a blue input bar contains the query: 'What's driving the high weekly growth in SLA violations for Hardware requests?'. The response area features a chat icon and a title 'Reasons for growth in Hardware request SLA violations'. It lists four key drivers for growing SLA violations: 1. Increase in volume of Stale requests from Requests, 2. Growth in unassigned requests from Requests, 3. High on-hold time from Resolution, and 4. Delay in routing to other service levels or functional teams from Request. Below this list is a section titled 'Key drivers and their impact' which contains a table.

Potential drivers	Contribution to increase in SLA violation	Impact on SLA violation
Stale requests	12	98%
Unassigned requests	9	84%
On-hold time	20	120%
Request routing delay	4	78%

The analytics assistant doesn't stop at RCA. Zia can further synthesize and correlate this data with information from various knowledge sources, such as KB articles, past resolution notes, and resource documentation, to deliver data-backed recommendations designed to curb the growth of SLA-violated requests, transforming raw data into actionable strategies.



Every SLA breach is a story that reveals more than just inefficiency; it highlights hidden operational friction and technician bottlenecks. By proactively surfacing these signals and quantifying their impact, ITSM leaders can move beyond simple reporting. They can:

- Pinpoint the exact areas of friction.
- Understand their true impact on the business.
- Trigger targeted remediation using integrations (like MCP server integrations) for automated actions in respective applications and modules.

This entire process transforms SLA management from reactive firefighting to intelligent prevention, creating a seamless bridge between insight and action.

Optimizing technician performance

While avoiding SLA breaches is critical, the unavoidable truth is that technician performance bottlenecks affect the service desk around the clock, not just during an SLA crisis. Technician productivity is the most crucial lever for service desk excellence, directly influencing resolution efficiency, SLA compliance, user experience, and operational stability.

When service desk technicians grapple with issues like disproportionate workloads or inefficient routing, the negative consequences ripple across the entire ITSM ecosystem. This manifests as slower resolutions, inconsistent SLA adherence, higher escalation rates, and a decline in employee morale. For service desk managers, real-time, proactive tracking of these performance gaps must be a top priority.

Modern GenAI-powered IT analytics platforms are fundamentally changing how IT leaders monitor and address technician performance. Instead of manually navigating complex dashboards, service managers can now gain real-time, AI-backed insights on technician performance directly within their preferred LLM window.

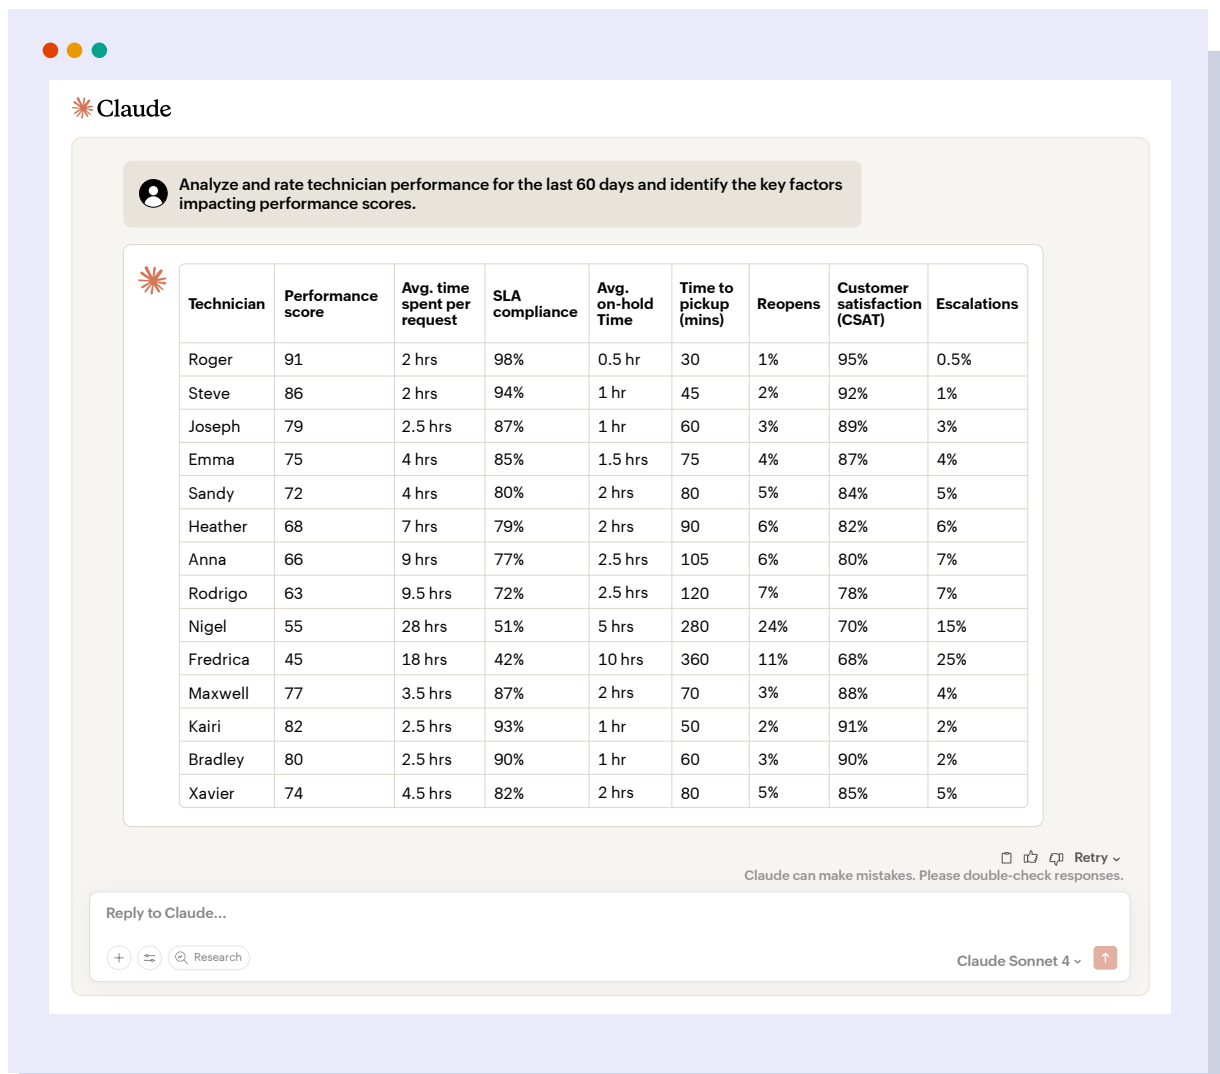
This leap is facilitated by the MCP server setup that allows LLMs to interact with and leverage the AI-powered analytics and no-code ML models within IT analytics platforms like Analytics Plus.

The power of agentic AI in mitigating technician performance gaps is best illustrated through a simple, yet transformative, scenario.

A service desk manager is looking to understand and monitor technician performance and factors driving performance issues. They simply need to pose their requirement to the LLM interface:

"Analyze and rate technician performance for the last 60 days and identify the key factors impacting performance scores."

In mere seconds, the LLM links up with the MCP-enabled IT analytics platform like Analytics Plus. These agentic AI platforms automatically analyze and correlate multiple skillset factors to calculate a precise performance score for every technician account. The LLM then returns quick, actionable insights in natural language focused on the underperforming technicians.



The screenshot displays the Claude AI interface. At the top, the Claude logo is visible. Below it, a prompt is entered: "Analyze and rate technician performance for the last 60 days and identify the key factors impacting performance scores." The response is a table of technician performance data.

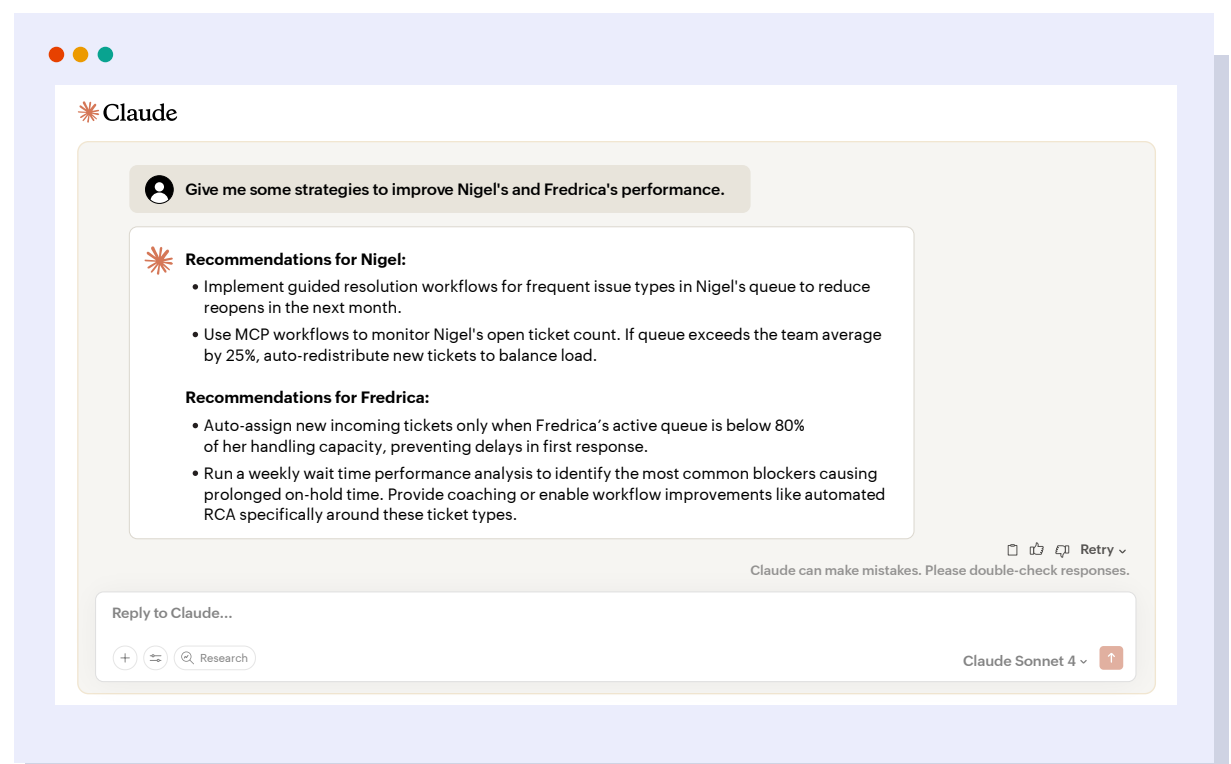
Technician	Performance score	Avg. time spent per request	SLA compliance	Avg. on-hold Time	Time to pickup (mins)	Reopens	Customer satisfaction (CSAT)	Escalations
Roger	91	2 hrs	98%	0.5 hr	30	1%	95%	0.5%
Steve	86	2 hrs	94%	1 hr	45	2%	92%	1%
Joseph	79	2.5 hrs	87%	1 hr	60	3%	89%	3%
Emma	75	4 hrs	85%	1.5 hrs	75	4%	87%	4%
Sandy	72	4 hrs	80%	2 hrs	80	5%	84%	5%
Heather	68	7 hrs	79%	2 hrs	90	6%	82%	6%
Anna	66	9 hrs	77%	2.5 hrs	105	6%	80%	7%
Rodrigo	63	9.5 hrs	72%	2.5 hrs	120	7%	78%	7%
Nigel	55	28 hrs	51%	5 hrs	280	24%	70%	15%
Fredrica	45	18 hrs	42%	10 hrs	360	11%	68%	25%
Maxwell	77	3.5 hrs	87%	2 hrs	70	3%	88%	4%
Kairi	82	2.5 hrs	93%	1 hr	50	2%	91%	2%
Bradley	80	2.5 hrs	90%	1 hr	60	3%	90%	2%
Xavier	74	4.5 hrs	82%	2 hrs	80	5%	85%	5%

Below the table, there is a text input field with the placeholder "Reply to Claude..." and a "Research" button. At the bottom right, it says "Claude Sonnet 4" with an upward arrow icon. A small disclaimer at the bottom right reads: "Claude can make mistakes. Please double-check responses."

To provide a comprehensive view, LLM has displayed this deep skillset measure in an intuitive table, assigning each technician a tailored performance score based on how they fare across seven critical health segments.

From the visualization, it's clear that for technicians Nigel and Fredrica, these combined factors have directly led to lower satisfaction scores and, consequently, the lowest overall performance scores among the team.

Service desk managers can also leverage the LLM to synthesize the identified performance drivers and gain tailor-made, actionable recommendations to improve each technician's performance.



The difference between an efficient and a struggling service desk isn't just the tools; it's the technician enablement.

By seamlessly integrating advanced no-code ML capabilities with the narrative context of LLM-based reasoning, IT leaders gain unparalleled clarity into roadblocks or challenges affecting individual technician performance. They don't just see who is underperforming, instead they finally gain insight into the why, and what specific steps to take to correct it.

The outcome is a service desk that evolves from being data-informed to decision-intelligent. By establishing the analytics platform as the unified data source and using flexible LLM integration for conversational diagnosis, IT leaders gain a powerful and scalable means to close technician performance gaps, transforming raw data into strategic IT workforce management decisions.

This approach sets the stage for the next phase: predictive incident management.

04

Predictive incident management

Incidents like downtimes, system slowdowns, and crashes are inevitable in any IT landscape. However, the critical realization is that incidents rarely occur without warning. They are almost always preceded by subtle operational indicators or signals. Unfortunately, most service desks are reactive, only detecting these signals after an incident.

This leads to a chaotic cycle of firefighting—a scramble to resolve the crisis and respond to a massive surge of incident-related tickets. In today's complex, interconnected, and hybrid environments, this reactive model is unsustainable, costly, and cumbersome.

Even ignoring a small signal like a minor configuration change followed by a system slowdown, or a resource consumption spike, can cascade across departments, leading to a massive surge in incident volume within hours.

Service desks must have a mechanism to foresee potential incident surges and act before they snowball into a huge pile of tickets that impact SLAs and user experience. With unified, AI-powered predictive incident management, ITSM teams can identify and resolve issues proactively, even before a single ticket is logged.

AI-powered analytics platforms are the foundation for this shift. They correlate massive volumes of data, including:

- Service desk requests and alerts
- Performance metrics from NOC or ITOps tools
- Change records and system logs
- CMDB data

By analyzing these interconnected sources, the analytics platform identifies patterns that allow it to forecast when, where, and how a potential incident storm will surface—days or even months in advance. IT managers can access these critical insights using simple natural language queries, empowering them to neutralize threats before they reach landfall.

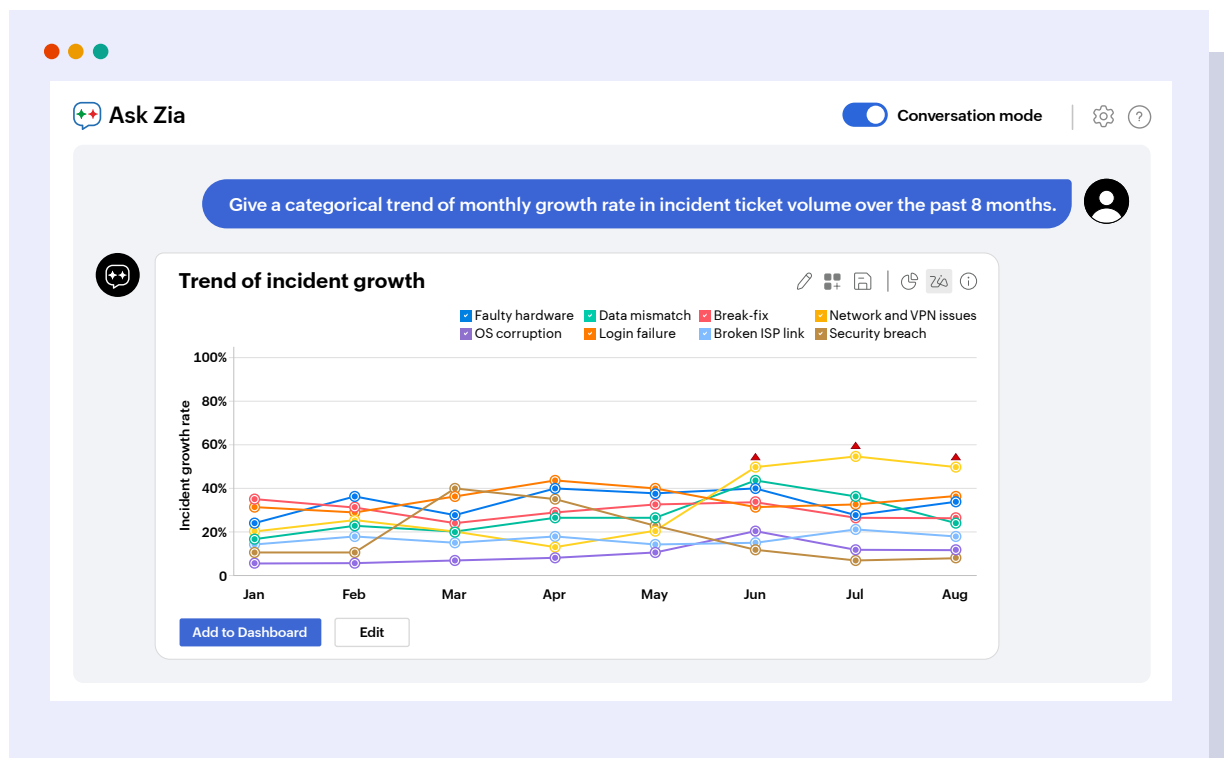
This crucial transition from reactive to predictive operations hinges on two complementary capabilities, powered by GenAI analytics:

- **AI-powered predictions:** Using real-time, unified analytics and ML-driven multivariate forecasting to anticipate a future spike in incident volume and preemptively address the underlying system issue.
- **Diagnostic insights:** Employing cross-functional data correlation and advanced GenAI-powered reasoning to understand the specific drivers behind an incident surge and stop the root cause.

When a service desk faces an unexpected spike in ticket volume, leaders need immediate and deep answers. Consider a service desk manager who notices a significant month-over-month increase in incident tickets.

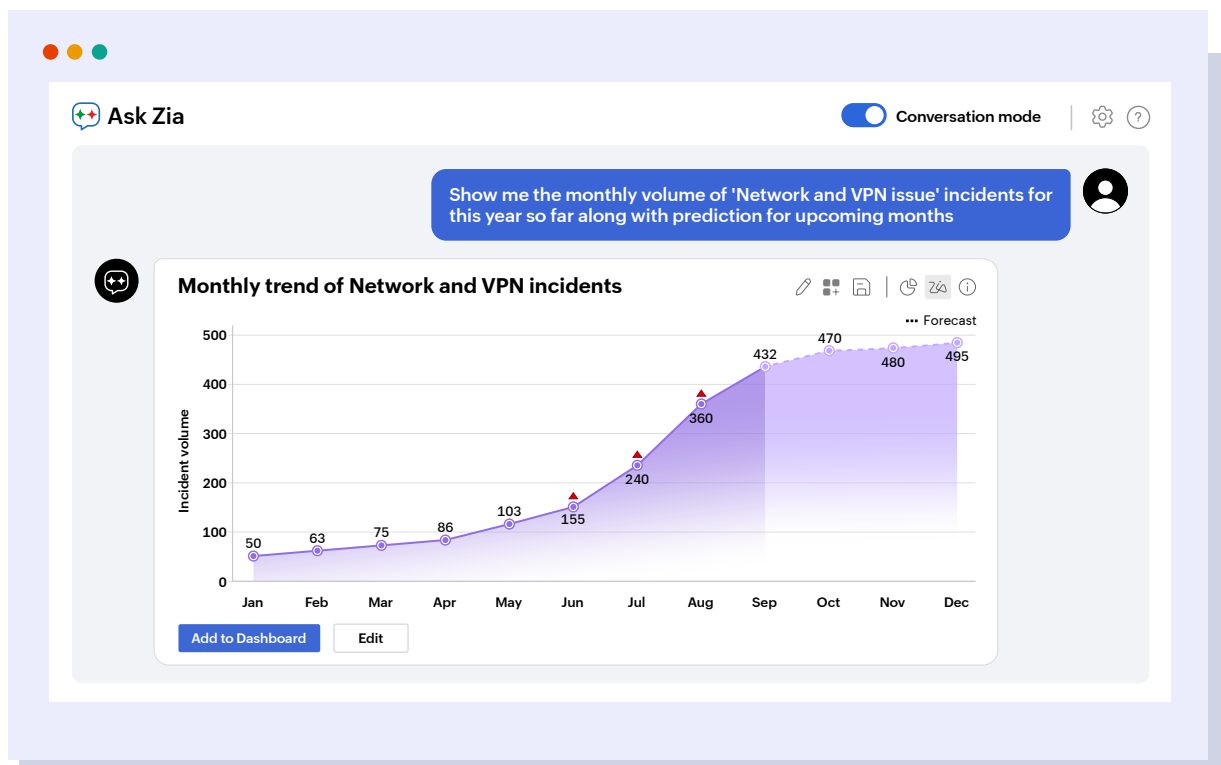
Leveraging Zia's conversational intelligence, they can perform a comprehensive analysis with a simple query:

"Give a categorical trend of month-on-month growth in incident ticket volume over the past eight months."



The Ask Zia snippet above instantly reveals a significant surge in the incident growth rate for "Network and VPN Issues" during specific months. Crucially, Analytics Plus' automated anomaly detection flags any surges as deviations from the expected pattern of incident growth.

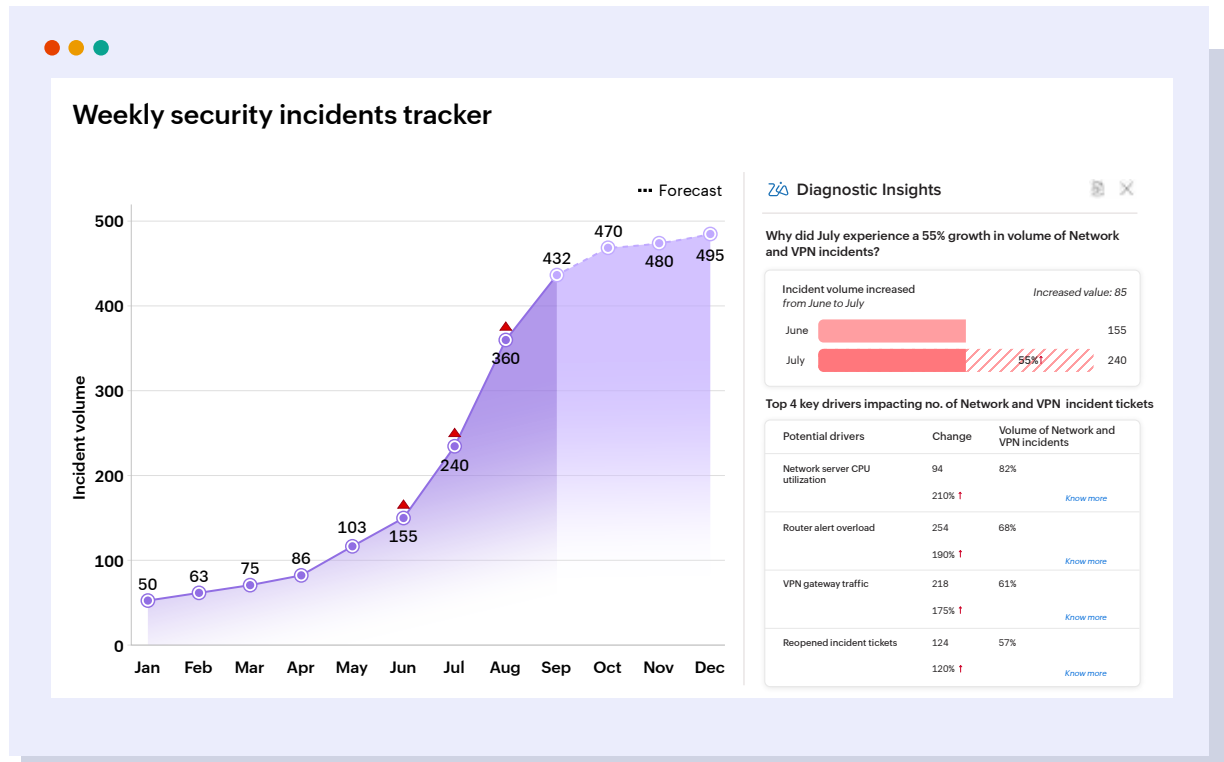
Service desk managers can now focus on this single problematic category to forecast its future trajectory.



The visualization shows the forecasted incident volume for upcoming months, highlighting a dire scenario: each month is expected to show a consistent growth in monthly incident volume.

While further RCA can be performed using conversational queries as demonstrated in the SLA section, GenAI-powered summarization of Zia insights deliver another visual and intuitive path to RCA insights.

Embedded directly within the report or dashboard, Zia's insights provide simplified visual and narrative summaries of the analysis, instantly uncovering the key drivers that led to the breach indicators. This enables intuitive and actionable RCA without typing out further questions.



The visualization uncovers the four key vulnerabilities driving the anomalous increase in Network incidents. IT managers can run similar diagnostics for both actual and predicted spikes to identify the specific root causes that have or will drive incident growth.

Finally, just as with SLA escalations, the Zia's advanced intelligent recommendation engine provides actionable, automated recommendations. These suggestions enable IT managers to implement targeted, proactive measures to curb network-related incidents for upcoming months, before a user is impacted.

The combination of precise data-driven forecasting, advanced diagnostic analytics, and conversational intelligence dramatically elevates enterprise incident management. This powers the shift from reactive ticket resolution to predictive problem anticipation.

By acting on these predictive signals, the service desk moves from fielding thousands of crisis calls to managing incidents preemptively. This prevents technician bandwidth exhaustion and counters potential SLA hits. The result is an intelligent, self-aware ITSM ecosystem where every incident spike can be understood, traced, and preemptively mitigated before users ever notice or a ticket is logged.

Conclusion

The next evolution of ITSM won't just be smarter, it'll be self-optimizing. With GenAI and Agentic AI advancements, service desks no longer need to wait for incident surges or performance drops. They'll sense, reason, and act autonomously.

By strategically adopting unified, AI-powered analytics, enterprises can turn every ITSM insight—from predicting SLA risks and technician inefficiencies to preventing incident surges and asset failures—into an intelligent action, efficiently transforming the service desk from a reactive support center into a strategic partner that drives uptime, resilience, and business growth.

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/

