Netreo

WHITEPAPER

Increasing Enterprise IT Infrastructure & Operations Efficiency

JANUARY 2019 - AN INDUSTRY PERSPECTIVE

Contents

Introduction

Maintaining efficiency in IT operations and infrastructure is a challenge for organizations of all sizes, and becomes an even greater challenge as organizations scale. In this whitepaper we share lessons learned and best practices gathered over nearly 20 years of working with IT leaders from all industries. As you read further you'll discover key tactics for Increasing Enterprise IT Infrastructure & Operations Efficiency, spanning topics from chaos in the network operations center to breaking silos to improving crisis communications. Enjoy!

Addressing Chaos in the Network Operations Center

The operational success of any large-scale enterprise IT environment begins and ends with the Network Operations Center. In spite of this fact the NOC is often littered with inefficiencies and chaos.

A NOC is, by its nature, a chaotic place. It's an interwoven and interdependent orchestra of people, process, and tools: Flat-screen monitors (sometimes 15 or more deep), real-time flash, actionable notifications and critical alerts. All of this data is presented in various cryptic ways from a passel of non-integrated tools. Already overworked NOC personnel struggle to keep their eyes trained on this dizzying array of content in the hope of spotting outages and anomalies as they occur. At the same time they must also manage the flood of informational and related data streams into their inboxes and SMS feeds. Rarely is there any intelligent filtering or aggregation on this inbound tsunami of data. Industry studies show that a high percentage of this deluge is erroneous, extraneous, and unnecessary.

In most cases scripted protocols for incident management must be strictly followed as problematic issues are distilled out of the fray and isolated for remediation. Service Level Agreements (SLA) depend upon such a regimented approach. In the midst of the aforementioned chaos user, customer, and partner communities rely on the efficient and secure operation of the overall environment. SLA attainment necessitates remediation be transparently completed as close to instantaneous as possible.

In spite of these seeming insurmountable challenges, success is attainable. When reviewing the efficiency of your NOC consider the following points:

- A fully-integrated or consolidated tool-suite, optimally designed by a single manufacturer or at minimum integrated via APIs to the other 'cogs' in the wheel. This integration will minimize the optical touch-points, eliminate finger-pointing between products, and ensure a consistent interface into collected data.

- Evaluate technologies that have the ability to intelligently filter, sort, and aggregate large amounts of data. The goal is to separate out elements you need to draw attention to. Examples are technologies such as root-cause analysis, event correlation, and automatic anomaly detection. However, it's imperative these technologies avoid the trap of requiring more time to implement than they save during the troubleshooting process. These technologies must be fully-automatic.

- Look toward dynamic products that are able to roll-up enterprise-wide visibility into a finite visual presentation. You can always drill-down into the weeds later (and when you choose to do so it should also be intelligent).

- Stretch your NOC personnel by making them subject-matter experts on the new set of integrated tools. Radically improve their individual efficiency statistics by distilling the content they receive. Remove the noise.

- Where practicable electronically embed process information across the fabric of the NOC. Naturally, hard-copy backup documentation is OK, but to the extent that you are able to intelligently and automatically embed required information such as workflow, procedures, and NOC wikis for anytime, anywhere access, crisis scenarios become a lot more manageable.

A chaotic NOC environment doesn't have to be status quo. There is a better way. Based on our experience, we've seen that with the right combination of tools, process and personnel they can be well-oiled machines.

Eliminating Silos and Blind Spots in IT Management

Today's Infrastructure and Operations (I&O) leaders have a challenging path ahead of them. Service Level Agreements (SLAs) which have been committed to the user community must be not only met, but measured, validated and communicated across the enterprise. Geographic expansion, M&A, and other key business initiatives must be factored into I&O decisions. Vendor and product selection, management, migration and right-sizing create a ceaseless treadmill of activity. Compounding the typical workload, leaders also need navigate the nonstop flow of cutting-edge trends that must be considered - how will Artificial Intelligence affect my sector? IoT? Virtual Reality? Digital Twins? How about Cognitive Computing? Amidst this entanglement of competing agendas and pressures, I&O stewards must find ways to elevate themselves and their operation via streamlining and optimization efforts.

The Truth About Silos

Consider silos - those computing environments that are directed towards a very specific function or environment and work right alongside other silos which perform that are performing similar tasks. Silos tend to evolve over time as a form of entropy. How? Think about that business unit that got folded-in or acquired and used their own suite of tools. Or think about smaller departmental groups who insisted on more specific tools for their non-standard needs. Another common scenario: The team that built homegrown open-source suites and refuses to let them go.

Sometimes silos do make sense - there might be very specific individual business unit requirements that drive the need. However, that case is more the exception than the rule. From non-shared data, multiple points of administration, and dissimilar operation in most cases silo-specific tools introduce unnecessary redundancy and waste.

As an alternative, consider tool-suites that are integrated. In exchange for potentially losing a minuscule amount of highly-specific functionality, your team saves time not having to redundantly manage data and processes across disparate platforms. Integrated, non-silo product suites make a huge impact in terms of reducing overhead and introducing an overall streamlining effect to your I&O team.

Blind Spots

Blind spots are areas of the I&O environment where staff has no visibility. This lack of insight will inevitably lead to last-minute fire drills and unnecessary surprises. We've all witnessed scenarios where IT production levels seemed to be running at optimal levels. However, in reality latent problems had been forming for months. The situation suddenly degraded to the point of imminent failure and your team has no advanced warning. The method of discovery? Angry users. Not a great situation to facilitate your optimization efforts.

The solution is adopting a proactive approach to service delivery. Assemble the right stakeholders, identify possible blind spots, and build automated tools and processes to simulate and evaluate performance. Run persistent synthetic checks on all aspects of I&O to isolate bottlenecks and congestion. Set up a protocol to routinely run spot-checks to help ensure that you're ready and in a proactive mode.

The pressures of running an I&O organization needn't be daunting. Finding areas to optimize and simplify you can reclaim time, effort, and resources and make the task more manageable.

Maintaining Visibility in Hybrid, Transitional, and Private Cloud Environments

It's no secret that most IT services are either hosted in the cloud, or they are headed that way. In fact, the proliferation of cloud-based services in today's corporate infrastructure environments will be looked back upon as one of the most compelling and important trends the IT marketplace has ever witnessed. While the massive advantages gained in terms of scaling, agility, and elasticity are impossible to ignore, cloud computing can (and typically does) introduce some visibility and control challenges. Whatever the flavor (IaaS, PaaS, SaaS, STaaS) or architecture (public, private, hybrid, transitional), Infrastructure and Operations (I&O) leaders are finding out that most cloud providers don't provide (or even allow) adequate visibility. Going forward, how will I&O leaders commit, measure and attain given user-experience SLAs as they have done in the past? How will they accurately assess utilization levels to determine over-subscription or under-subscription of capacity? How will they spot potential problems and bottlenecks as they are forming, versus after they have had a chance to fester and ultimately affect production traffic?

Dealing with the Opacity of Most Infrastructure as a Service (IaaS) Offerings

Typically, IaaS vendors tend to cloak activity 'behind the curtain', and assumes that you will be satisfied with vague, aggregate roll-up data of web workloads. You may have to take matters into your own hands. Therefore, consider tools that allow access directly into your core cloud providers via REST-APIs, Webhooks, or other means of integration. However, even the best tools won't be up to the task if certain requirements aren't met. First, you verify that your existing management and cloud vendors support these methods of data collection. Next, start persistent polling of sample transactions and workloads. Lastly, tie it all together. You'll need to model these polled samples as they occur in your actual production traffic, incorporating each step, call, and hook to give a complete picture. This picture will show you exactly what components need to be monitored to give the visibility you need.

Filtering Out the Noise

Too much data is worse than not enough data. When presented with a tsunami of inbound notifications I&O personnel will go into 'alert fatigue' mode. They'll begin to ingore important problems as well as the false alarms and redundant alerts. Therefore, it's critical to keep the signal-to-noise alert ratio very high and avoid this syndrome. It is crucial to note that many isolated events are benign when they happen on their own, but when they occur in concert with other specific events, they can spell trouble. You don't want to flood your team with these isolated and inconsequential events, so start to consider which need to be monitored as a collective grouping. Map out all the application dependencies, and strategically tie together related events via conditional, sequential or other patterned trigger statements. The net effect is to only receive filtered, intelligent alerts that mirror what your actual stakeholders are experiencing when transacting with your cloud-based applications.

Single Screen for All Elements

Keep the visibility that you are able to achieve into cloud-based resources front-and-center, stacked alongside other I&O metrics. If personnel must sift through various screens to get to cloud-specific content, then those metrics won't get the attention it needs. Today's cloud environments can be tricky to keep control of but putting the right visibility strategy in place can set you on the right path.

Next-Gen Crisis Communication for IT

The importance of clear communications during a crisis is something emergency planners and first-responders live with every day. They spend countless hours planning for contingencies and working out the worst-case scenarios. They make sure that they don't lose the ability to coordinate. In contrast, Information Technology professionals rarely engage in this type of activity, although it's clear from our experience they should.

Crisis Policy Development

When we're asked to develop crisis policies for our clients in isn't uncommon to find they lack any plan whatsoever. Further, it isn't uncommon to find during a mass outage that's impacting a company's most critical operations, the IT staff is improvising, trying to make the best of a bad situation. Simultaneously, a large number of different subject matter experts are brought in, each looking at different alerts, different toolsets, and coming to (understandably) different conclusions as to the root cause of the problem. The takeaway is that resulting recovery inevitably takes much more time and effort than anyone anticipated.

Any organization that relies on IT for daily operations must spend some time developing a crisis communications and alerting policy. What elements comprise a good alerting policy? A good starting point is to define "What" variables should be monitored, "Who" should be alerted in case of a problems, "How" should personnel be notified, and "When" it is appropriate to notify personnel. be used to alert those persons. Of course, the biggest question to ask when setting up an alerting policy is "Why" are we sending out alerts and deploying an NMS system at all? The answers to the "Why" question will help frame your answers to the others. This document then provides a baseline to determine where capability gaps are present and can be used as a blueprint for how to improve outage response and reduce mean time to repair.

Why Not Monitor Everything?

Often the question of "why not monitor everything" is posited. In our experience, monitoring everything only results in causing excessive noise and makes real problems harder to spot. Instead, focus monitoring on the mission-critical services, applications, and infrastructure. Over-monitoring masks real issues and causes operators to 'tune out' alerts or even configure filters to block them. While collecting information on many aspects of your systems and networks can be beneficial from a forensic or troubleshooting standpoint, notifications should be strictly limited to only actionable items - things that will cause an operator or engineer to take corrective action. Anything else should be limited to on-demand retrieval or scheduled reports.

Once you have identified what to monitor, centralizing notifications and workflow through a ticketing system or alert manager is the best course of action. There are two primary benefits here: First, it will reduce the number of disparate points of administration required as personnel changes occur. Second, there is a singular source for IT service-level metrics that are crucial tracking the efficiency of you I&O team. Whatever the selected system it must be simple to integrate both upstream and downstream with other IT components, offer flexible notification options, and intelligent filtering to group related notifications or alarms together to avoid creating a flood of alerts during a mass outage. It's also very helpful if it automates the process of keeping everyone updated as people take ownership of issues or begin working the problem.

Internally-hosted email isn't sufficient as a method of communication during a crisis. Hosted email providers are generally better as they are less dependent on your organizational infrastructure. However, hosted providers do require Internet access to be used successfully. Therefore, thought must be given to creating redundancy for the communications channels used during a major incident. Redundant Internet connectivity or providing true off-the-wire backup (such as 4G access points) can go a long way to creating more options during a crisis. Make sure your monitoring and automation systems can easily integrate with whatever platform you choose.

Group communications with ad-hoc organization are a powerful way to improve your crisis response, provided your team can access them when the network is dark. There are a number of these services available now, ranging from extremely simplistic (e.g., Jabber) to much more advanced and complex options that offer voice or video and can integrate with your organization's unified communications system.

No matter which method is chosen many requirements need to be met: They must be clearly understood, well-documented, communicated internally, thoroughly monitored, and redundantly connected. The last, and most important requirement, is that the system get used in crisis situations. If your organization does not regularly run crisis drills, then consider doing adding such a drill to the operational plan. Practice makes perfect and it's always preferable to get surprised in a controlled situation rather than during the heat of an actual outage.

Managing Application Performance In the Hybrid Cloud

Hybrid cloud adoption more than tripled during 2017. As enterprises move more of their resources to cloud infrastructure, the need for physical infrastructure monitoring begins to diminish. However, there are a number of good reasons from a security and architecture perspective for keeping some resources internal, especially during a transition period. Even if the goal is to move all IT resources to the cloud, that change doesn't happen overnight and the challenges of managing this hybrid environment would still exist. How else can user service levels and operational efficiencies be maintained sans this visibility even if control of the physical infrastructure is outsourced?

Many cloud providers offer their own monitoring and reporting services. Some are rudimentary while others are more advanced. Unfortunately, this scenario creates multiple points of administration, potential data-sharing problems, and necessitates training users in multiple tools. If this narrative sounds familiar that's because it is. However, instead of having an "in-house" silo-specific problem it's been shifted to the cloud services provider. It doesn't matter where the silo is it's still redundant and wasteful.

It's a well-known tenet of enterprise management that the fewer places you need to look for information, the faster encountered problems are identified and resolved. Translation: IT teams are more effective when a consolidated approach is employed. Visualizing and managing local and cloud hosted platforms in the same toolset should be the goal. However, cloud providers don't always offer explicit availability guarantees for application-level functionality. Therefore, the burden falls on the I&O team to choose applications that can handle both environments well.

Hosted virtual machines and storage do offer operational redundancy. However, crafting truly resilient application architectures within the cloud can be tricky, and can still carry the risk of degraded user performance once complete.

Managing Application Performance In the Hybrid Cloud, continued

The key to identifying and tackling this challenge is integrating application performance monitoring combined with a resource-focused monitoring tool. Ideally, your application monitoring platform supports an integrated global view of all your enterprise resources regardless of their location. The ability to monitor the performance of applications from various locations around the Internet, as well as from where it is hosted is paramount. Only then can you ensure the metrics shown give you an infrastructure as well as a "user-experience" perspective into your application. The best of these tools can integrate all of that functionality and visibility directly into a single console.

Many cloud providers like Amazon AWS have advanced tools for controlling the security and traffic flows within their cloud, but these highly granular controls are complex to configure and create the potential for misconfiguration, which can cause degradation and service interruption.

All organizations must have a plan to ensure reliable application delivery with rapid notification of performance issues and failure. An effective application monitoring platform must provide detailed validation of content, including the ability to look for XPath or regular expressions to fully validate content and guarantee that applications are performing as required.

Additional key functionality is the ability to integrate with a cloud provider's APIs to automatically detect new services, consolidate events, and alerts and logs. When this capability is present the time and effort saved monitoring, administering, and on-boarding new resources is incalculable. In contrast, sans this capability, a monitoring platform that fails to adjust dynamically has more dire consequences: unreported/undetected user-facing errors and rapidly degrading user confidence in services - and much bigger problem to rectify.

While the move to a hybrid cloud architecture is inevitable for most organizations, replete with its many advantages, a consolidated monitoring approach will help to mitigate the inherent performance risks that can occur if care is not taken from the outset.

Find the Low-Hanging Fruit: Making Your IT Workforce More Efficient

Throughout this piece, we've attempted to illustrate ways in which you can avoid common pitfalls of IT management to improve your organization's efficiency, and to avoid costly and time-consuming mistakes. To sum it all up, here are several conclusions and specific best practices we've seen in our customer base that have proven to be both effective and efficient.

1. Make IT Management as Simple as Possible, but No Simpler

It's long been a goal of most organizations to try and simplify IT operations, but as time passes the organization, the systems it must support, and the methods used to do that increase in complexity anyway. This change isn't inevitable. The problem can be mitigated by choosing tools that carry a minimal burden to implement, manage, and gives you visibility into the widest range of systems and applications. Choose the tool that allows focus on the information you really need, rather than drowning the team in a sea of raw data.

Look for tools that have intuitive and clean user interfaces with minimal clutter help to manage the 'attention budget' of your over-taxed IT team. IT is already a noisy environment with many other things fighting for your team's attention. It makes sense that keeping things simple helps them focus on what's important. However, it's imperative deeper information is quickly available. In this manner your team members' focus on the problem at hand, while at the same time allow them to drill into historical data for capacity planning, correlation, or forensics, where more data is needed. Similarly, also validate your tools keep enough history to show meaningful change over time without becoming so slow as to be unusable or becoming a huge cost in resources and storage.

2. Consolidate Your Tools

Whenever possible, reduce the number of places your team has to go to get information. Although, a true 'single pane of glass' is not really possible due to the number of disparate systems to monitor, the closer you can get to this goal, the faster you'll be able to restore service to your users.

In order to cover everything, a single tool may not be enough, though it should be easy to get 90% of everything integrated into a primary console. Supplement it as needed with 'microscope' tools that get very deep into specific environments such as sniffers, SQL, and mainframes, which extraordinarily detailed visibility is required. Move everything else to a system that gives you the quickest possible access to the data you need, while streamlining operations and response time. This approach is key to maintaining a fast, DevOps-oriented continuous deployment and integration.

3. Stop the Fire Alarms

Identify the specific reasons for each alert you get from your systems. If the problem is not something you will take action on immediately, consider replacing that real-time alert with an intelligent scheduled report. Instead of relying on the system to wake you in the middle of the night to deal with a hard drive that's 95% full, isn't it preferable to get a report every day containing the servers with the least free space? Isn't it preferable to fix the problem proactively, instead of waiting for the fire alarm to go off? An ounce of prevention here is worth more than a pound of cure. Your staff can get a good night's sleep.

4. Spend Less Time Configuring Your Management Platform

To minimize the time required to set up and maintain your systems, look for tools that allow you to manage by exceptions and automatically detect anomalies and unusual behavior. Especially focus on eliminating tools where each monitored service or application must be individually configured.

Consider the maintenance load of the tool itself. Does it require frequent updates that have to be applied to multiple systems in a precise order to be effective? Or can you simply click a single button to update everything? Does it require secondary application support, such as databases that must be maintained and patched, or an OS that requires weekly reboots to stay current? Consider ways in which you can move to more stable platforms and reduce the time investment of keeping the system functional.

Following these four best practices will net your IT workforce fewer distractions and more time to focus on what's important, in easy stages.

Netreo

Product Overview

Observe. Analyze. Act.

Building the Single Source of Truth

Observe.

Full-stack visibility into your IT infrastructure, systems, applications, and user experience — in a single customized dashboard.

Analyze.

Netreo's AIOps engine delivers simple answers from mountains of data leveraging over 20 years of historical baselines and trends.

Act.

Netreo provides real-time dashboards, dynamic automations, and extensible ITSM integration so you can focus on taking intelligent actions.

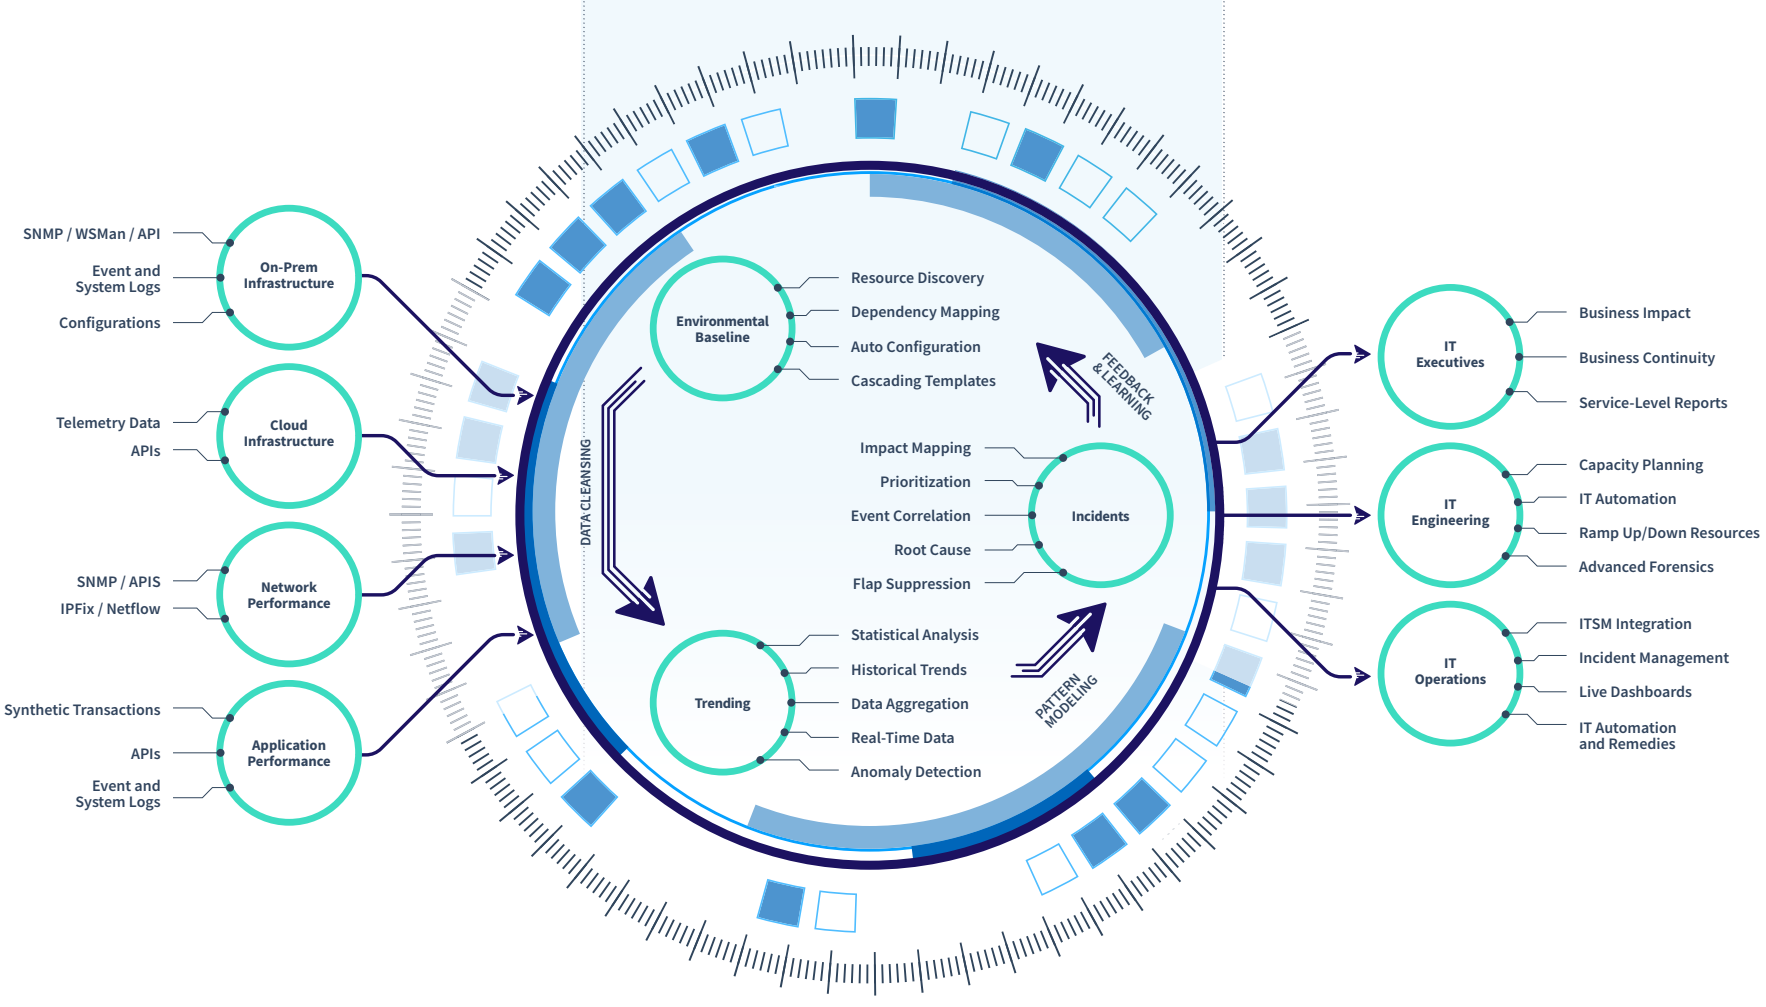

Visit netreo.com to learn more

A Single Source of Truth

Full-Stack Monitoring and AIOps

Observe and automate everything across the enterprise in a single, unified dashboard — from wherever you happen to be.

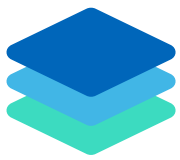

Full-Stack Monitoring

Netreo is a full-stack IT management solution that measures the state, operational status, and business impact of all the components in your technology stack. Full-stack monitoring begins at the infrastructure level and includes servers, application performance, system integrations, and digital user experience. Only by adopting a comprehensive visibility approach can you guarantee that your teams will always have the information they need when they need it.

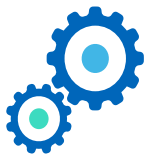

Incident Automation

Your IT Engineers are extraordinarily busy. Netreo can automatically respond to actionable incidents and take corrective action faster than humans can react, or provide 'click to fix' buttons for your operators and engineers to simplify their life. And of course, Netreo can also generate trouble tickets, send notifications via email, SMS, Slack, OpsGenie, Pagerduty, or any of dozens of other notification platforms.

Mobile-Enabled Management

Netreo's mobile app (IOS/Android) allows your teams to manage your environment from anywhere they happen to be, using a patented security architecture that ensures reliable, secure access even from across the country or around the world. Teams now have full access to all of your monitoring data, history, incidents, events, and reports, and can even initiate corrective actions, all from their mobile device.

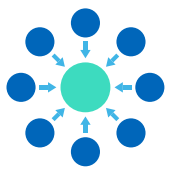

Zero-Touch Onboarding and Life Cycle Administration

Netreo is designed from the ground-up for scalable, easy administering. New devices can be onboarded and monitored automatically without any operator action, so you never miss a problem due to lack of monitoring. Our automatic discovery and classification systems enable your configuration to stay in-sync with your environment by integrating with your cloud providers, SDWAN, virtualization, and CMDB platforms.

Intelligent Incident Detection and Management

No one can work through a storm of constant alerts. Netreo ensures that your teams only see the alerts they need. Using dynamic topological data and machine learning algorithms, Netreo is able to detect unusual behavior to correlate events. Problems can be prevented before they affect users, root causes are found faster, mission-critical services are more reliable, and SLA attainment becomes the rule rather than the exception.

Deep Insights, Instead of Floods of Raw Data

Filtering through huge lists and logs is a poor way to find issues, and as IT increases in complexity, this problem only gets worse. Netreo makes sure your technical personnel and IT leadership have the data they need to make informed decisions, including predictive analysis, three years of on-board historical data, and instant reporting and dashboard options that don't take special training or scripting knowledge to use.