

How to declutter performance metrics

and gain insights through the noise

- Essential strategies to avoid alert fatigue and ensure seamless IT operations.

Table of contents

■	Introduction	3
■	Crafting a single pane of clarity for comprehensive intelligence	4
■	Proactively triage recurring alerts	6
■	Mitigating alert noise from faulty devices	10
■	Revisiting outdated thresholds that undermine alert quality	12
■	Embracing MTTR and MTBF for continuous operational improvement	15
■	Conclusion	18
■	About ManageEngine Analytics Plus	19

Introduction

As the enterprise IT landscape continues to evolve and become more dynamic, the sheer volume of data generated from IT operations can overwhelm even seasoned IT leaders. With observability going mainstream, daily IT operations are inundated with a barrage of alerts that leaves IT teams drowning in a sea of information.

Discerning meaningful patterns and critical insights amidst the chaos of information overload has become paramount for IT teams in their pursuit for seamless operations. To alleviate the challenge of alert noise and data deluge, IT managers need to overhaul their operations—which is reactive by design—to a proactive approach that monitors usage and performance trends to preempt potential faults and down times in IT infrastructure.

This e-book looks at detailed analytics-powered strategies to help organizations sift through alert noise, uncover hidden patterns, and gain unparalleled clarity into their IT operations.

Crafting a single pane of clarity for comprehensive intelligence

RReal-time, in-depth operational intelligence around application health has emerged as the primary way for IT teams to streamline application functionality and avoid needless alert chaos. Individual IT infrastructure components form the backbone of an application and play an important role in determining its performance. Therefore, to gain a holistic understanding of an application's health, it is necessary to monitor the performance and health of the components that shape its infrastructure.

Comprehensive dashboards driven by advanced analytics have emerged as an indispensable tool for accurately gauging application performance in real time by visualizing key infrastructure components in a single interface. Dedicated application-specific dashboards allow for targeted monitoring of important metrics associated with these infrastructure components, enabling IT teams to gain better clarity into application performance and enhance operational efficiency.



Unified application health and performance tracker

App performance

End-user experience

Middleware

DB cluster health

Application availability

70% ↑

Last month: 55%

Avg. application uptime

60% ↑

Last month: 48%

Alarm volume

2% ↑

Compared to last month


Daily outages

4 ↓

Last month: 7


Server – CPU and memory utilization

CPU utilization Memory utilization




Server – Disk I/O

Disk read OPs (avg. operations/sec) Disk write OPs (avg. operations/sec)




Network device – Traffic trend

Network in (bytes/sec) Network out (bytes/sec)




Application server – Response time




DB server – Cluster query performance

No. of query executions Wait time per query execution



Application server – Downtime and request volume correlation



This consolidated application health and performance monitoring dashboard provides nuanced insights into multiple applications parameters, with individual tabs demonstrating critical health insights for a specific parameter (the projected view demonstrates insights on app performance).

Each tab of the dashboard monitors the associated health and performance metrics of individual elements that make up the application's infrastructure, including servers, network devices, application servers, and database cluster. Operational intelligence from the dashboard can be leveraged to effectively optimize application performance by identifying patterns, anomalies, and correlations in infrastructure data.

Extending such focused analyses to other core business applications gives IT teams end-to-end and detailed insights on the overall IT infrastructure. This consolidated dashboard also eliminates redundant alarms and notifications, decluttering IT operations

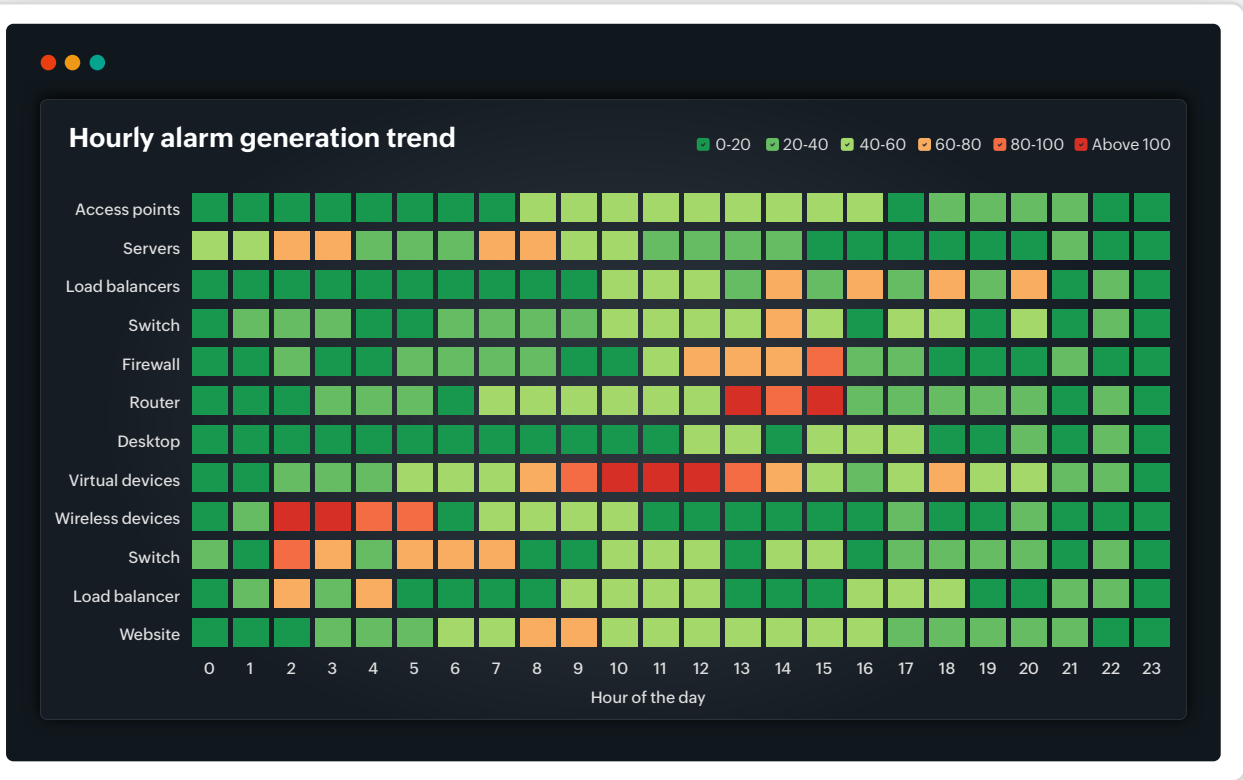
02

Proactively triage recurring alerts

Once IT teams gain detailed visibility into and an understanding of their IT infrastructure, they can use this insight to identify and manage recurring alerts stemming from common IT events.

Events like downtime and scheduled maintenance can create major disruptions in an organization's day-to-day IT operations by introducing a deluge of alerts. Unless IT teams find a way to proactively identify and mitigate these alerts, they will be stuck in a perpetual firefighting mode, further hampering other strategic IT growth initiatives that require their attention. Advanced analytics can be a beacon of hope for IT teams in managing recurring alert noise by establishing patterns in alarm generation that enable proactive triaging of potential alerts and incidents.

The below visualization tracks the alarms generated across device categories in an organization's IT infrastructure during a typical day.



This analysis clearly pinpoints peak alert periods throughout the day. The visualization reveals an uncharacteristic spike in alarms generated by virtual devices. A similar careful evaluation of the alarm patterns enables IT teams to better prepare for and prevent future spikes by proactively taking preventive measures such as configuring application redundancy, re-allocating technicians to handle imminent alert onslaughts, and equipping teams with advanced tools—like automated alert queue management, and smart alert routing—required to address alerts quickly. Taking these steps ahead of time helps prevent a sudden influx of alarms from disrupting NOC teams.

While analyzing alarm patterns and implementing fail-safe procedures can promote smooth IT operations, unless the root cause of these alarms is identified, it is difficult to determine whether alarm spikes are a temporary or recurring phenomenon.

If peaks are caused by one-time activities, establishing elaborate fail-safes would be a waste of time, effort, and resources. However, if regular activities trigger peaks, IT teams should investigate the reason behind the alarms and implement preventative measures to minimize them during those activities. This requires further analysis of the peak alarm periods and a log of critical concurrent activities to unlock insights to implement targeted, effective solutions.

Daily log of organizational events

S.no	Timeline	Event	Device impacted
1.	12:00 am to 3:00 am	Apache server maintenance	Application Server
2.	7:30 am to 10:00 am	DB server migration	DB Server, Corporate website
3.	8:00 am to 9:00 am	Windows 8 update	Corporate website
4.	10:00 am to 12:00 pm	DELL Poweredge R740 server maintenance and system update	VM
5.	2:00 pm to 2:30 pm	Firewall installation	VM, Corporate website, Payroll app, DC
6.	5:00 pm to 6:30 pm	Asset scanning	Corporate website
7.	6:00 pm to 6:30 pm	VMWare ESX/ESXi	VM, Payroll Apps
8.	8:00 pm to 9:00 pm	Tomcat server maintenance	ERP
9.	9:30 pm to 10:00 pm	URL Monitor	Web server

From this analysis, it can be observed that between 10am and 12am, there was a scheduled system update and maintenance on the host server where the virtual machine (VM) resides. During such maintenance activities, the host server may require a restart or undergo configuration changes, which can temporarily impact the VM's availability and provided services, triggering service down alerts.


Often a time-consuming activity, similar updates can often last a few hours or even several days. In such scenarios, NOC teams need to implement appropriate actions such as employing a load balancer to distribute incoming traffic across multiple VM instances, or implementing automated failover mechanisms to redirect traffic to healthy instances upon unavailability. This ensures continuous service availability and subsequently curbs alarm noise.

03

The above analysis exemplifies the value of performing a thorough investigation into anomalies in alarm patterns and fluctuations in alarm volume. The insights from this analysis—combined with the operational knowledge of events causing these alarms—will empower IT teams to control recurring alerts and preempt failures that could disrupt daily business operations.


Mitigating alert noise from faulty devices

To further reduce alert noise and increase alert quality, NOC teams need to isolate devices or nodes that clutter the alert landscape by generating high alert volumes, until the underlying issues are addressed.



The visualization illustrates the trend of alerts across different components of an organization's IT infrastructure, providing insights into where the majority of alerts originate. IT managers can observe that the Cisco router generated an alarming number of alerts, contributing to alert clutter. If left unattended, these alerts pile up and disrupt IT operations, snowballing into a multitude of issues that could result in a total shutdown of business operations.

Isolating such faulty devices and subjecting them to an in-depth investigative analysis will reveal their originating incident.



04


This analysis isolates and categorizes every alert generated by the Cisco router, revealing that service down notifications make up a substantial portion. NOC teams can then run comprehensive health checks on the impacted router, then address the outage as a priority. It is good practice to also verify whether other routers in your network are facing similar outages and implement the identified fix across the organization. In a similar manner, IT managers can identify and isolate other repetitive alerts to effectively reduce alert fatigue.

Revisiting outdated thresholds that undermine alert quality

Aside from faulty devices and IT events, outdated thresholds are a major harbinger of alert chaos in the ITOps landscape, and any endeavor to declutter IT operations would be incomplete unless the obsolete threshold problem is resolved.


Pre-configured thresholds play a decisive role in IT operations as the majority of performance alerts are triggered based on threshold breaches. In today's dynamic IT landscape, as the business conditions evolve, performance thresholds need ongoing updates to reflect changing environments and curb alarm noise from obsolete parameters.

The below analysis of application response time and latency clearly showcases the hazard of operating under outdated threshold parameters.



In pre-pandemic scenarios, wherein most users worked from office premises, most organizations' infrastructure was designed to support on-premises workloads. Therefore, the latency and response time thresholds might have been configured by taking into account the on-premises setup (in which the computing and storage are physically closer to the operating environment). But post-pandemic, as organizations moved to a hybrid work mode, IT infrastructure has evolved to support an equal volume of users connecting remotely. Therefore, the performance thresholds configured in the pre-hybrid mode would not hold true, resulting in unnecessary alerts and contributing to alert noise.

Failure to revisit and update pre-configured thresholds can lead to alert fatigue and inefficiencies in incident response. By leveraging AI-driven analytics to track alert trends over time, organizations can gain insights into the effectiveness of existing thresholds and implement corrections as needed.




This intuitive analysis isolates the alert volume resulting from older thresholds, emphasizing the need to rework threshold configurations across the organization. NOC teams must periodically analyze threshold relevance against organizational practices, and rework them accordingly. This helps organizations mitigate false alarms, improve alert quality, enhance system reliability, and streamline operational efficiency.

Embracing MTTR and MTBF for continuous operational improvement


IT teams are often stuck in a reactive cycle, constantly putting out fires and resolving a barrage of alerts that overwhelm IT operations. This flurry of break-fix events often overshadows key IT operational drivers: MTTR (mean time to repair) and MTBF (mean time between failures). NOC teams often dismiss these metrics as vanity measures since they have little direct impact on resolving daily outages or providing insights into why alerts arise and escalate.

While these metrics do not directly impact alert noise, they serve as a baseline for numerous critical indicators such as service or application availability, uptime, reliability, and technician efficiency—all of which are essential for ensuring minimal alert clutter and efficient IT operations. As a result, focusing on the core metrics like MTTR and MTBF will enable NOC teams to centralize various essential IT infrastructure and operational KPIs as well as alerts, eliminating the onerous chore of simultaneously and separately tracking them. Additionally, monitoring the mean time between application failure accelerates the identification of faulty infrastructure, thereby playing a key role in mitigating alert noise.



The visualization above correlates two IT operations performance metrics: service availability and MTBF. As service availability decreases, alarm count increases drastically. The analysis clearly indicates that a longer MTBF can significantly improve application or device service availability, leading to enhanced end user experience.

Similar to MTBF, MTTR can also have a critical impact on IT operations, as it is a good indicator of how efficient NOC teams are at resolving infrastructure issues and preventing operational disruptions.



As evidenced by the above analysis, tracking the historical MTTR data helps NOC teams generate an accurate forecast of resolution time and corresponding repair cost for upcoming months. Its benefit is two fold: It helps NOC teams implement strategies to improve MTTR, and ensures IT managers set aside a sufficient buffer for repairs in future budget plans. Furthermore, the analysis demonstrates that lowering MTTR significantly reduces monthly operational costs, validating the impact of MTTR on key operational metrics like return on investment and cost efficiency.

With a proactive focus on IT operations strategies—like preventive maintenance, problematic device identification, and alarm pattern detection—IT teams can reduce failures and improve MTBF and MTTR. Doing so further enables them to streamline processes, enhance operational efficiency, and achieve optimal end-user satisfaction.

Conclusion

The practical strategies listed in this e-book can assist organizations in navigating the maze of alert noise and streamlining their IT operations with clarity and focus. With advanced analytics as their guide, IT teams can adopt a proactive strategy that fosters continuous improvement, unlocks untapped potential, and drives transformative change in the future of IT operations.

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 help desk 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.

