

5 HIDDEN PITFALLS IN YOUR hybrid infrastructure management strategy

- A deep dive into the asset management inefficiencies sabotaging your hybrid landscape

Table of contents

■	Introduction	3
■	Not tailoring infrastructure configuration to usage needs	4
■	Failing to anticipate usage spikes in the cloud resource procurement plan	7
■	Overpaying for licenses and subscriptions	10
■	Ignoring department-level usage patterns	13
■	Using multiple tools for inventory monitoring	17
■	Conclusion	19
■	About ManageEngine Analytics Plus	20

Introduction

Asset management is a cornerstone of the IT service landscape—a vital responsibility that's been safely entrusted to and seamlessly carried out by IT technicians for decades. Yet, as cloud adoption surges and organizations transition to hybrid infrastructures, managing IT assets, particularly software, is becoming more complex.

The distribution of software modules and workloads across on-premises, public, and private cloud environments adds a layer of intricacy to enterprise digital transformation efforts. To excel in managing these assets, IT teams must evolve their approach beyond procuring, deploying, and managing hardware and software assets. IT teams must now also oversee underlying infrastructure, monitor their usage, and control costs effectively.

However, this modern hybrid landscape is rife with inefficiencies that drain costs and hinder operational efficiency and productivity. While most IT teams have grasped the importance of a robust hybrid asset management strategy, many are still navigating the challenges.

This e-book delves into common inefficient asset management practices in hybrid IT landscapes and offers insights on overcoming them to establish streamlined and effective IT service operations for the modern digital enterprise.

Not tailoring infrastructure configuration to usage needs

Developing a prudent and efficient software asset procurement strategy is a critical responsibility for IT teams. A robust procurement strategy that factors in organizational considerations, like usage frequency and duration, allocated budget, and more, can streamline an organization's overall IT spending and asset management. However, organizations' efforts to optimize asset procurement will fall short unless the prevalent inefficiencies related to underlying infrastructure are addressed.

While there are various approaches to streamline infrastructure procurement, the primary focus should be addressing the root cause of cost drains—a lack of visibility and alignment between software usage and its underlying infrastructure.

Software applications typically comprise multiple interconnected components, including compute, storage, access, network, database, and more, which were traditionally deployed on dedicated on-premises infrastructure like application servers, web servers, and virtual machines, or hosted entirely on public or private clouds.

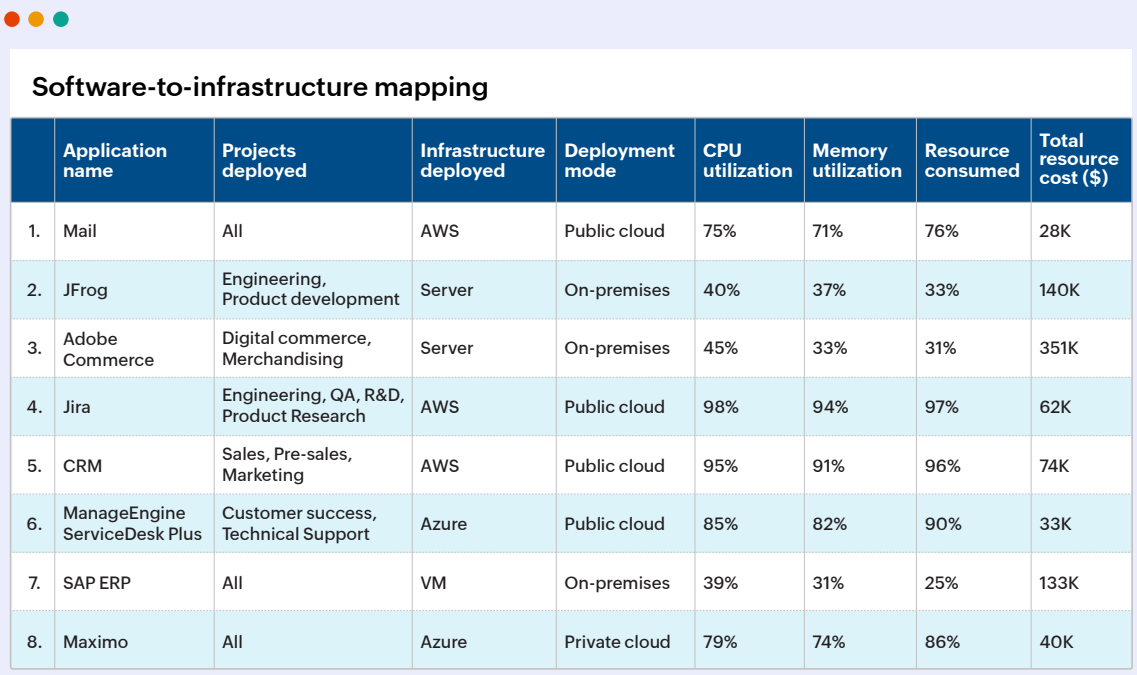
This traditional approach has now evolved in keeping with the flexibility and scalability of modern hybrid environments. To harness the full potential of this evolving IT landscape, IT teams are beginning to distribute software workloads across hybrid resources for optimized results and cost savings.

However, lack of visibility into the infrastructure elements consumed by software and the subsequent usage patterns often leads IT teams to procure and allocate more resources than necessary. This results in overprovisioning, particularly in hybrid environments.

Overcoming provisioning issues with software-to-infrastructure mapping

Overprovisioning can significantly inflate costs and undermine the benefits of a hybrid, multi-tenant infrastructure. On the flip side, resource management inefficiencies can also lead to underprovisioning, which can cause performance issues with the associated software.

Creating a detailed software-to-infrastructure mapping can provide IT teams with full visibility into their software assets and associated infrastructure, ensuring complete visibility into how the infrastructure and apps work together for a particular technology.



The illustration shows a window titled "Software-to-infrastructure mapping" containing a table with 8 rows and 9 columns. The columns are: Application name, Projects deployed, Infrastructure deployed, Deployment mode, CPU utilization, Memory utilization, Resource consumed, and Total resource cost (\$). The rows list various applications like Mail, JFrog, Adobe Commerce, Jira, CRM, ManageEngine ServiceDesk Plus, SAP ERP, and Maximo, along with their deployment details and resource costs.

	Application name	Projects deployed	Infrastructure deployed	Deployment mode	CPU utilization	Memory utilization	Resource consumed	Total resource cost (\$)
1.	Mail	All	AWS	Public cloud	75%	71%	76%	28K
2.	JFrog	Engineering, Product development	Server	On-premises	40%	37%	33%	140K
3.	Adobe Commerce	Digital commerce, Merchandising	Server	On-premises	45%	33%	31%	351K
4.	Jira	Engineering, QA, R&D, Product Research	AWS	Public cloud	98%	94%	97%	62K
5.	CRM	Sales, Pre-sales, Marketing	AWS	Public cloud	95%	91%	96%	74K
6.	ManageEngine ServiceDesk Plus	Customer success, Technical Support	Azure	Public cloud	85%	82%	90%	33K
7.	SAP ERP	All	VM	On-premises	39%	31%	25%	133K
8.	Maximo	All	Azure	Private cloud	79%	74%	86%	40K

As shown in the above illustration, by mapping the software to their corresponding infrastructure components and analyzing their usage patterns, IT managers can gain a clear picture of the resource demands and inefficiencies in the current resource allocation for each application. Armed with this insight, IT teams can enable better resource deployment and cost management practices.

For instance:

- **Consolidating application workloads or moving to the cloud:**

Applications like SAP ERP, Adobe Commerce, and JFrog that are currently running completely on-premises with less than 50% CPU utilization can benefit from consolidation onto fewer servers, which can result in reduced hardware installation and maintenance costs. Or, the organization can move the workloads of these applications to the cloud. The pay-per-use model of the cloud provides IT teams with the flexibility to quickly scale resources up or down based on application usage or workload fluctuations. This helps organizations save on the upfront sunk costs and maintenance overhead associated with underutilized on-premises hardware.

- **Aligning cloud resource procurement and allocation with usage:**

On the other hand, applications like JIRA and CRM, which are hosted on public AWS Cloud instances, seem to have insufficient resources allocated, with average consumption exceeding 95%. This can lead to performance degradations. IT managers should consider purchasing additional instances and setting optimum resource benchmarks after careful assessment of performance changes post infrastructure upgrades.

- **Migrating critical high usage applications to private cloud:**

Considering the high resource requirements of these applications, an alternative approach is migrating them to a private cloud. This allows in-house IT teams to have better control over usage tracking and optimization, resulting in savings in the long run. Private cloud environments strike a balance between the flexibility of public clouds and the control of on-premises systems, making them ideal for high-priority applications with higher utilization, which could inflate costs when run in a public cloud.

Choosing the right infrastructure based on resource consumption and criticality of software applications can help organizations optimize infrastructure performance while achieving significant cost savings, ensuring that resources are allocated based on actual needs rather than assumptions.

02


Failing to anticipate usage spikes in the cloud resource procurement plan

Despite allocating appropriate cloud infrastructure based on intended software usage to prevent resource wastage and infrastructure cost overruns as recommended in the first strategy, IT teams can still encounter an unforeseen increase in cloud asset and infrastructure costs.

This can happen when the deployed software or applications suddenly experience unplanned spikes in usage. A commonly observed reason for this phenomenon is when IT teams fail to account for major planned business events and strategic initiatives in their procurement planning and budgeting. This inefficient practice can result in additional last-minute cloud resource procurement to handle increased workloads during these events. Such ad-hoc purchases can strain already stretched IT budgets and have longer lead times, disrupting application performance and further increasing costs.

Optimizing cloud spending with RIs and SIs

To overcome such budget inefficiencies, IT managers can use cost saving options like reserved instances (RIs). These are offered at a discount compared to standard instances due to long-term commitments and can be deployed with shorter lead times. With RIs, organizations can ensure optimal resource availability without disrupting application performance or the IT budget.




This analysis provides valuable insights into using RIs to reduce cloud costs and optimize resource utilization. It examines cloud infrastructure costs for an e-commerce application that experiences consistent traffic throughout the year, with peaks during the holiday and sale seasons in November and December. This is evidenced by the anomalous increase in cloud costs during these months, which can be tied back to the last-minute, large-scale purchase of additional on-demand instances.

Therefore, to manage increased workloads during peak periods without incurring unplanned costs from on-demand instance purchases, this organization should proactively plan and utilize RIs for their seasonal workloads.

Employing RIs efficiently can help stabilize cloud spending to acceptable levels and ensure seamless operations, as clearly indicated by the drop in cloud instance costs during November and December 2023, post RI usage.

RIs work best in scenarios where IT teams have a fair idea of peak usage periods and usage is fairly periodic otherwise. If teams are unsure of workloads during the peak periods, RIs can go underutilized and result in cost wastage, which can be monitored through the following analysis.



To overcome this, IT teams should reserve RIs for critical applications with consistently high usage, predictable traffic peaks, and scenarios where any downtime would cause significant business impact.

For non-critical workloads with fluctuating usage, spot instances (SIs) are a cost-effective solution. These are available at substantial discounts but can be terminated with little notice, making them ideal for flexible tasks that can handle interruptions.

Deploying a balanced combination of planned on-demand instances, RIs for critical applications, and SIs for non-critical tasks ensures efficient infrastructure spending, optimizes resource allocation, and reduces budgetary risks.

Overpaying for licenses and subscriptions

As seen in the previous section, the purchase of software and associated infrastructure resources constitute major capital expenditures (CapEx), with inefficiencies in procurement processes often leading to budget wastage.

An associated operational expense that frustrates IT teams during the budgeting phase is the recurring costs involved in managing licenses throughout the software life cycle. Effective software license management (SLM) requires proactively handling all aspects of software assets, from subscriptions and renewals to upgrades and more.

Key challenges or inefficiencies in SLM include:

- Lack of comprehensive visibility into all licenses and software.
- Gaps in license purchase and actual usage.
- Complicated licensing models and confusion over choosing the best fit.
- Choosing the right vendors, deployment modes, and negotiating optimal agreements.
- Misalignment between software licenses and required cloud service subscriptions, leading to the overallocation or underallocation of resources.

Addressing these inefficiencies is crucial for efficient and cost-effective SLM, helping IT managers streamline budget allocation and reduce overall asset spending.

Conversely, neglecting these inefficiencies can drive up the costs of both on-premises and cloud software licenses, as well as cloud service subscriptions, further contributing to the two major cost-drivers in SLM:

- The purchase, maintenance, upgrade, and renewal of software licenses.
- The acquisition and management of cloud infrastructure to host applications.

To reduce licensing costs, IT teams need to refine their license management approach for the hybrid ITSM landscape. While there are several strategies to address the ever-increasing software license costs, one effective, but often overlooked strategy is carefully evaluating each software's functionality and intended use, and then procuring licenses based on specific needs.

By aligning license purchases with actual usage requirements, organizations can avoid redundant or unnecessary software, leading to substantial cost savings. This strategy can significantly curb both of the major cost drivers listed above.


● Addressing software redundancy to eliminate unnecessary licensing spends

Software redundancy is widely observed inefficiency across organizations.

The enterprise software market is saturated with a proliferation of new tools and applications, many of which are close copies of existing tools with only a few additional features that add little value, if any. The abundance of choice can lead teams to experiment with multiple options, often stemming from a fear of missing out on the latest and greatest solutions. This ultimately results in redundant software purchases that create unnecessary costs for the organization.

Organizations frequently buy multiple software applications with overlapping functions for different projects or departments. IT teams, responding to user requests, acquire licenses for new software that duplicate existing capabilities. This results in wasted licenses, reduced usage of existing software, and limited exposure to the full capabilities of existing tools—all contributing to a significant drain in IT budget.

The below analysis outlines the impact of software license redundancy on the IT budget. Here, we analyze the monthly consumption of all actively deployed IT assets, categorized as Collaborative design tools.



The organization has multiple product teams that rely on critical collaboration, wireframing, and mind mapping tools. In response to user requests, the IT team has procured and deployed licenses for multiple collaborative design tools across the organization.

The analysis clearly shows that the IT team is expending significant annual licenses on Miro and Whimsical, applications that are commonly used for white-boarding and wireframing, despite their minimal usage. However, these functionalities can also be found in Figma, an application predominantly used by the design and UX teams, which currently has a significant volume of unused licenses.

To avoid redundant license purchases and the under utilization of existing tools, IT teams need to categorize and catalog software based on the functionality offered. Running an analysis under each category will easily highlight redundant software purchases. The audit will help teams make informed decisions on new software license procurement and help guide users towards existing tools.

Overall, the analysis shown above can be a great step towards helping IT teams inculcate efficient software license management practices and optimize license costs.

04


Ignoring department-level usage patterns

Clearly, IT asset management in a hybrid infrastructure landscape is highly complex, not unlike navigating a running track riddled with hurdles of inefficiencies. IT teams have to navigate these numerous hurdles quickly and efficiently to establish an optimized IT asset landscape.

While solving procurement and licensing inefficiencies removes some hurdles, there are more still. Deployment optimization, the third stage in IT asset management strategy, warrants IT teams to continuously monitor and streamline asset utilization as well as costs across hybrid landscapes.

Inefficient utilization of deployed resources can lead to unnecessary wastage and increased operational costs. According to **Flexera's recent state of cloud report^[1]**, nearly 32% of cloud spending is wasted or underutilized, highlighting the need for precise cloud asset management. Unused assets, whether hardware, software, or infrastructure, are major cost drains, silently wasting IT dollars and leading to duplicate purchases.


By leveraging AI-powered analytics, IT teams can track asset utilization, set automated alerts for underutilized or overutilized resources, and implement timely adjustments. This proactive approach ensures efficient and cost-effective asset management, transforming the hurdle-ridden IT landscape into a streamlined and optimized environment.



The visualization above tracks the department-wise usage, user activity, and license allocation for a project management application and compares it with the overall cost expended from the IT budget for the same.

The sales and marketing department has lower active users and usage compared to other teams, despite raking up a high license fee, while the IT department shows significant users but low activity. This discrepancy may stem from using other tools daily and the project management tool sporadically for specific tasks like month-end audits.

The inefficiency becomes glaringly profound when analyzing the department-wise influx of license requests in the IT service desk. This highlights the need for better resource allocation to improve utilization efficiency and cost-effectiveness.




This analysis shows that departments with high software usage, such as engineering, R&D, and product, have a high number of new license requests. This indicates the software is under-licensed in these areas, creating a supply-demand gap with potentially catastrophic business consequences.

Implementing department-specific budget limits to reduce IT resource wastage

To address this, IT teams can adopt a modified version of the chargeback model for resource budget allocation. The chargeback model involves directly billing departments or business units for their specific IT resource usage. This approach encourages departments to be more responsible with their IT budgets and resource planning, promoting greater accountability for resource utilization and reducing waste and inefficiencies.

However, it is recommended to implement department-specific budget limits when following this model. For departments where historical data shows wastage of assets like software licenses, more stringent limits can be implemented. This means that when they reach this limit, future allocation or requests for new assets will not be approved.



The above visualization offers transparency into department-specific budget usage patterns over the last few quarters, enabling more efficient resource and budget allocation.

By comparing this analysis with the previous insights, IT leaders can uncover departments that underutilize allocated resources yet consume a significant chunk of their allocated funds. Departments that consistently fall under this category, like sales, would benefit from lower budget thresholds. Conversely, departments that have optimized their asset and budget utilization, like R&D and engineering, can maintain higher available budgets.

By gaining a detailed breakdown of department-wise asset usage and associated expenses, and by establishing budgetary constraints aligned with actual asset consumption, IT teams can promote a more accountable and effective utilization of IT resources. It aids in pinpointing and remedying instances of underutilization and budgetary inefficiencies.

Using multiple tools for inventory monitoring

Regular monitoring and timely maintenance is a critical part of IT asset life cycle management. During this fast-paced journey, multiple maintenance challenges can arise.

Regular maintenance ensures the critical assets that are part of an organization's infrastructure remain available and efficient throughout the life cycle with minimal disruptions, downtimes, or performance issues. Neglecting this can lead to breakdowns and costly repairs, similar to how in a race, a lack of regular preventive pit stops can lead to longer repair periods at necessary pit stops, completely derailing the driver's momentum.

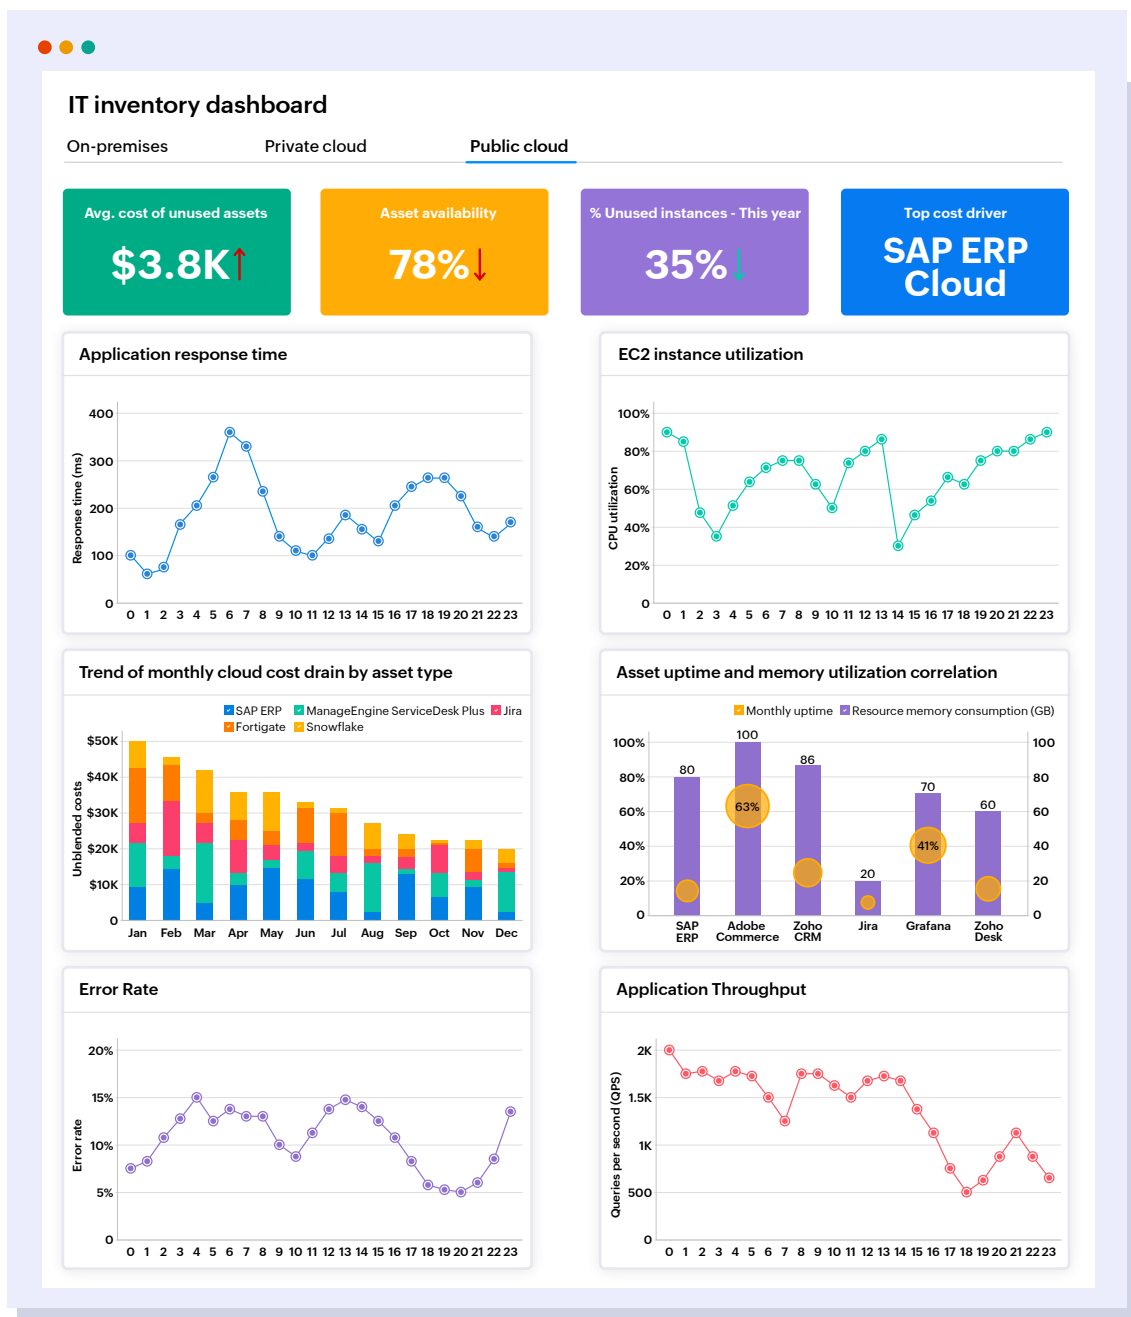
Monitoring key asset health and performance parameters in real time is one of the most effective and widely adopted maintenance techniques by IT teams. This proactive approach helps identify potential issues before they occur, or at least reduces their impact and associated downtime through timely detection and resolution.

Traditionally, IT technicians will toggle between multiple monitoring applications to track software, hardware, and infrastructure performance. However, this approach eats up valuable time, and vital insights indicative of impending disruptions often slip through the cracks. In addition to slowing down technicians, by operating with a siloed approach, IT teams often fail to leverage the collected data to its full potential.

Analyzing asset data in isolation without considering the broader context delays the resolution process, as IT teams struggle to dig into underlying infrastructure data and correlate it with organization-wide insights to truly understand and address underlying problems.

Creating a unified dashboard for comprehensive IT inventory monitoring

To effectively monitor and maintain asset health, IT teams need a single, centralized dashboard that consolidates data from all assets across the IT landscape onto a single platform. This single source of truth allows teams to analyze, visualize, and correlate asset data, providing the actionable, real-time insights required to manage and optimize their asset footprint.



Unified IT inventory dashboards offer end-to-end insights into the health and performance of IT assets across the hybrid infrastructure. The above dashboard seamlessly aggregates data from the IT service desk, network and application monitoring apps, and other infrastructure management tools. By correlating this data, it delivers critical real-time insights into the entire application stack, providing unprecedented visibility and transparency into the underlying infrastructure.

By eliminating the need to switch between multiple, siloed infrastructure-specific monitoring tools or windows, unified IT inventory dashboards can save time and effort. This allows IT teams to seamlessly and easily monitor and manage their IT assets.

Conclusion

Inefficient hybrid asset management practices can deplete an organization's IT budget, hamper operational efficiency, and undermine productivity, preventing the achievement of both operational and business goals. However, by implementing the above five strategies and leveraging AI-powered analytics, organizations can effectively manage IT assets in a hybrid infrastructure, optimize spending, and ensure optimal resource utilization.

About

ManageEngine Analytics Plus is a self-service, AI-driven IT analytics solution that helps organizations implement complex initiatives to address the requirements of expanding businesses. Available on-premises and in the cloud, Analytics Plus visualizes IT data from several applications and integrates out of the box with several popular IT applications such as ManageEngine ServiceDesk Plus, Jira, ServiceNow, Zendesk and ManageEngine Endpoint Central. Analytics Plus features an AI-powered analytics assistant that responds to voice and text prompts to provide meaningful visualizations. This eliminates the need for a data analyst to aid IT managers and reduces report building time while enabling organizations to make faster, data-driven decisions.

Kick-start your IT analytics journey with a free trial of Analytics Plus.

Want to learn more about the product before giving it a try?

Sign up for a free, virtual tour with one of our solution experts.



Reference

1. <https://resources.flexera.com/web/pdf/Flexera-State-of-the-Cloud-Report-2022.pdf?elqTrackId=414badd9b3cd4eee979d7f8bbfa8269e&elqaid=6925&elqat=2>



© ManageEngine, a division of Zoho Corporation