


5 ways to create a **chaos-free NOC**





5 ways to create a **chaos-free NOC**

The Network Operations Center, or NOC, of any organization is the cornerstone that connects it people, processes, devices, tools, and technology. All of these generate a staggering amount of data—and alarms. Walk into your operations room, and you're bound to find too many monitors filled with alarms and notifications from applications, networks, and servers.

That's not all! In most organizations, incident management service-level agreements (SLAs) are regimented with strict guidelines to get systems and applications back online quickly, leaving NOC personnel with barely time to fix the issues on hand. The result? Chaotic operations centers often without enough time to follow scripted protocols to avoid downtime, while vast amounts of NOC data, reflecting helpful insights, remains untouched in its databases.

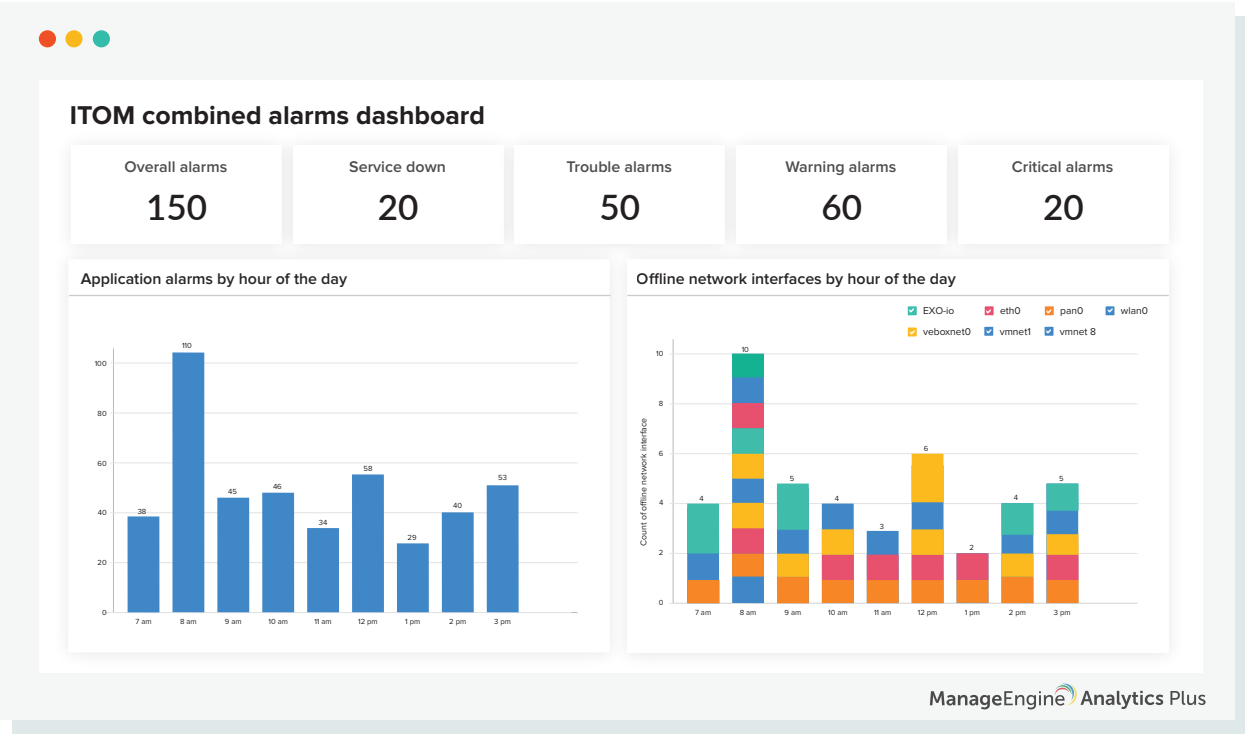
Despite these seemingly insurmountable challenges, NOC centers can break free of this chaos, and run efficiently by following a few simple steps:

Establish a unified view of all NOC applications.

NOCs primarily handle fault logging and performance management of resources such as networks, servers, applications, virtual and cloud resources, and more. Each of these resources deploy independent tracking and monitoring applications. For example, to monitor apps, you may be using ManageEngine Applications Manager. For virtual hosts, you may be using ManageEngine OpManager, or another network monitoring software.

In effect, more technology means more visibility into problems, issues, or bottlenecks in your resources. However, too many monitoring solutions create complexities because too many dashboards make it impossible to visualize end-to-end processes. Also, conflicting status updates often lead to a blame game, making it difficult to efficiently trace the root cause of issues.

The need of the hour is to establish a single source of truth that can integrate data from several resources and applications into a unified console or dashboard. This will help you achieve full transparency over the entire NOC.



The dashboard above, built using data from Applications Manager and OpManager Plus, enables you to visualize the impact of network interface-related issues over application alarms. You can see spikes in applications alarms whenever applications go offline due to the error-prone network interface devices.

Similarly, you can create unified dashboards to visualize several interrelated NOC processes, such as the impact of bandwidth congestion on application responsiveness, the impact of server availability on application accessibility, and more.

Establishing a unified view of all NOC operations empowers you to achieve full transparency over processes, technology, and people, enabling you to make better, faster business decisions to streamline processes, eliminate bottlenecks, and ensure smoother IT operations.

Streamline incident management processes with automation.

Automating redundant operational tasks is the most effective solution to reduce the chaos brimming within the NOC. So, where do you start? Build a laundry list of typical tasks carried out in your NOC, and put them into an Eisenhower's matrix based on their importance and urgency. Start from tasks that are urgent and important; and then move through the grid vertically. This type of grid-based categorization forces NOC personnel to view tasks objectively, and question the time spent in resolving them, while also providing greater clarity of operational processes.

To illustrate, here's a sample matrix that shows the split of NOC tasks by urgency and importance.

	Urgent	Not-urgent
Important	Alert resolution, service restart, free-up disk space on sever, Managing VPN issues	Strategic tasks such as resource or capacity planning, deployment of new resources, scheduled maintenance of infrastructure
Not-Important	Handling false-positive alerts from networks and applications, server patching	Tracking resources, provisioning new equipment, creating dashboards for senior management

Based on the matrix, most of the important-urgent tasks that can be automated are:


- **Alert resolution** - Draft workflows to automatically assign alerts to technicians based on priority. Use automated workflows to deal with routine alerts, or L1 alerts.
- **Service restart** - Configure automatic service restart for Windows, Linux services, such as backup services, antivirus, spooler, Internet Information Services (IIS), and Apache services.
- **VPN tasks** - Automate VPN management tasks such as VPN account unlocks, issuing, extending or replacing authentication tokens, or resetting VPN.

Using automation, many redundant tasks can be streamlined, making way for better and faster incident management processes, and liberating NOC personnel from mundane tasks, so they can focus on other important tasks.

Filter out the most meaningful alarms.


The longer you work in a NOC, the less those flashing yellow, orange, and red alerts capture your immediate attention due to the overwhelming tide of alarms you see from monitoring applications. Too much noise (alarms) drowns the critical ones amidst an ocean of alarms, making it impossible for NOC personnel to identify critical issues that could create a tsunami of costly and show-stopping outages.

Take, for instance, the SQL database server. The CPU usage might hit 100 percent several times a day, but the NOC team doesn't need to be alerted every single time. You only need to receive notifications when the CPU usage hits 100 percent when it regularly hits only 20 percent at a certain time of the day.



Anomaly detection in application or network behavior should be built on actual anomalies, and not predictable behavior. This will help isolate out-of-the-ordinary behavior, and eliminate redundant and false alarms from flooding your inboxes.

The solution is to eliminate static thresholds, and use dynamic thresholds that can detect anomalies and show you only those alerts that matter.




The report above shows you dynamic thresholds for CPU usage created during operational hours in a given day.

Partner with AI to facilitate continuous service improvement or faster incident resolution.

IT infrastructure is highly complex and interwoven, it involves several resources so that even a single isolated incident will trigger a tsunami of alerts from several monitoring applications. Owing to the sheer volume of this data flow, it's futile to rely only on manual reporting and analysis to parse through incident data. Partnering with artificial intelligence (AI), machine learning (ML), and natural language processing (NLP) methods will help you sift through the data expertly while enabling you to identify patterns and recurring behaviors hidden deep within the data.

A classic example of ML-based pattern identification involves the clustering of issues and applying event correlation logic to identify the impact on IT issues of seasonal changes, or maintenance schedules. ML algorithms can quickly discover patterns and seasonal changes in your data, such as seasonal maintenance schedules, server migration schedules, and new technology implementation.


ManageEngine Analytics Plus takes ML and AI-piloting a step further with Zia, it's built-in AI-assistant. You can just ask or type in a question, such as "Can you show me the combined alarm trends for the past two years?", and easily gain insights from Zia.



The report above shows consistent spikes in alarms every six months. Maintenance activities on servers or networks that are carried out yearly or bi-annually typically cause some disruptions in network or application availability, leading to an increase in issues during this period.

Now, imagine NOC technicians gaining this insight using AI. It could help them better prepare for an onslaught of issues during seasonal maintenance.

The same report in the hands of operations managers can arm them with enough ammo to forecast capacity and resource requirements. Using the report alarm, NOC managers can perform casualty analysis, and understand the range of infrastructure devices that are likely to go down during scheduled maintenance. With this information, NOC managers can develop alternate plans, such as sharing server load with other servers located elsewhere, etc. NOC managers can also use a version of this report to predict future alarm trends and proactively make suggestions, such as adding temporary resources to combat incoming incident during scheduled maintenance.




The report above shows the historical trend of alarms as well as a forecast of alarms for the next year.

Boost productivity by training technicians to become subject matter experts.

While it's impossible for any organization to achieve zero issues, it's important that organizations of all sizes aspire to achieve this goal. They can start by documenting the incident resolution process.

NOCs frequently operate in high-stress environments where the sense of urgency to fix issues quickly and by using all means possible prevails. However, a poorly thought out fix can quickly cascade into bigger problems. The best action for NOC personnel is to document all resolution processes. Create and store all necessary information, such as workflows, procedures, and knowledge base articles for anytime, anywhere access so that crisis scenarios become a lot more manageable.

If needed, break down incident resolution processes into several steps. Thoroughly analyze the root cause of issues, and document actions that will resolve or reduce the effect of issues, or prevent it from occurring in the future.



The report above shows that the most frequently occurring issues are a good place to look if you're just beginning to document IT issues. If you're already documenting these issues, explore other areas that can help establish your stronghold and build resilience to IT issues.

Conclusion

Despite the inherent chaotic nature of NOCs, adopting certain best practices—establishing a unified view of NOC, automating tasks wherever applicable, deploying analytics to identify critical alarms, partnering with AI, ML, NLP, and documenting incident resolution processes—it's possible to break free of this chaos. This will help establish a seamless, stress-free operations center that not only focuses on delivering exceptional customer services, but also strives to continuously improve and level-up their processes.



About

ManageEngine Analytics Plus

Analytics Plus is a self-service, IT analytics solution that enables you to visualize your system data in the form of colorful charts, reports, and dashboards. It offers out-of-the-box integrations with ManageEngine Applications Manager, OpManager, and other ManageEngine tools to give you an in-depth look at your IT infrastructure. It features a built-in AI-assistant, Zia, that eliminates the need for query-based reporting, and enables everyone in the NOC team to gain instant insights by asking or typing in questions. Analytics Plus also enables you to forecast capacity and resource requirements accurately to help optimize operations and improve service delivery.

180K
customers
across the world

18+
years of IT
management experience

90+
products
and free tools

190+
countries
served