



 Investigating  
pocket...

Whitepaper  
**Optimizing Application  
Experience, Monitoring and  
Security: A Holistic Look**

With the sudden shift to remote working, applications have become critical to organizations, their employees, and their customers. Applications that used to have 10-20% loads are now experiencing loads of 90% or more, while the demand for better application experience is unprecedented.

More than just managing loads and satisfying user demands, the ability to provide a satisfactory application experience has become critical to business continuity. Today, slow or faulty applications can break businesses. Conversely, a great application experience can propel a business through these troubled times, even helping to increase its bottom line.

IT teams have been pressed to provide and maintain seamless application experience. It is paramount that applications are secure without compromising on experience. And as digital transformation has accelerated throughout the world, new challenges have arisen, such as dealing with complex networks at the edge and on the cloud.

IT teams are now looking towards new tools, technologies and approaches that can improve application experience. And the first key step in this is to be able to monitor the network in a clear and holistic manner, and diagnose issues as soon as they arise.

## Applications: The Key to Business Continuity in the New Normal

An IT operations team needs the right people, the right skillsets, the right processes, and ultimately, the right tools. At the same time, most organizations face budget or other constraints, restricting the ability for IT teams to scale up – whether in terms of increasing the size of the team or hiring the right people. With these constraints, IT teams – especially operations teams – have to work as effectively as possible in order to deal with the challenges that come with providing a seamless application experience.

IT teams often spend much of their time troubleshooting, sifting through a tremendous amount of noise before finding the root causes of issues. Such work is often tedious and routine, involving manually monitoring all aspects of the network and capturing possible pain points. The time and energy could be better spent solving bigger, more complex challenges.

Automation tools can help. By taking over menial tasks such

as routine network monitoring, automation can reduce the noise-to-signal ratio and identify the more serious issues that require human intelligence. They can also provide a holistic view of the network, making IT teams more aware of parts of the network that were previously ignored. The right tools and infrastructure can connect different domains, extract pertinent information, and allow operators to quickly triage issues in application experience. Investing in them can save an organization a significant amount of time,

### Example:

*Email is typically looked at as non-critical infrastructure. However, when networks are disrupted, and email becomes slow or inaccessible, the importance of email suddenly becomes evident. Suddenly, it becomes necessary to quickly address issues that were previously ignored because email was not seen as critical infrastructure.*

money, and effort.

## Getting business on board

IT teams might be well-convinced about the benefits of the latest platforms and technologies, and be eager to invest in the right tools and infrastructure. However, acquiring the resources to do so will require getting business on board, which is not always an easy task.

IT departments should think about the business objectives that the technology in question can help achieve. Having a conversation with business to understand what the rest of the organization needs, and then framing technology in terms of non IT-related objectives, can go a long way towards convincing business on the benefits of investing in technology.

Questions that IT can ask include:

Does the technology

1. *Reduce cost of the IT department, or even the entire organization?*
2. *Reduce the time it takes for a problem to get resolved?*
3. *Help the organization become more agile?*
4. *Increase the ability to get services online as quickly as possible in the event of a breakage?*

*Tip: Start with a prototype*

*Once IT is able to set clear goals for the desired toolset or infrastructure (technical as well as business goals), a good first step towards implementation would be to work on a small and achievable project or a prototype. IT teams should ask themselves:*

*How do we apply this technology to a very small area, to see if it's even viable within our operating environment? How will it meet the objective?*

## Defining a good application experience

Today, a good application needs to be secure, scalable, and available 24/7.

For this, traffic flows need to be efficiently distributed, getting clients from the application to the best possible back-end target. Intelligent decision-making, whether manually or through automation, is critical when it comes to traffic distribution, and security and integration with identity systems. Multiple architectures may be involved, including cloud and on-premise servers. Additionally, microservices container ecosystems need to be mapped, and the same process of distribution, security, and integration applied.

The objective is ultimately to connect the dots. If the upper layers of the application stack can be connected, end-to-end, with the lower layers of the network, then a picture can be drawn of the activity of the entire network. The actual application experience a user goes through can then be thoroughly understood. It also gives new insights into security and mitigation, because new workflows can be triggered to circumvent performance problems, or new rules applied to mitigate security issues.

For an application to be smooth and secure, two domains have to work hand-in-hand: network operations (NetOps) and security operations (SecOps).

NetOps' duty is to monitor the performance of the network, detect anomalies, and find ways to fix them. SecOps, on the other hand, is crucial to ensuring that the network and application are secure, and must be on the lookout for potential threats and breaches.

## Network Performance Monitoring & Diagnostics (NPMD)

Before any issues in the network can be fixed, operators need to monitor the network and diagnose any anomalies that may arise. This is known as Network Performance Monitoring and Diagnostics (NPMD), which plays a key part in application experience because

it helps frame the context in any disruption of the application experience.

If an application is having a performance issue, there are several angles from which to analyze it. One way would be to look at the end-point, and another, the server. However, these pictures are disparate and fragmented. To give a holistic picture, an organization needs to also monitor the communication flows between the various entities in the network.

NPMD is meant to create this holistic picture for the network operator. Traditionally, a network operator would need to stitch and string things together manually across captures and logs to get a clear overall view of the network.

Such tedious manual captures are time-consuming, and most organizations do not regularly receive reports that can give them relevant, customized data pertinent to their specific needs. In a poll of approximately 100 senior IT executives in APAC, fewer than 30% said they were able to get tailored network reports.

### **Are you able to get tailored network traffic reports?**



The outcome of the poll attests to the genuine challenge IT faces in understanding every relevant aspect of the environment. Some organizations do not have the right tools to help with this, and others do not know how

to make the best use of automation tools. On top of that, the environment typically sees many exceptions, such as special projects or legacy systems. For most organizations, getting a segmented capture, for example from Point A to Point B of the network, is possible, but having a clear map of the entire ecosystem can be very challenging.

There are platforms and tools to do this, helping network operators understand issues more quickly, predict them, and plan around them. For example, NPMD can help an operator predict how the network will behave when a traffic surge is expected, such as during a big sale on an ecommerce platform.

- **Amorphous Networks and the Edge**

Compared to a decade ago, network architecture today is very amorphous and fragmented. With the edge and the cloud, the number of servers and connections have multiplied and reside further and further away from the core. Critical entities such as identity systems may be in the cloud ecosystem. Many aspects of the network, such as the devices connecting into the ecosystem, are out of an organization's control, and security needs to go beyond borders and firewalls. Without a fixed perimeter, ensuring a smooth application experience is more challenging than it used to be.

The result is a lot of unpredictability, especially in terms of what a customer or user is experiencing. The clearest example is in nationwide rollouts of applications, for example, a banking or a government application. A user in Florida might be experiencing a network slowdown, which might not affect an operator in California. Different geographies have different edge and service providers, some of whom might have an issue, making application experience beyond an organization's control. And yet, the application is expected to deliver consistently throughout. In such cases, network architects and operators should be mindful of ways to make application experience less volatile.

- **Setting the right policies in place**

To begin managing the amorphous network, the first step is to create the right policies. These should be policies

an organization will be able to apply regardless of where entities are connecting from, and what they're connecting to.

Some questions to ask when setting policies include:

1. *Can consistent rules be put in place?*
2. *How do I enforce things across my entire ecosystem?*
3. *Can I arrange my network elements so that the same policies can be used regardless of where the endpoint and client are?*
4. *Do I have an example of a service or set of services that spans beyond the border?*
5. *What's the outcome I'm looking for?*

Work on a small area to see what works and what doesn't, and then scale that across the ecosystem. When policies are finally put in place, all the monitoring tool sets should be chosen carefully and analyzed, to make sure that the infrastructure can be integrated into the cloud.

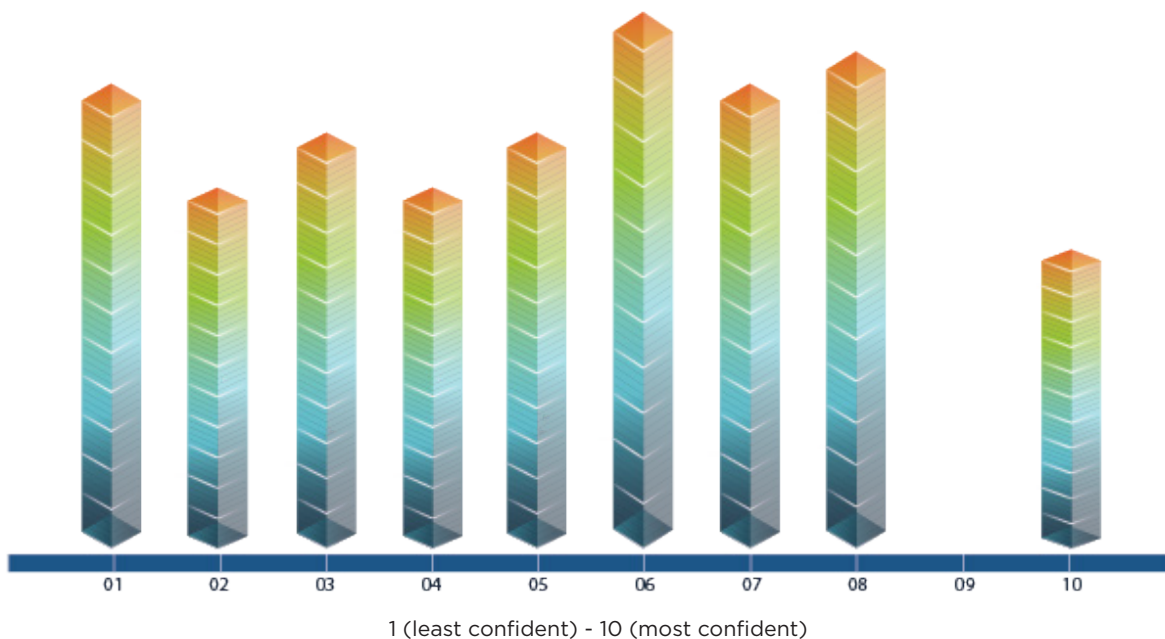
## Emerging Security Threats

Attacks today can be zero-day and polymorphic, meaning enterprises can no longer predict the nature of the attacks or the form they take. For example, attackers today have the means to gain privileged access, replace files, erase logs, and change the behavior of malicious codes over time. Given the ever-changing nature of threats, organizations that have a set of stale definitions, even if they are updated on a semi-regular basis, cannot be said to be well-protected. The only way to remain secure is by persistent and deliberate monitoring and management, and a conscious effort to enhance and improve security.

Yet, large security teams are rare, even in large MNCs. Usually, a single individual within the security team has to juggle multiple workloads, such as managing both the firewall and the Security Operations Center (SOC). Responsible for monitoring so many aspects of the network, an individual is overworked and distracted, and cannot be expected to remain alert to unpredictable threats within the network.

In a poll conducted by Kemp, out of 63 senior IT officials, only 3 respondents felt fully confident that their organization was not experiencing any ongoing cybersecurity attack in their IT network. Twice that number, i.e., 6 respondents, were not confident at all, with the average level of confidence ranging somewhere between 3-4 on a scale of 1 to 10.


## How confident are you that there is not any ongoing cybersecurity attack in your IT network right now?



- **Ransomware, the No. 1 threat today**

In the poll, more than half the respondents said their biggest security concern was ransomware.

### What is your biggest security concern?



In the past year or two, ransomware has certainly been on the rise, with tens of millions of ransomware attempts detected last year. Today, there are even providers of Ransomware-as-a-Service platforms, which provide mechanisms for criminals to build their own ransomware practices.

The ever-changing nature of ransomware oftentimes allows it to slip under the radar of the perimeter, bypass

end-point security, and get lost in a bunch of false positives or hide in blind spots in network visibility. Network Detection and Response (NDR) and NPMD solutions address these issues and enable proactive protection. Kemp's Flowmon solutions are able to detect the ransomware attack several times during its deployment, using a combination of AI, Machine Learning and behavior patterns.



- **Compromised Applications: An Emerging Threat?**

The threat of compromised applications has emerged over the past year, with some high profile cases publicized in the past few months involving companies such as Microsoft and Citrix.

Compromised applications cause pain points for customers, especially since applications have now become critical to an organization's IT infrastructure. Whereas in the past, applications were project-based and worked on by small teams, today, they are core business operations.

Preparing to deal with a compromised application means organizing the rest of the network infrastructure, such that other elements in the network can be used to mitigate the situation. Load balancers, firewalls, next-gen firewalls, and other components should be activated to work around a compromised application so that pain points can be quickly isolated. Also, organizations should work with vendors to detect and mitigate potential compromises that may exist in their application services.

## Bridging NetOps and SecOps

An IT department consists of many individuals – NetOps, SecOps, network engineers, application developers, data architects, and so on. With so many domains involved, an organization needs to build teams that are conducive to ensuring that the application experience remains smooth, while being secure. Different groups need to work together as a team, even if they are in different parts of the organization.

- **The Danger of Silos**

In the traditional organizational structure, an IT department has totally separate and isolated groups. Architects sit at the drawing board, furthest away from the end-user. Then come the engineers, who sit closer to the infrastructure, but still do not interact with users. Finally, operations teams are responsible for running the environment. In other words, there is a vertical, hierarchical structure within an IT organization.

Such an approach is no longer practical today, because teams that do not interact directly with the end-users tend to detach themselves from day-to-day occurrences in the network. Even the various operations teams might focus only on their specific domains or silos, and might not be motivated to help fix problems in other areas. For example,

access management teams might not concern themselves with network performance, and so on.

The trend today is to blend the teams more, so that those doing the architecting and engineering are closer to the end-user, and all the domains contribute to overall performance.

- **Breaking down the silos**

Blending teams is sometimes easier said than done, because doing it successfully requires a shift in mindset. To effectively break down silos and have different parts of the organization collaborate as an organic unit, the entire organizational structure of an IT department should

be re-thought. Rather than the traditionally vertical, hierarchical structure, there should be a horizontal, cross-functional structure. Moreover, the challenge

that stands in the way of merging teams to functional effectiveness is aggravated by the fact that they rely on their disparate tools, which rarely enable cross-functional collaboration by default.

*Example:*

*The effectiveness of blended teams shows itself in a simple example:*

*A security problem is typically addressed by making some sort of a networking change. For example, a server might have to be locked down. Even if something needs to be unplugged in a critical scenario, or a security group modification needs to be implemented, it may entail making a change at the network level. Since both NetOps and SecOps are necessary for threat mitigation, it makes sense that they should collaborate and duplication of effort minimized.*

NetOps and SecOps often use the same data to extract insights, but examine it through different lenses. Although it might not appear so to specialized teams, network problems are often security problems, and vice versa. The same data sets can be consolidated, then interpreted in ways relevant to both domains.

An important aspect of blended teams is to have shared metrics so that all the groups work towards the same outcome. Networking teams need to be motivated to help security teams, and vice versa. Having the right tools can make it easier for these metrics to be measured, and for collaboration in general.

Ultimately, the different teams should have joint responsibility and accountability. When a team is tasked with monitoring, for example, it should also be responsible for fixing problems. Every possible vector represents a possible breach, and the network should be looked at holistically.

*Example:*

*Organization A experienced a ransomware exploit, wherein a bad actor was attempting to exfiltrate data. At the same time, NDMP detected ping traffic that had payloaded it. Simultaneously, there was unusual escalation on the EP server, and then admin access to a Samba share. By looking at all the communications happening over the network, the organization was able to stitch the picture together quickly.*

## Automation in workflows

Many enterprises are uncomfortable with automation, partly because it is so new. They are unwilling to put complex decision-making out of the hands of trained personnel - especially when it comes to areas that could potentially negatively impact application experience. Moreover, automation landscape is constantly changing with a rapidity that can be daunting; each day, new automation tools are developed that need learning. However, keeping up with the new developments is imperative - these tools can increase efficiency and leave IT teams free to work on higher-value problems.

- **What should and should not be automated?**

Automation should primarily be used in repeatable and predictable tasks, on workflows with a high level of context. Rules can then be applied, the initial direction set, and subsequently automated.

*Example:*

*A simple example would be a log event of a specific level of severity on a firewall. To address this, there may be some additional information that should be captured, or perhaps, a full packet capture between the two communication entities identified in the log. Normally, an operator will have to perform these tasks manually. This is where automation can help.*

In spite of the rapid progress in AI/ML, there are some tasks that are still better done manually. Complex and unpredictable scenarios are an example, especially when issues cannot be foreseen. In this case, operators are often unsure of the exact type of data or information needed for analysis and mitigation. In other words, there is no playbook in place. In this case, automation can be challenging.

- **An Automation Strategy is Crucial**

Preparation is crucial before starting to automate. In order to get the best results, the process has to be structured and well-mapped out, even before automation tools are implemented. The capabilities of human intuition are sometimes taken for granted, and the effectiveness of AI overestimated. While automation can achieve significant results, the behavior of AI/ML does not replicate human behavior. If proper preparation is not done, a “garbage-in garbage-out” scenario might arise.

## Application Experience in a Multi-cloud Environment

Whether an organization is operating in a hybrid cloud or a multi-cloud environment, ensuring a stable, consistent application experience can be tricky - especially when migrating to the cloud. Legacy applications built decades ago are moved to the cloud with great difficulty, only to experience a decline in performance afterwards.

Before moving a legacy application to the cloud, an analysis needs to be conducted first, to measure readiness and compatibility. This is because cloud ecosystems are built to optimize certain application architectures. Legacy three-tier applications, for example, are not built for the cloud - they do not have real-time connectivity to the backend database, are stateful in nature, and cannot handle lossy connectivity to the storage ecosystem. In such cases, doing a straight lift-and-shift from on-premise to the cloud is likely to result in poor performance. Before the migration, some level of optimization - even if the whole application is not rewritten to be cloud-native - may need to be done.

When migrating to the cloud, security also needs to be carefully considered. Although every type of situation is different, generally, the most crucial part is access management. Applications that are moved to the cloud need to be managed carefully in terms of who is

accessing that application, and proper encryption needs to be done to make sure the application can be trusted by all the users.

## Diagnosing and Remediating Slow Applications

Diagnosing and remediating slow applications can be troublesome, given the number of factors that could possibly be causing the issue.

In a poll conducted by Kemp, the majority of respondents pointed to an overloaded server as the primary cause of slow applications, followed by slow or poorly written database requests. For nearly half the respondents, the issues lay at the network level, namely, due to problems with the load balancer, or network performance issues.

### When you get a slow application, what usually turns out to be a reason for it?



<b>Overloaded server</b>	32.4%
<b>Database request too slow to execute or poorly written</b>	27.0%
<b>Network performance issue</b>	18.9%
<b>Issues at the load balancing level</b>	13.5%
<b>Slow application code execution</b>	5.4%
<b>Slow web request</b>	2.7%

Given the numerous entities driving the application, there can be many reasons why an application is slow. In identifying the problem, the first step is to do a process of elimination to start ruling things out as a problem. For example, to check whether the load balancer is causing the problem, a client connection could be issued directly to one of the application servers. Or, a firewall could be bypassed and the speed monitored. A process of elimination gives only some information, but it helps unravel the mystery and indicate, for example, whether it is only a subset of the application servers that is experiencing the issue.

Although this may seem like a daunting task, there are automation tools to do such routine testing. Because there is a fixed set of steps and a playbook of rules, and is predictable, it is possible to automate the process. Additionally, NPMD tools can also help by baselining the server response to the network response, and to the jitters or retransmissions occurring. The goal of all this, ultimately, is to reduce the number of possible causes so that manual operators can hone in on a few specific possibilities.

Another approach would be to first look at the entire infrastructure and go through every layer. Here again, the same process of elimination is conducted. Typical investigations would include checking if the server is timing out or overloaded, or if the memory is on an optimal level. Such an approach is more traditional, and is especially helpful for organizations with on-premise infrastructure. If the application is cloud-native, the ability to scale nullifies the problem of overloaded servers, and a network-first approach would probably be more fruitful.

It is not typical that all the servers within a cluster face issues at the same time, unless there is a systemic problem. Usually, a slowdown is caused by a specific instance within the ecosystem. Having detection methods to narrow down to that level of granularity

is key, and can be achieved at the architecting stage, when mechanisms should be put in place to be able to extract that information. They could be placed on the application ecosystem and the elements of the network that connect to it, be it via network performance or via the configuration of the load balancer.



*TIP: It's not always a hardware issue*

*Many IT organizations take for granted that issues with application experience, such as the speed of applications, arise because of hardware problems. Therefore, to fix the problem, they add more hardware to their networks. However, in many cases, all that is required are some minor tweaks and enhancements.*

## Measuring Success

Even after the appropriate NPMD and security tools are put in place, IT teams still frequently face complaints relating to application experience. For example, an employee in another department might complain that the system is slow, or that customers are dissatisfied with the application experience.

Listening to such feedback is important because it often comes from the field and may result in a huge difference in the revenue and bottom-line of the company. Dissatisfied customers can sometimes mean the loss of millions of dollars in potential revenue.

IT teams can often be ill-equipped to respond to such complaints. One way to overcome this, and to build a bridge between IT and business, is for IT teams to use their own applications and to go on the customer journey themselves. IT teams need to act as “mystery shoppers” – if they have created an application for car loans, for example, they would do well to apply for a car loan themselves, to understand it better from the customer’s point of view.

## • Coming up with Metrics

Once an IT team has gone through the user experience, they will be in a better position to have a frank discussion with the rest of the company.

This discussion should start with IT establishing a common understanding on what might be a good or bad user experience. As experiences can be subjective, disagreements may arise unless there are clear metrics on what is expected of the application.

Such metrics can be crafted by asking specific questions, for example:

1. *How long does an application need to take to load for it to be considered “slow”?*
2. *Are there too many steps to authenticate the user? How many is too many, and can it be reduced?*






Traditional ways of approaching issues relating to applications and security are no longer as effective because of the acceleration of digital transformation, Work From Home, and other shifts in the way the world does businesses.

Applications have become the conduit through

which customers experience the value businesses or organizations have to offer, and can mean the difference between winning or losing in the marketplace. IT executives need to rethink their strategies and initiatives as they plan for the future.

## About Progress

Dedicated to propelling business forward in a technology-driven world, [Progress](#) (NASDAQ: PRGS) helps businesses drive faster cycles of innovation, fuel momentum and accelerate their path to success. As the trusted provider of the best products to develop, deploy and manage high-impact applications, Progress enables customers to build the applications and experiences they need, deploy where and how they want and manage it all safely and securely. Hundreds of thousands of enterprises, including 1,700 software companies and 3.5 million developers, depend on Progress to achieve their goals—with confidence. Learn more at [www.progress.com](http://www.progress.com)

-  /progresssw
-  /progresssw
-  /progresssw
-  /progress-software
-  /progress\_sw\_

2022 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved. Rev 2022/01 RITM