

# Cybersecurity reimaged: **4 ways agentic AI strengthens enterprise defenses**

✦ GenAI-powered strategies for a resilient  
IT landscape

# Table of contents

|   |   |    |
|---|---|----|
| ■ | Introduction  | 3  |
| ■ | Reimagining security health monitoring for comprehensive cyberdefense | 4  |
| ■ | Proactively conquering shadow IT                                      | 8  |
| ■ | Ending privileged access misuse                                       | 12 |
| ■ | Streamlining patch deployment   | 16 |
| ■ | Conclusion  | 19 |
| ■ | About ManageEngine Analytics Plus                                     | 20 |

# Introduction

## The agentic AI revolution

GenAI in enterprise IT is evolving at a breakneck speed. What began with basic chatbots and evolved into conversational assistants has now culminated with the rise of agentic AI, one of the most disruptive forces in modern technology.

At its core, agentic AI is defined as GenAI-powered systems built to autonomously achieve complex goals, executing multi-step workflows with little-to-no human intervention. Unlike other AI technologies that relied heavily on predefined scripts or static models, agentic AI thrives on autonomy, deep reasoning, and adaptability, making it uniquely suited for dynamic IT environments.

Within enterprise IT, agentic AI can be purpose-built to break down complex incidents or critical security event response into smaller, actionable, and automated tasks. It performs deep reasoning across complex, disparate datasets, evaluates the findings, and can then either recommend or autonomously carry out the best course of action. This level of self-directed intelligence is a fundamental game-changer, especially in cybersecurity.

## Contextual intelligence for critical security challenges

For too long, traditional cyberdefense models were characterized by alert overload, frantic firefighting, time-consuming manual investigations, and static policies that consistently lagged behind evolving threats.

Agentic AI flips this model on its head. With its advanced autonomy, sharper monitoring, and deep reasoning capabilities powered by AI-driven analytics, it can transform cybersecurity from a brittle, rule-based system into a proactive, resilient, and adaptive defense framework.

This e-book deep dives into four critical cybersecurity strategies, illustrating how agentic AI transforms complex security analytics into simple conversational queries. You'll discover how advancements in GenAI can elevate cybersecurity into a proactive enabler of enterprise IT resilience.

01

## Reimagining security health monitoring for comprehensive cyberdefense

**E**nterprise cybersecurity is not a series of isolated battles; it is a continuous war on threats and vulnerabilities. While critical breaches and intrusions demand the immediate attention of security teams, one of the most underrated and often toughest challenges fought daily within any organization is the relentless, proactive effort to maintain a resilient security posture.

In today's hyper-connected cyber landscape, no enterprise can afford blind spots or hidden vulnerabilities that compromise its defenses. Yet, gauging an organization's true level of protection is notoriously difficult. Security teams spend hours buried in siloed tools—like SIEM logs, access reports, and vulnerability scanners—trying to piece together an answer to the question: "How secure are we?" The outcome takes the form of fragmented dashboards or complex reports that lack context and fail to provide a comprehensive view of enterprise risk.

This is a major area where advancements in GenAI, especially agentic AI and conversational analytics, are poised to revolutionize enterprise cybersecurity.

GenAI transforms security health monitoring into a continuous, intelligent, and action-oriented process. It leverages custom, no-code ML that intelligently correlate critical security parameters from across the entire IT landscape. This allows IT teams to:

- Visualize and predict the growth rate of various vulnerabilities and potential incidents.
- Assess their precise impact on the organization's security posture.
- Rank and score the organization's security health at any point in time.

What was once a reactive, manual process is now a continuous, intelligent cycle driven by AI-powered decision intelligence.

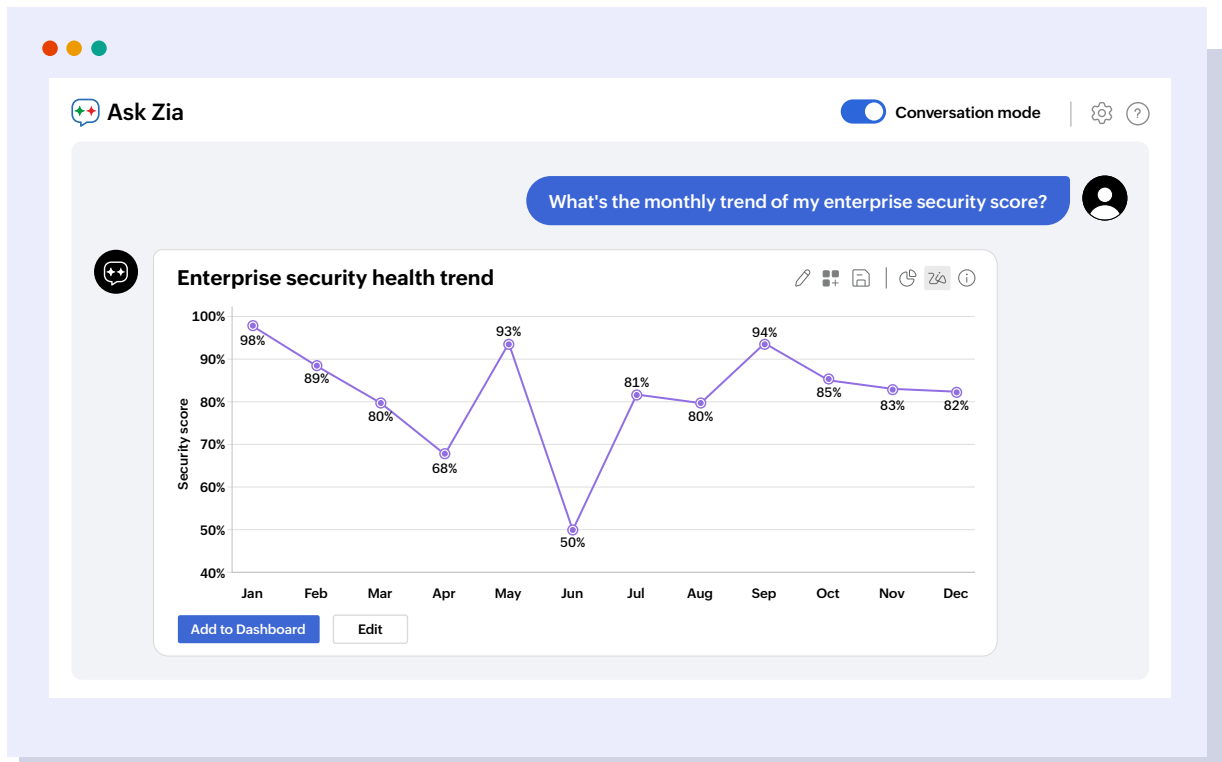
### ● **Ask Zia: The AI-powered security advisor**

Ask Zia, the conversational GenAI assistant in Analytics Plus, acts as a virtual cybersecurity strategist. Zia uncovers vital trends in infrastructure health and vulnerability status by tracking and analyzing key dimensions such as incident patterns, user access behavior, asset vulnerability, and policy compliance—all through a conversational interface.

By employing advanced no-code ML models, Ask Zia surfaces and tracks a tailored security score for any enterprise IT landscape. It also dynamically pinpoints the specific factors that are weakening your security posture over time. This crucial insight fuels tailored security strategies and dramatically improves your overall security footprint.

For example, a security analyst at a global enterprise can simply ask: “What’s the monthly trend of my enterprise security score?”

Within seconds, Zia generates a trend line analysis of the organization's overall vulnerability score. This score is calculated using a tailored ML model that integrates and correlates multiple factors (e.g., patch deployment status, incident density, configuration compliance) from all security tools. It provides an immediate sense of the organization's security health on a percentage scale, indicating the extent of exposure to critical threats and vulnerabilities.



The decision intelligence and GenAI capabilities of Zia go far beyond simple report generation. When a risk or anomaly is identified, Zia provides data-backed, accurate insights into the key factors driving the security score.

For instance, in the above analysis, the security score dipped below 50% in June. Now, the security analyst can ask: "Why is the security score low for June?"

Zia instantly analyzes the multitude of influencing factors and quickly surfaces the top three or four drivers bringing down the overall security score.

**Reason for decrease in organizational security score (Month: June)**

The key drivers that might have influenced the increase in no. of downtime incidents are:

1. Increase in risky unlicensed commercial software installations from *Asset vulnerability audit*.
2. Non-business hour access volume from *Identity and Access Management logs*.
3. Spike in repeated failed MFA attempts within privileged accounts from *Identity and Access Management logs*.
4. Elevated volume of critical endpoints without up-to-date security patches from *Endpoint Management*.

| Potential drivers                            | Contribution to decrease in Vulnerability score | Impact on Vulnerability score |
|--|---|-------------------------------|
| Unauthorized software installations          | 25%   | 90%                           |
| Non-business hour privileged access requests | 20%   | 80%                           |
| Failed MFA attempts                          | 18%   | 72%                           |
| No. of unpatched endpoints in last 30 days   | 30%   | 95%                           |

By gaining quick, data-driven insights into these key drivers, the IT team can accelerate resolution and focus their efforts on the root causes that truly matter, instead of investigating disparate issues and stretching their resources thin. However, conversational analytics with Zia doesn't just leave decision-makers lingering with raw, data-backed insights; it can further propose actionable, prescriptive strategies to improve the security score.

|  |     |     |
|--|-----|-----|
| Failed MFA attempts                        | 18% | 72% |
| No. of unpatched endpoints in last 30 days | 30% | 95% |

**Recommendations**

1. Expand endpoint discovery and remove unauthorized applications. Ensure all new hardware and software comply with the authorization protocols and baseline security criteria.
2. Conduct a privileged access audit, applying geo or time restrictions and password resets to prevent anomalous privileged account behavior.
3. Enforce multi-factor authentication organization-wide, with close monitoring of high-privilege accounts.
4. Automate patching for critical assets including servers, routers, and workstations, with special focus on high-risk nodes.

This progression, from uncovering what is happening and why to prescribing what to do about it, is the power of AI in cybersecurity.

While traditional security assessments merely document weaknesses, conversational analytics transforms them into a continuous cycle of intelligence and action. With Analytics Plus and Ask Zia, enterprises can finally shift from a reactive, alert-driven approach to a strategic, AI-enabled defense strategy. This isn't just assessment—it's decision intelligence in motion, shaping a secure foundation for the enterprise of tomorrow.

## 02

# Proactively conquering shadow IT

**E**very modern enterprise grapples with the chaotic dilemma of shadow IT. Driven by the rapid proliferation of AI tools and SaaS applications, employees today often install unapproved applications. While they help employees work better and faster, these shadow IT applications, when deployed outside organizational policies, governance, and necessary security configurations, introduce dangerous vulnerabilities that can result in data theft, compliance violations, unnecessary cost drains, and a vastly expanded attack surface vulnerable to insider threats and external attacks.

The biggest challenge for large enterprises isn't just that shadow IT exists—it's that most don't have comprehensive visibility into them. Shadow IT footprints are scattered across system logs, cloud services, network endpoints, and departmental assets.

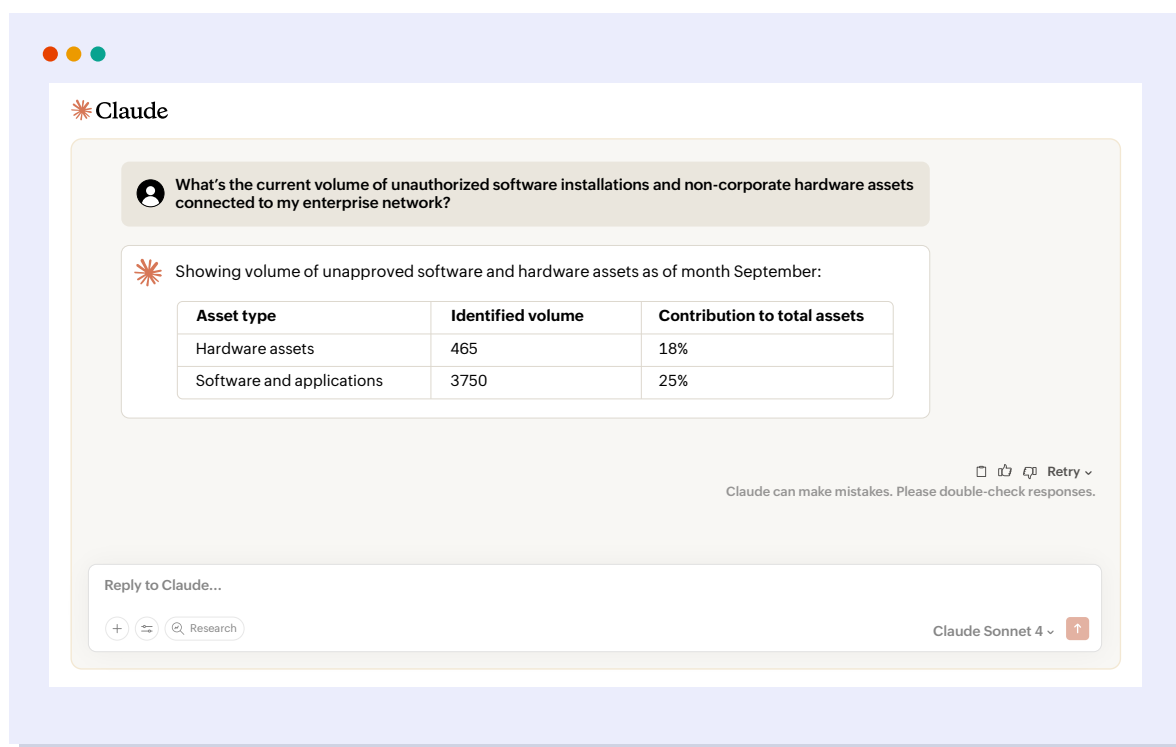


Even when some insights are available via siloed reports, correlating cross-domain and cross-departmental information to get a true picture of the shadow IT landscape can be painfully complex and time-consuming. By the time IT manages to uncover instances of shadow IT, the footprint may have already expanded dramatically, putting the organization at greater risk.

## Turning visibility into control with MCP and GenAI

This is where the Model Context Protocol (MCP) server in GenAI-powered IT analytics platforms like Analytics Plus changes the game. MCP support unifies intelligence from cross-domain reports, asset inventories, service desk tickets, network logs, usage patterns, and privileged access systems and makes them accessible through natural language conversational queries from the LLM interface of the user's choice.

Instead of manually digging through countless reports, IT leaders can simply ask: “What’s the current volume of unauthorized software installations and unapproved devices connected to my enterprise network?” Within seconds, the LLM via Analytics Plus surfaces precise, contextual insights.



Conversational analytics enables IT leaders to dig deeper and perform a granular dissection of shadow IT penetration within the enterprise. The visualization below shows how to perform in-depth analysis and uncover the RCA of unsanctioned software applications using natural language queries.

The screenshot displays the Claude AI chat interface. At the top, the Claude logo is visible. The first query asks, "Which departments are driving the most unsanctioned software installations?". The response provides a summary of the top 5 departments with the highest volume of unsanctioned software installations, categorized by department: HR and Recruitment, Marketing, Sales, Finance, and R&D. Each category lists the total number of unauthorized software, the average duration of usage, and the most popular unsanctioned software. The second query asks, "Which month saw the highest growth in unsanctioned software downloads?". The response states that after reviewing the unauthorized software growth trends in the Analytics Plus workspace, two months with the highest growth in download rate were identified: June and December, both showing a 108% M-O-M growth in downloads. A table is provided for these months. At the bottom, there is a text input field labeled "Reply to Claude..." and a button labeled "Research". The interface also includes a "Claude Sonnet 4" label and a "Retry" button.

**Which departments are driving the most unsanctioned software installations?**

Based on asset management data, here are the top 5 departments with the highest volume of unsanctioned software installations:

**HR and Recruitment**

- Total number of unauthorized software: 1240
- Avg. duration of usage: 365 days
- Most popular unsanctioned software: Zoom Meeting

**Marketing:**

- Total number of unauthorized software: 740
- Avg. duration of usage: 240 days
- Most popular unsanctioned software: Zapier

**Sales:**

- Total number of unauthorized software: 650
- Avg. duration of usage: 208 days
- Most popular unsanctioned software: Gong

**Finance:**

- Total number of unauthorized software: 300
- Avg. duration of usage: 108 days
- Most popular unsanctioned software: Microsoft OneNote

**R&D:**

- Total number of unauthorized software: 220
- Avg. duration of usage: 80 days
- Most popular unsanctioned software: Github Copilot

**Which month saw the highest growth in unsanctioned software downloads?**

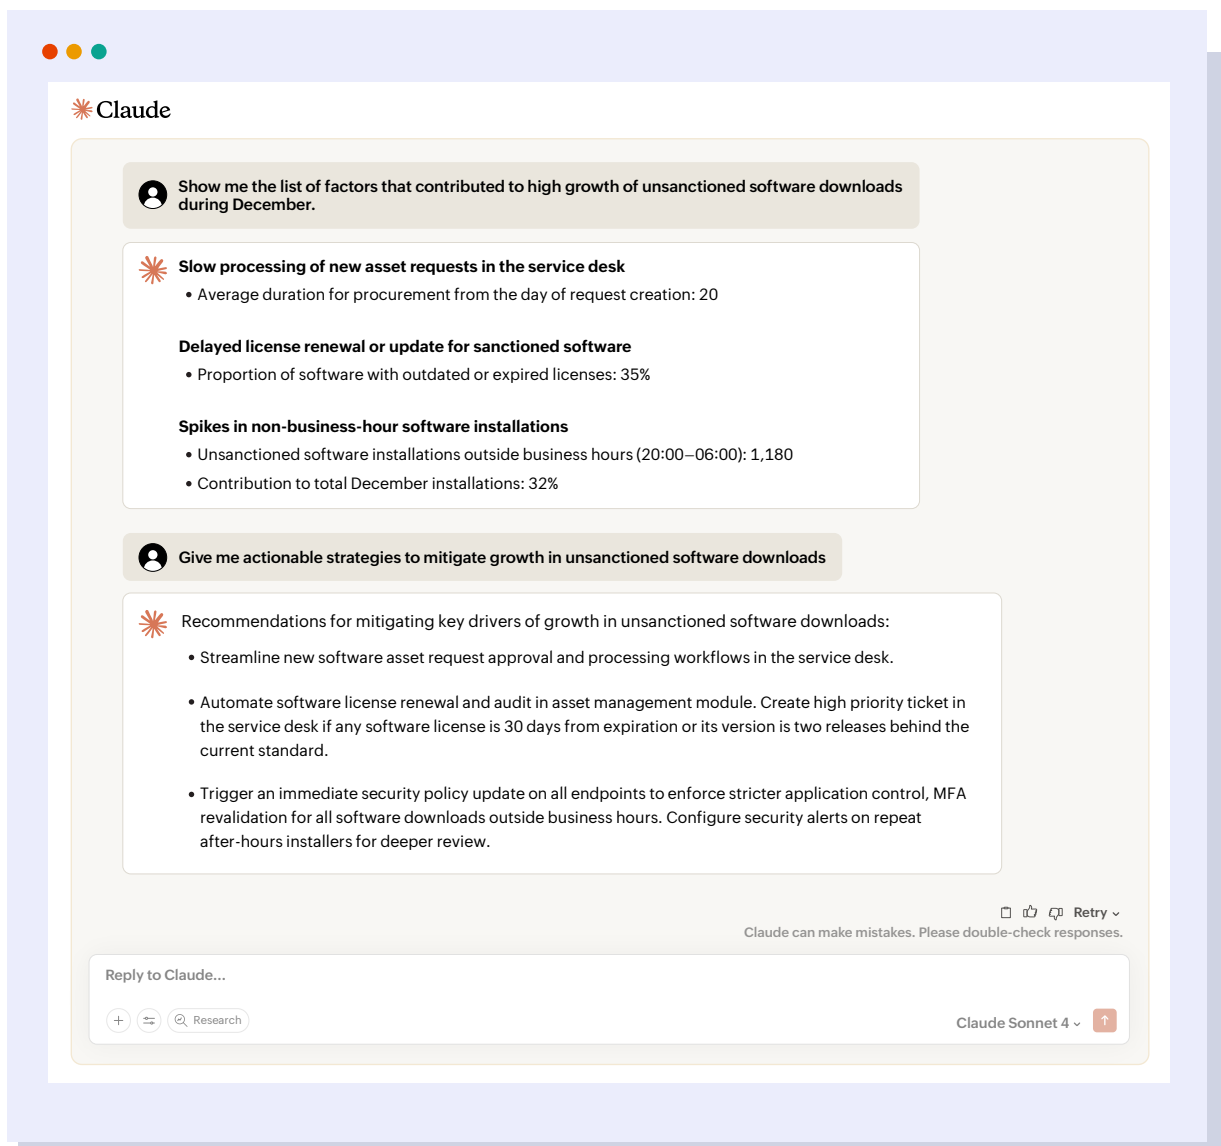
After reviewing the unauthorized software growth trends in your Analytics Plus workspace, I've identified two months with highest growth in download rate:

| Month    | M-O-M growth in downloads |
|----------|---------------------------|
| June     | 108%                      |
| December | 108%                      |

Reply to Claude...

Claude Sonnet 4

Similar to Ask Zia, the agentic AI and LLM integrations allow IT leaders to identify the driving factors or root causes behind the explosion of shadow IT and instantly gain tailored, data-backed recommendations to proactively prevent future incidents. RCA can be performed for shadow IT hardware assets, too.



These in-depth recommendations and data-driven insights form the foundation for automated shadow IT prevention. Provided these necessary applications and platforms are MCP-enabled, IT teams can execute these corrective actions—such as blocking an application or updating a policy—directly from the LLM screen. This creates a 360-degree, closed-loop automation, including insights, recommendations, and remediation, without ever switching to a separate application or performing time-consuming analysis and triage.

## Ending privileged access misuse

**P**rivileged access management (PAM) is the bedrock of enterprise IT security. Privileged accounts control the keys to critical systems—servers, databases, and enterprise applications—and they remain one of the most coveted targets for threat actors and the hardest to secure for IT teams. A single misuse or compromise can pave the way for devastating breaches, insider attacks, compliance failures, and costly business disruptions.

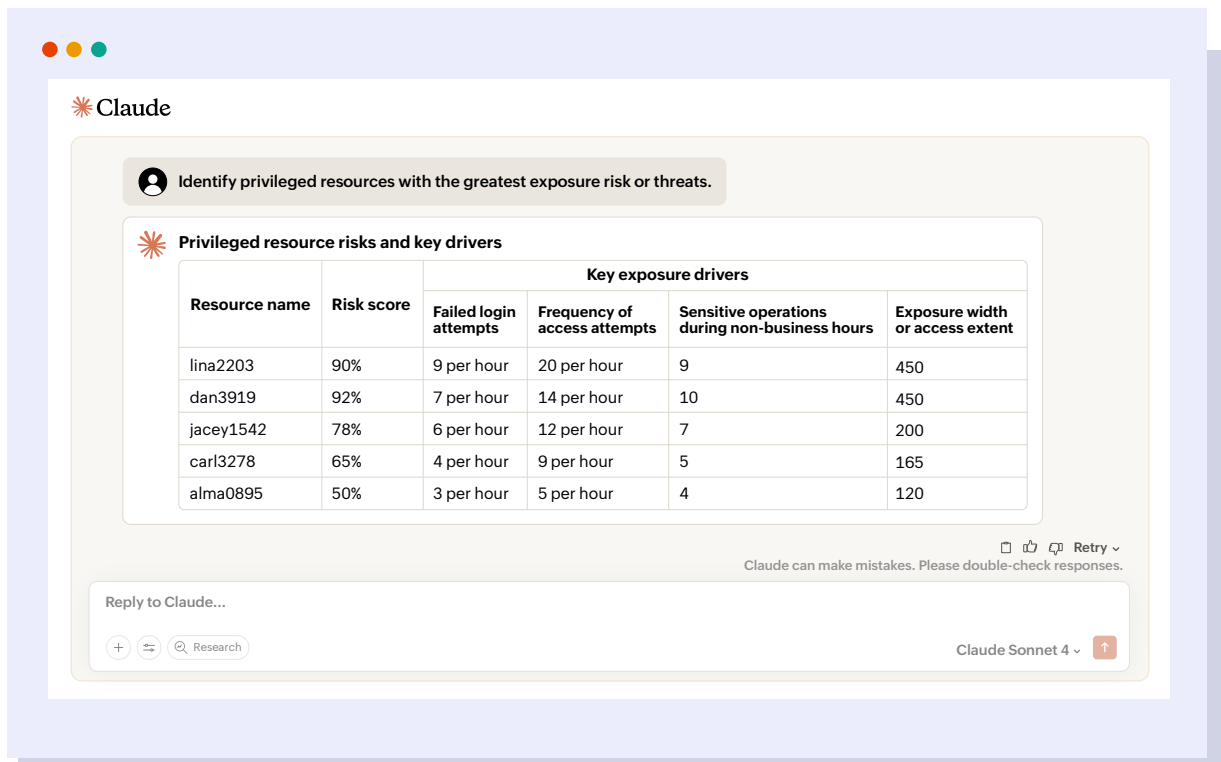
Dedicated PAM tools are vital; they are designed to spot policy breaches, accounts bypassing standard controls, and threats from within. However, the alerts and findings from these tools are often too little, too late. They typically lack the predictive intelligence required to correlate hidden behavioral patterns in advance and identify the resources or activities most likely to be exploited.

In today's dynamic IT environments—where privileged policies evolve quickly and access is distributed across on-premises, hybrid, and cloud infrastructure—organizations need PAM strategies that are context-aware, adaptive, and self-reliant.

Agentic AI reshapes PAM. By continuously monitoring privileged user behavior and learning what's normal, AI can flag deviations in real time, spotting malware injections, insider threats, and fraud attempts.

Instead of building complex dashboards or manually running scripts, security teams can now interact with an LLM interface to gain instant, forward-looking insights derived from complex cross-domain analysis.

Through the MCP server, these LLMs can link up with IT analytics platforms like Analytics Plus, creating a powerful agentic AI setup that can auto-analyze user access patterns and contextual operational data to detect subtle anomalies in privileged account behavior. By correlating multiple cross-functional factors, these AI agents can arrive at a comprehensive risk level for any privileged account. The LLM then returns quick, actionable insights in natural language on the top high-risk accounts.

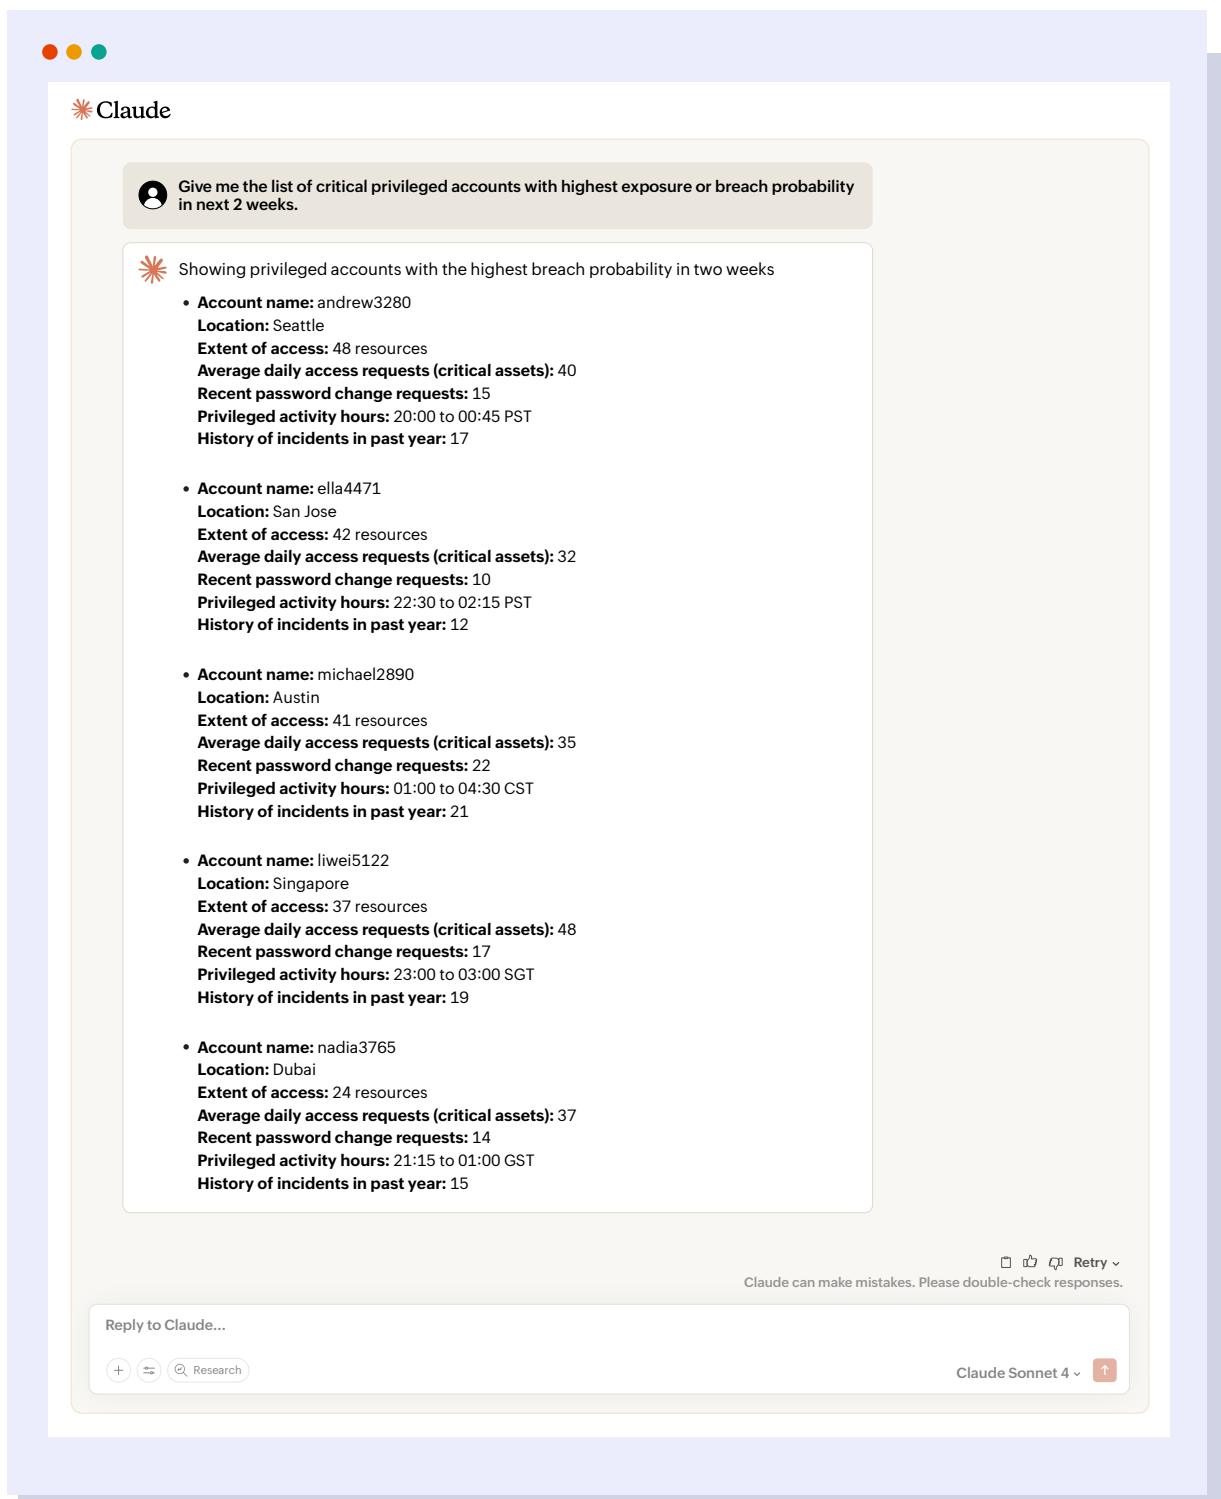


The screenshot shows the Claude AI interface. At the top, there's a header with the Claude logo. Below it, a prompt box contains the instruction: "Identify privileged resources with the greatest exposure risk or threats." The response from Claude is a table titled "Privileged resource risks and key drivers". The table has six columns: "Resource name", "Risk score", and four columns under the heading "Key exposure drivers": "Failed login attempts", "Frequency of access attempts", "Sensitive operations during non-business hours", and "Exposure width or access extent". The table lists five resources: lina2203, dan3919, jacey1542, carl3278, and alma0895, with their respective risk scores and exposure metrics. Below the table, there's a "Reply to Claude..." input field and a "Research" button. The bottom right corner shows "Claude Sonnet 4" and a "Retry" button.

| Resource name | Risk score | Key exposure drivers  |                              |  |                                 |
|---------------|------------|-----------------------|------------------------------|--|---------------------------------|
|               |            | Failed login attempts | Frequency of access attempts | Sensitive operations during non-business hours | Exposure width or access extent |
| lina2203      | 90%        | 9 per hour            | 20 per hour                  | 9  | 450                             |
| dan3919       | 92%        | 7 per hour            | 14 per hour                  | 10   | 450                             |
| jacey1542     | 78%        | 6 per hour            | 12 per hour                  | 7  | 200                             |
| carl3278      | 65%        | 4 per hour            | 9 per hour                   | 5  | 165                             |
| alma0895      | 50%        | 3 per hour            | 5 per hour                   | 4  | 120                             |

Having deep, instant insights into privileged account risks can significantly reduce the mean time to respond (MTTR) to potential malware or insider threats, minimizing cost and productivity drains.

Agentic AI goes further and brings advanced multivariate forecasting directly into a conversational LLM interface. This allows IT teams to predict which accounts are most likely to be compromised in the near future.



As seen here, in response to a query, the LLM not only pinpoints at-risk accounts but also identifies the driving factors of user behavior, contextual operating conditions, and historical anomalies, helping IT teams prioritize remediation with laser precision.

The entire predictive loop is achieved seamlessly within the LLM interface—from determining suspicious activity by correlating complex privileged user data to running a complex multivariate forecast, and finally, gaining insights in conversational language.

By deploying such intuitive agentic AI workflows that integrates LLMs with Analytics Plus workspace via the MCP server setup, enterprises can establish a proactive shield against privilege misuse.

Privileged accounts will always be tempting targets for threat actors. The key to resilient security lies in discovering misuse or anomalous user behavior to predict and stop any critical vulnerability before damage occurs.

By combining anomaly detection, forecasting, and automation through agentic AI, IT leaders can:

- Reduce MTTR to privileged threats.
- Cut costs and productivity losses from breaches.
- Build a resilient, predictive defense shield around the most sensitive accounts.

In short, agentic AI elevates PAM from a reactive control into a resilient defense shield for enterprise IT.

# Streamlining patch deployment

**R**egular patching cycles are the foundation of a resilient IT security and endpoint management strategy. When executed well, patches seal critical vulnerabilities, shrink the attack surface, and strengthen endpoint resilience. But when poorly managed, patching can just as easily become a security and operational liability.

## The hidden risks of patch mismanagement

Untested patches—those not thoroughly evaluated for compatibility, reliability, and security—can trigger cascading failures. A flawed patch not only disrupts the target endpoint but also destabilizes entire systems, causing outages, application crashes, and frustrated end users.

The challenge for IT teams lies in visibility. Traditional patching practices often measure success by whether the deployment was completed, not by its real impact on system health, application health, or end-user experience. Without deeper analytics, the hidden link between a patch rollout and the sudden rise of downtime or performance issues remains invisible until it's too late.

With conversational AI assistants like Ask Zia, IT teams can instantly correlate patch data with downtime logs, help desk incidents, and endpoint health trends via natural language queries. This delivers real-time visibility into how patches affect system stability without having to dig through siloed tools and complex system logs.

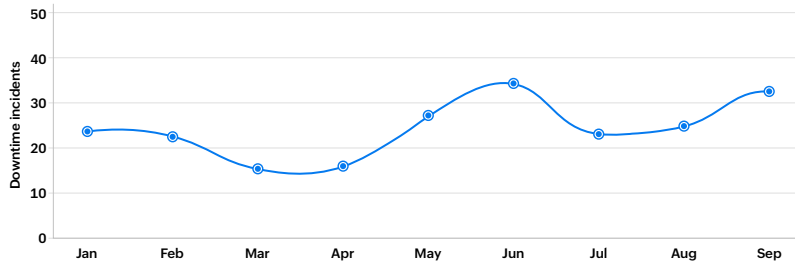


Show downtime incident trends for the cloud drive application over the past 9 months.



Trend of cloud drive downtime

✎ 📊 📅 📄 📈 ⓘ



Add to Dashboard

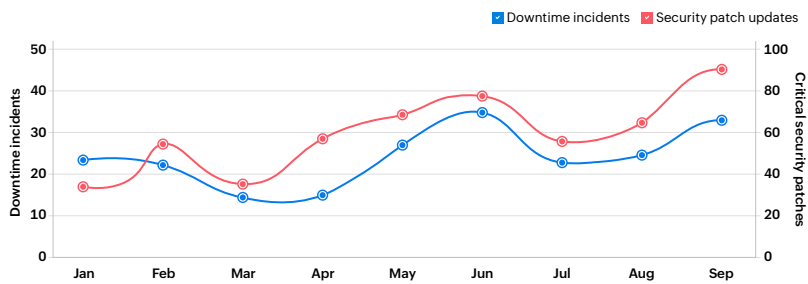
Edit

Compare this with the number of security patches deployed.



Cloud drive: Trend of downtime incidents vs. Security patch volume

✎ 📊 📅 📄 📈 ⓘ



Add to Dashboard

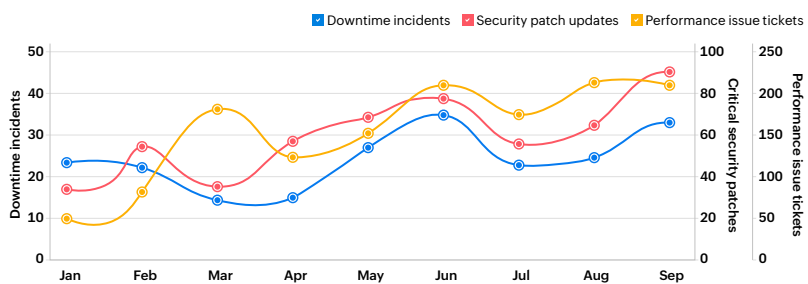
Edit

Add service desk incidents tagged under performance issues.



Impact analysis of cloud drive security patch deployments

✎ 📊 📅 📄 📈 ⓘ



Add to Dashboard

Edit

From the above visualization, the direct, positive correlation between three disparate metrics is clear. As patch volume increased, downtime and performance issue tickets also spiked, highlighting how flawed patch deployments are contributing to deteriorating application health.

With the correlation clearly established, IT teams can now leverage Zia's predictive prowess to observe how continued flawed deployment practices can impact future performance.



The forecast above shows a high probability that the number of high-priority downtime incidents will significantly exceed the average in November and December. This spike coincides directly with the increase in planned security patch rollouts, clearly indicating a crucial need for better pre-deployment testing before year-end activities.

Predicting future risks in patch deployment allows IT teams to strategically plan improved patch testing and streamline deployment. As seen in the case of organizational security health monitoring, Zia's advanced key driver analysis and intelligent recommendation engine helps IT leaders proceed with further steps: performing a key driver analysis to pinpoint the root causes (e.g., lack of sandbox testing) and delivering actionable, automated recommendations to streamline the process.

When integrated with MCP-supported tools, Zia's insights trigger immediate corrective actions, creating a closed-loop automation cycle. The result? Patch management evolves from a hidden risk into a predictable, resilient process—ensuring applications stay future-ready, stable, and secure.

## Conclusion

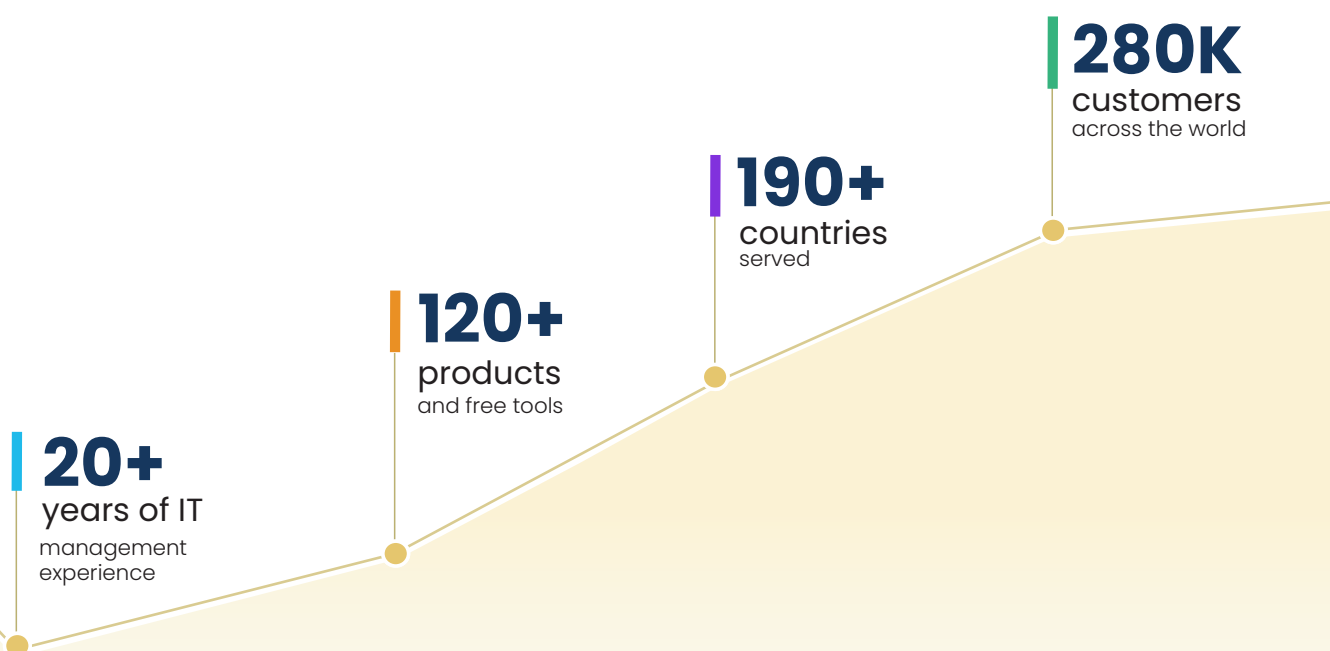
The ultimate value of agentic AI in enterprise IT security is not just speed of resolution; it is unburdening your security team to focus on strategic initiatives while the AI handles the complexity of correlation, prediction, and instantaneous remediation. This is the future of the resilient enterprise: A secure environment where threat intelligence is not a report to be read, but an action to be executed.

# 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/](http://www.manageengine.com/analytics-plus/)





© ManageEngine, a division of Zoho Corporation