BROADCOM

WHITE PAPER

Validating Cloud Connections: 3 Keys to Success for Network Operations Teams

TABLE OF CONTENTS

OVERVIEW

Virtually every enterprise in the world relies on the cloud in some fashion. Because of that, these organizations are also fundamentally reliant on third-party networks to access these critical cloud resources.

Cloud adoption decisions are largely made by executive teams. However, when it comes to managing access to cloud services and addressing issues when they arise, the responsibility falls squarely on the shoulders of network operations teams.

While spending on the cloud increases, network operations teams struggle to contend with visibility gaps posed by the cloud services their organizations have adopted. This lack of visibility results in wasted time and resources. Teams spend significant time trying to resolve application and network performance issues. Lengthy troubleshooting prevents network operations teams from spending time on more valuable strategic initiatives. Ultimately, this lack of visibility presents increased risk to the business.

To overcome these challenges, teams need to employ a sound network operations strategy for reestablishing comprehensive visibility, including into cloud and transit networks. They need to be able to ensure that applications and networks meet internal and external service level agreements (SLAs) centered on performance.

As teams look to optimize connectivity for internal or external consumers of cloud services, they need to factor in two key aspects:

- **Cloud migration.** Cloud migration is the process of moving existing apps and services to cloud environments. This process ultimately introduces logical and physical distance between users and the apps and services they access, gaps that didn't exist before. To ensure optimal performance, teams need to establish new approaches for monitoring these connections and remote services. If network operations teams fail to budget or plan for these new requirements, the progress and potential benefits of these transformations can be stifled.
- **Cloud adoption.** Cloud adoption is the process of replacing historically on-premises apps with cloud-hosted services, such as SaaS apps, which are outside of network operations teams' control. While many cloud services provide very basic up/down metrics, these offerings lack in-depth performance insight from the end-user perspective. Consequently, these offerings present significant limitations for teams doing troubleshooting.

Understanding when a service has an outage or is underperforming is only useful if teams can determine whether their organization is affected. In most enterprises, when performance issues arise in applications hosted partially or fully in the cloud, and these issues can't immediately be explained by cloud outages, a "war room" style response is initiated. This means members of disparate network operations, security operations, and development teams have to come together to identify the root cause of an issue.

These calls are costly and often result in finger pointing and increased animosity. Exacerbating matters is that each team only presents data that confirms the issue is not in their sphere of responsibility—rather than pinpointing what the actual problem is and where it is located. Typically, each team has their own siloed toolset that provides them with the insight they need for their roles, but no central view is available.

Increasingly, network operations teams are shouldering the responsibility of network performance, even when they don't own the infrastructure. It's important to gain visibility in order to validate cloud connections that support business-critical apps and services. To do this, network operations teams need to invest in a solution that provides these key capabilities:

- **Deliver end-to-end visibility across both internal and externally managed networks.** Teams need evolved tools. Traditional passive monitoring tools need to be augmented with active monitoring capabilities. It is only with these combined capabilities that teams can gain the true, end-to-end insight required to map where traffic is going and determine where the root cause of performance bottlenecks is located.
- **Add application context to network insights.** Network operations teams can no longer simply rely upon the basic network data of the past. As applications become more diversified and users become more mobile, it is more critical than ever that teams understand the nuanced business context of complex applications.
- **Track SLAs for networks and services.** With most modern apps and third-party networks, SLAs are no longer guaranteed. Therefore, when performance issues arise in external service provider environments, network operations teams need objective data to prove that.

The following sections provide an in-depth look at each of these capabilities.

ESTABLISHING TO END-TO-END NETWORK VISIBILITY

Cloud migration and adoption transformations present unique challenges for today's network operations teams. Significant infrastructure components that are outside of the organization's control now sit between business-critical services, end users, and ultimately productivity. This physical and logical distance limits the visibility of network operations teams, impeding their ability to isolate and identify issues.

When prolonged issues occur, it undermines the rest of the organization's confidence in the network operations team. This is true even if issues arise that are due to third parties or networks outside of the network operations team's control, for example, the downtime of a SaaS application or an outage in an ISP network. What's worse is that, without end-to-end visibility, network operations teams cannot efficiently prove their innocence when pulled into war rooms for issues affecting large user populations.

Without a modern solution, the only way to achieve full end-to-end visibility is through manual correlation of pings and traceroutes. Using that rudimentary approach, insights are limited at best, and teams are stuck operating in a reactive mode. Additionally, with multiple ownership entities involved, isolating the root cause is often not possible through traditional means like SNMP or NetFlow, since third-party providers typically don't offer this data to external entities.

IP SLA monitoring has traditionally been another popular approach for active monitoring. This approach introduces a heavy load on networks. This overhead introduces the potential that teams will have to incur the added cost of deploying additional networking equipment, such as routers. This approach also has a number of limitations, including the fact that testing can only be done between two owned endpoints.

Fundamentally, this represents a legacy approach. SaaS and cloud applications cannot be monitored in this way, since visibility would be limited to the LAN side of the network path. Given this added resource and management overhead, IP SLA testing is not a scalable solution for gaining a meaningful, end-to-end perspective.

Employing traditional approaches, teams can't establish complete, end-to-end visibility. Through legacy approaches, visibility either ends at the firewall or offers ineffective data for troubleshooting purposes. These approaches typically fall into three categories:

- **Reactive.** Invariably, it is only when users have already been affected that teams can start to understand what happened. Reactive approaches are often based on manual troubleshooting. Relying on tickets to be entered before level-one network operations staff are aware of issues leaves the business exposed and undermines confidence in network teams.
- **High overhead.** Technologies like IP SLA and even some NetFlow implementations can add significant weight to the network in the name of monitoring it. They can only provide behind-the-firewall visibility and are subject to limitations due to security considerations.
- **Third-party reliance.** For outside-the-firewall visibility, teams must rely on the status pages offered by SaaS app and cloud providers or sites like DownDetector. What these solutions lack is insight into whether an organization's users are affected, and if so, which users. This lack of insights leads to longer mean time to resolution.

To overcome these limitations, teams must establish a more modern approach. Teams need a solution that combines both active and passive monitoring. Passive testing is employed to cover the internal local area network (LAN) infrastructure and active testing is used to cover external wide area network (WAN) paths. Active testing can provide synthetic transaction monitoring, which offers vital application context. The combination provides continuous insight for today's dynamic WANs. By establishing a centralized data repository, teams can layer analytics on top of network data to correlate issues from end to end.

NetOps by Broadcom represents a unified solution that brings together all types of visibility. The solution features automated discovery capabilities for the internally managed network and it brings active monitoring insights to the external networks that teams don't own or control. The core of that active monitoring is a technology called TruPath.

TruPath Packet Train Dispersion

TruPath sends and receives many varied short sequences of packets, which are referred to as packet trains. Packet trains are transmitted using Internet Control Message Protocol (ICMP) or User Datagram Protocol (UDP). Packets are sent to defined end hosts or targets, which can be any endpoint that can respond to an ICMP-based ping or can send back a Transmission Control Protocol (TCP) or UDP packet.

Using this technology, TruPath can build up a complete set of network statistics very quickly—in many cases, in just tens of seconds. TruPath uses special patterns designed to detect if instrumentation packets are interfering with each other. If that happens, it takes more varied samples over a longer time scale to ensure that the resulting statistics are clean.

By sending multiple sets of distinct packet sequences, TruPath can analyze a wide range of different traffic conditions that a user on a network path might experience. By probing the path repeatedly with the packet sequences, TruPath collects a statistically significant collection of responses for each type. TruPath will detect when samples are captured during times of rapidly changing conditions and adjust its measurement patterns accordingly.

Unlike so-called "packet flooder" technologies available on the market, this approach delivers high accuracy without requiring an intrusively high instrumentation load on the network path. With the technology's low-overhead approach, teams can run TruPath in production and through third-party networks for end-to-end visibility.

ADDING THE APPLICATION CONTEXT

While the focus continues to be on the network, application awareness is critical when issues arise. Teams need to understand what apps are running on the network and their impact on other network functions, including other apps. Ticketing and user complaints are often application-based, but it often becomes the network operations team's responsibility to manage troubleshooting. These teams often have a limited subset of data and are left searching for information on the user, app, and network connection involved. This leads to lengthy triage as network operations teams have to align apps to networks on the fly.

There are a couple fundamental issues teams face when trying to translate complaints about apps into network troubleshooting. First, trends like shadow IT mean that network operations teams don't always know what apps are traversing office networks. Identifying what those apps are is important for both performance and security concerns. Second, relying on traditional tools, network operations teams have a limited data set for isolating application issues outside of TCP packet data. By gathering new data sets that help isolate issues faster, teams can ultimately reduce troubleshooting time.

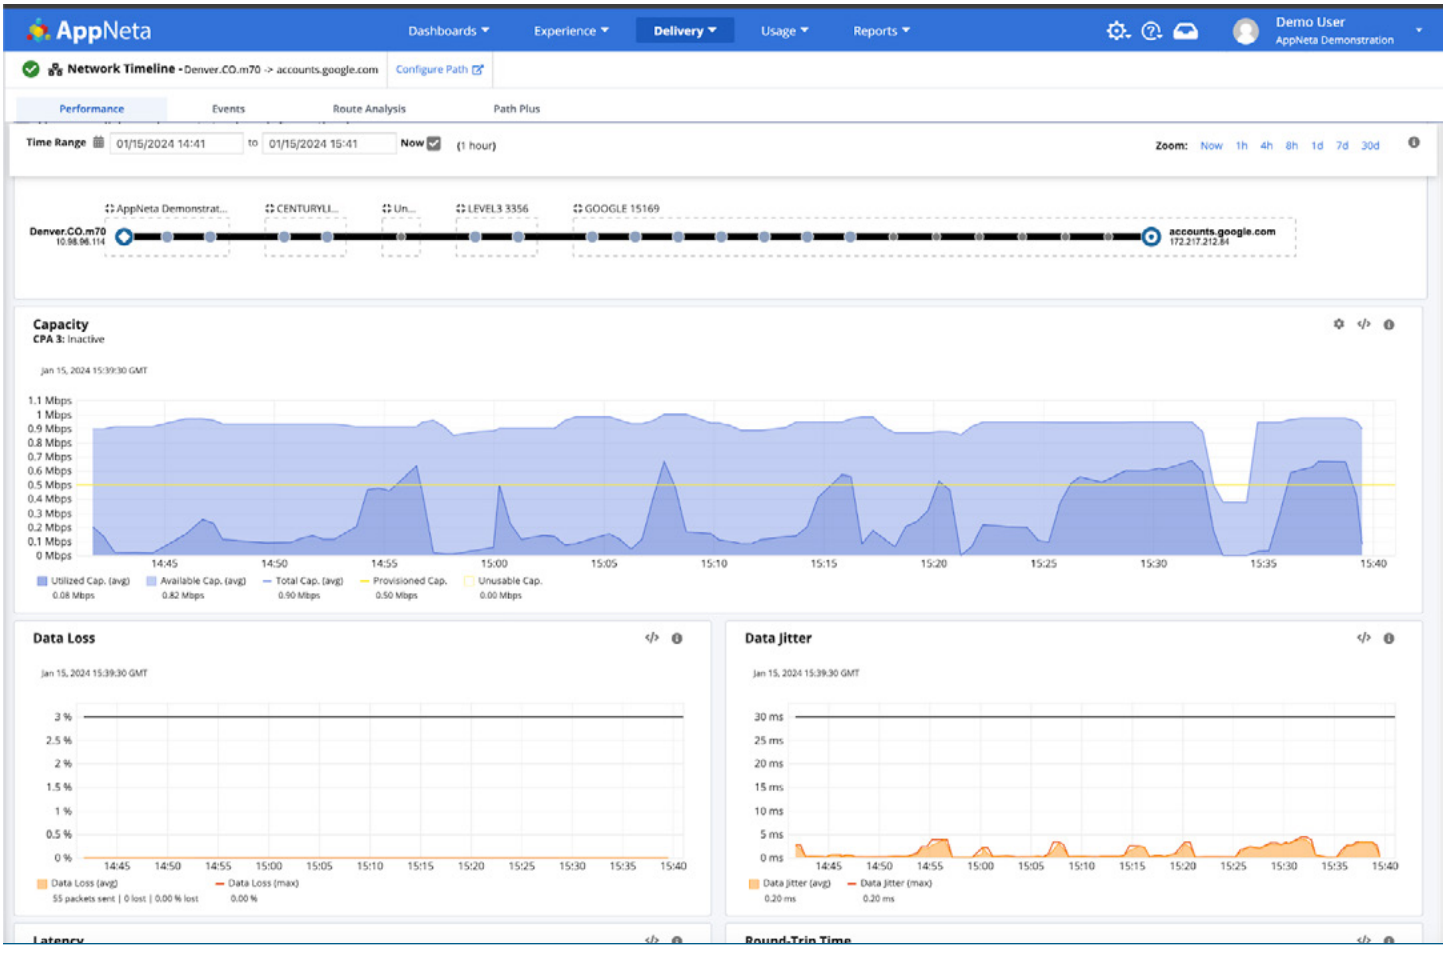

Figure 1. Reporting on continuous testing to a Google endpoint, offering capacity, data loss, and jitter metrics.

Current Approaches for Gathering Application Data

Understanding what apps are running on the network is easier said than done. Passive techniques like NetFlow allow for some protocol-based isolation of traffic. However, this technology's NBAR2 application identification engine still leaves large buckets of data in TCP, HTTP, and HTTPS groups, which makes it difficult for administrators to parse.

Identifying endpoints, either source or target IPs, can be useful in understanding traffic, but more sophisticated methods are needed. Additionally, most enterprise NetFlow implementations sample traffic instead of capturing 100% of traffic for analysis. This means that it's possible to have applications fly under the radar until they become a problem.

If the network is exonerated by network operations data, then understanding application performance is challenging, regardless of whether or not the app is owned by the enterprise. Here's more on the challenges each type of app presents:

- **Internal Apps.** Often, network operations teams have to rely on data from applications teams to get "user experience" insights. However, these teams often use typical application performance monitoring (APM) approaches, which only track TCP latency from the web server side.
- **External third-party Apps.** Teams can access status pages from the application service provider. Typically, these pages simply offer an availability icon for the app infrastructure that is red or green, depending on its status. These pages don't include any insights in terms of whether users can access the app from outside locations.

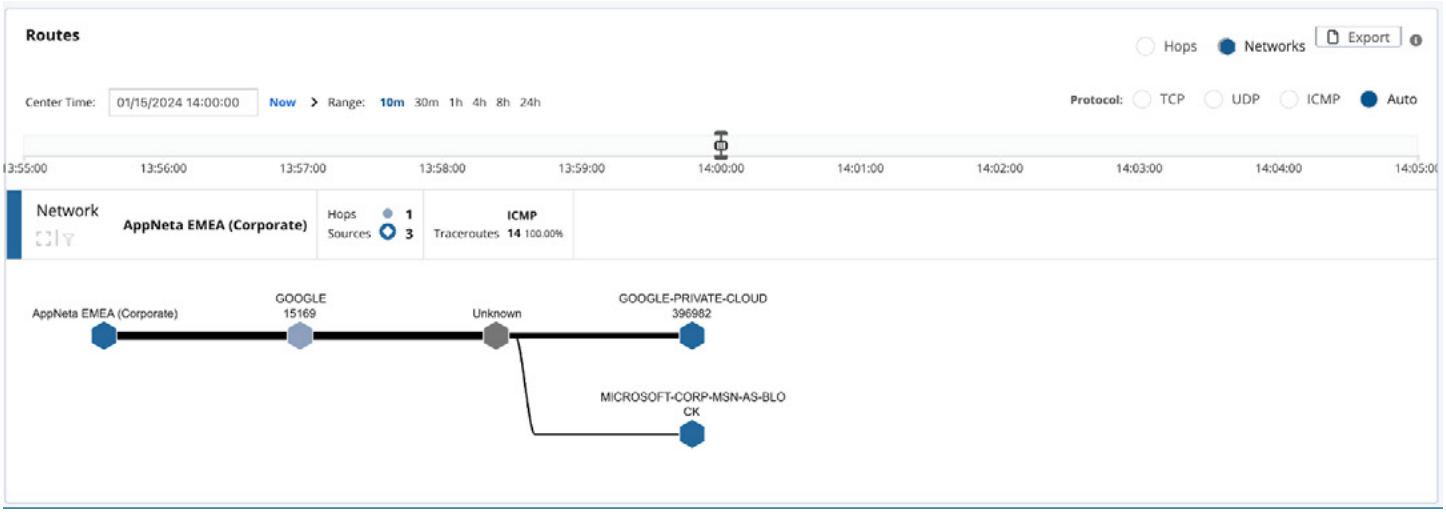

Figure 2. Macro-level Autonomous System network aggregation based on continuous testing to two cloud endpoints.

As a result of this challenge, when application issues arise, IT and network operations teams are stuck having to wait for status updates. These teams can provide updates to the user base, but they have no power to resolve the issue. Longer term, these issues may lead decision makers to change service providers, cloud providers, or application vendors. However, in the near term, these recurring issues can lead to eroded confidence in network operations, regardless of who's ultimately responsible for the issues. Fundamentally, network issues outside of corporate infrastructure represent a blind spot. General internet status can be identified for individual issues, but this is a retroactive approach that only works on a small scale.

Proactive Application Visibility

In order to move from a reactive stance to a proactive one, there are two requirements that need to be addressed. Network operations teams have to know what apps are traversing the network and understand which ones are critical to the business. This can be done by looking at network traffic. Once apps are identified, network operations teams can create active testing to those apps to understand the network delivery paths between apps and users. The primary way to achieve this is through synthetic transaction monitoring.

Deep Packet Inspection

Deep packet inspection (DPI) can be used to identify what apps are running on the network. DPI requires monitoring deployment at the network edge to generate a running list. This will ensure that network operations teams are on top of new applications running on the network, even those introduced by shadow IT. They'll also be able to determine which apps are consuming significant resources. By isolating top apps by traffic, teams can identify which apps are business-critical and which are not. In this way, they'll also be able to see whether low-priority apps are negatively affecting business-critical applications.

Isolating top traffic can provide a number of benefits, but chief among them is visibility into the experience of end users accessing business-critical applications. Here are key capabilities offered:

- **Isolating application versus network issues.** By looking at latency and TCP retransmits across multiple apps, teams can quickly identify when issues affect all apps, just some apps, or a single app.
- **Categorize impact.** Over 2,000 apps are included automatically in the application identification engine. Teams can use categories to understand if apps are critical to the business or recreational apps that are hogging resources. Advanced solutions enable teams to add custom apps or make changes in their organization's default settings.
- **Cross-location visibility.** This visibility enables network operations teams to dynamically visualize application use across locations for comparison purposes. In this way, they can isolate when apps are performing poorly for a particular employee cohort.

Web Synthetic Transaction Monitoring

Synthetic transaction monitoring is more widely known and used in the APM world, but unlike the DevOps teams that typically employ an APM solution, network operations teams have far less insight into an application's functionality. Further, these network teams often just need to isolate if and when the application is the root cause of a reported issue.

Unlike a simple ping of a web server or an HTTP GET request, web synthetics test beyond the login page of apps to mimic actual user workflows, enabling operators to better understand the performance of modern apps. Using logical milestones within scripting allows teams to align monitoring with end-user workflows. This helps in understanding the true performance of modern applications, which often fall into one of these two categories:

- **Microservice applications.** These apps have separate systems that typically speak to each other via APIs. This means that while monitoring the front end can provide some insight, there are additional network calls on the back end that might not be detected, but that may be responsible for slowing down performance.
- **Single page applications (SPA).** These apps do not load a new page when a user interacts with them. If monitoring is set up to use a client URL (cURL), then only one result would be available for all functions of the page.

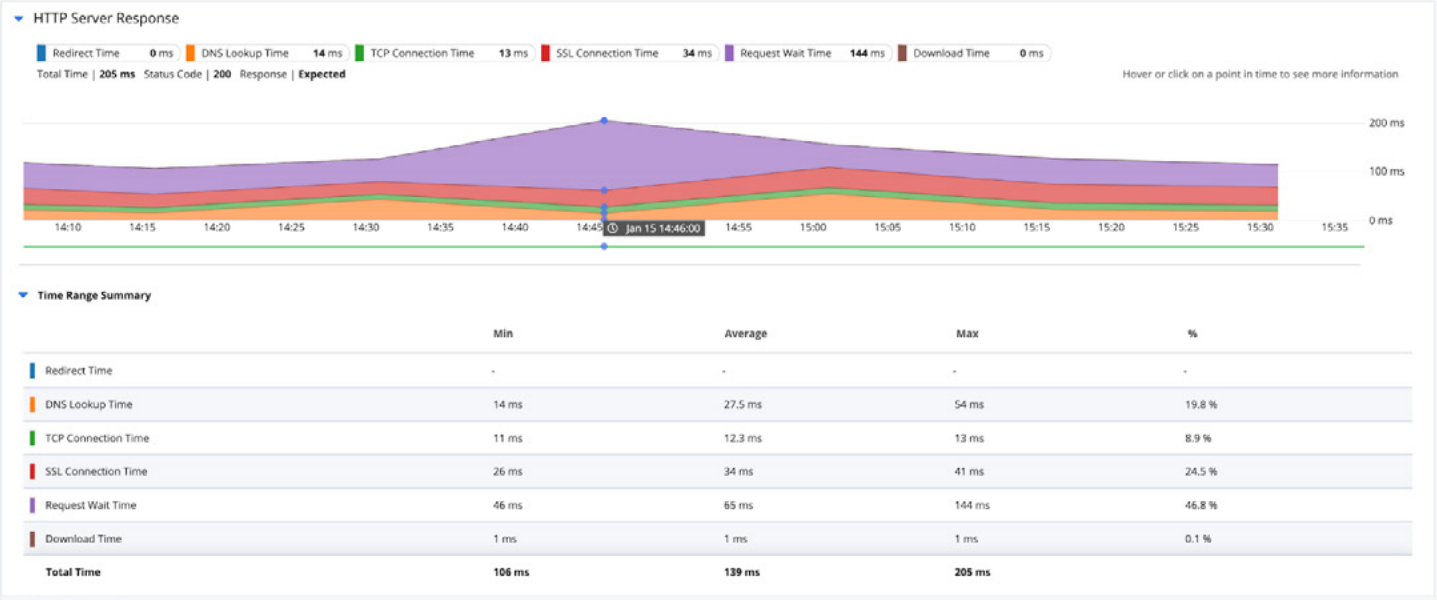

	Min	Average	Max	%
Redirect Time	-	-	-	-
DNS Lookup Time	14 ms	27.5 ms	54 ms	19.8 %
TCP Connection Time	11 ms	12.3 ms	13 ms	8.9 %
SSL Connection Time	26 ms	34 ms	41 ms	24.5 %
Request Wait Time	46 ms	65 ms	144 ms	46.8 %
Download Time	1 ms	1 ms	1 ms	0.1 %
Total Time	**106 ms**	**139 ms**	**205 ms**	

Figure 3. Chart depicting the timing information for an HTTP server test.

ESTABLISHING SLAS

When referring to network or cloud connectivity, a service level agreement (SLA) can mean a few different things. SLAs for traditional connectivity like MPLS were enforceable, monetary, and documented. Given the current reliance on third-party network connections, these types of SLAs are increasingly rare. Understanding when connectivity is falling below stated values becomes a burden of proof.

Internal SLAs for network operations teams and level-1 support revolve around ticketing and how quickly responses will start. However, due to the complexity of modern networks, these SLAs rarely guarantee a timely resolution. In some organizations, network and infrastructure teams provide SLAs to application teams.

External SLAs involve different connectivity types like direct connections and ISP links, which can come with SLAs attached to them. SaaS apps may have SLAs based on simple uptime metrics. Beyond those, cloud SLAs are uncommon. Typically SaaS and web apps do not include reimbursement or even guarantees of performance outside of a commitment to some form of uptime and availability.

The Challenge of SLAs

There is no one single route into a public or private cloud. There are many routes that can be taken, including through direct connections, tunnels over public internet, and public internet gateways. The design of each network dictates what the preferred route should be. Regardless, however, understanding what route the traffic is taking at any given time is always a challenge.

For application performance to be the same or better in public cloud or hybrid cloud deployments, there must be strict SLAs that govern connections between on-premises and cloud environments. For example, an organization in the financial sector had a banking application that had excellent performance globally. However, they started to see issues arise in specific regions. Latency between the application and database was too high, which created critical errors. For example, the account balance for online banking customers wouldn't display due to the lengthy delays.

It is easy to measure and enforce SLAs at a single point in time when administrators can flood the network, or average tests can be made at a set interval. However, these approaches aren't practical in modern application environments. For SLAs to be properly enforced, there needs to be continuous monitoring, but this monitoring can't have a meaningful impact on the values being tracked.

For many network operations teams, there is no viable alternative to what has been standard operating procedure. When outages occur, users contact support, teams issue communications about the outage internally, and start troubleshooting. Typically, in cloud and hybrid environments, there is a lot of finger pointing between application and network infrastructure teams when issues occur. Fundamentally, everyone struggles to determine whether it is the application that isn't working or if it is the network that's causing application performance issues.

A New Approach: Active SLAs

With active monitoring of any SaaS or web service, network operations teams can keep track of performance and report on SLAs by watching derivatives of performance. Much like financial markets pay attention to the rates of change, network operations teams benefit from looking beyond the trees to see the forest on a macro level. Once a team moves from reactive to proactive approaches, understanding the frequency and timing of aberrations represents the next phase of evolution. This will yield insights that help teams combat endemic issues that standard monitoring may not uncover. Active SLA monitoring helps in both of these arenas:

- **Internal.** For internally managed environments, active SLA monitoring offers insights for making more informed decisions on future technology purchases, doing capacity-minded resource augmentation, and planning effective rollouts of new technologies.
- **External.** In externally managed domains, active SLA monitoring helps teams isolate the source of an issue and identify who is responsible for resolution. The data captured can be furnished to vendors, enabling true bidirectional interaction, which leads to more sustainable and automatable long-term processes.

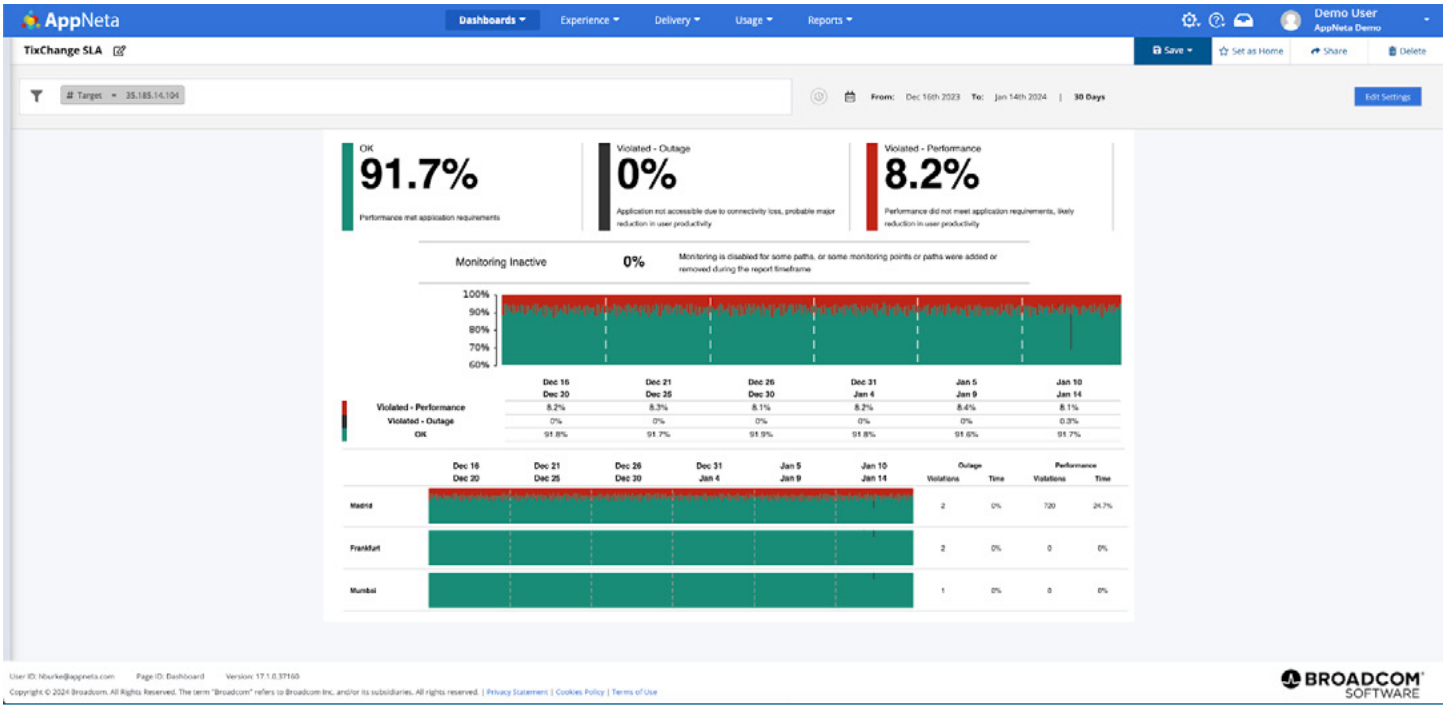

Figure 4. Report revealing the SLA performance of the TixChange app.

NETOPS BY BROADCOM

NetOps by Broadcom offers active monitoring and location quality reporting. Network operations teams set up alert thresholds in the solution and use them to determine when an application violates the performance boundaries established. Teams can also look at performance from across different offices to isolate issues that may be specific to a single user or location.

The Broadcom solution can help network operations teams create standard operational procedures by providing baseline data across locations, apps, networks, and devices. With this visibility, teams can understand what "normal" operation looks like. The solution can validate cloud connections in several ways:

- **End-to-end coverage.** Identifying and isolating issues is simple if teams have visibility into both internally owned networks and externally managed networks.
- **Continuous active monitoring.** The key to monitoring dynamic external networks is monitoring them over time and building up historical data on performance, routing, and the upstream providers involved.
- **Automated baselining.** Once monitoring is implemented, understanding what normal performance looks like is the next step. Once that understanding is established, alerting can be more intelligent, based on meaningful deviations from normal.

CONCLUSION

Network operations teams have learned all too well that, no matter which type of cloud service their organization employs, they still remain responsible for isolating and troubleshooting issues that affect users or business-critical apps. Traditional monitoring and management solutions can't provide end-to-end visibility across all the networks that organizations are now reliant upon. By adding active monitoring to their network operations toolkit, teams can increase visibility and also add application context to network insights. Once teams establish centralized visibility of their applications and performance data, they can more effectively monitor SLAs for networks and services.

NetOps by Broadcom can address the challenges of cloud connections from a few perspectives. From the network perspective, the Broadcom solution covers some of the largest environments in the world. The solution's end-user experience monitoring provides a unique view into external networks, enabling complete coverage across the footprint of organizations' hybrid infrastructure and cloud services. By gaining visibility into the end-user perspective from the network edge, network operations teams can better ensure compliance with performance SLAs, both for internal users and with third parties.

With the Broadcom solution, network operations teams gain the visibility required to enhance connected experiences for any user and any application, no matter who owns the network infrastructure. When issues arise and war room meetings are convened, network operations teams need to determine if their network infrastructure is to blame, and, if not, provide evidence that their environments are not the culprit. Gaps in visibility or data make this task impossible. By identifying where the gaps are, teams can be better prepared to pinpoint the location of issues and demonstrate innocence.

FIS EXPANDS SLA COVERAGE

FIS, a leading financial technology company, sought to ensure optimized network service delivery, including when user traffic is reliant upon networks that reside outside the traditional borders of the enterprise. Both FIS and its customers have grown increasingly reliant on cloud services, which means they're also increasingly reliant upon cloud providers' networks. Plus, with rising support of work-from-anywhere approaches, users now count upon a diverse set of networks, including local Wi-Fi and third-party ISPs.

"In recent years, we've been contending with an ever-expanding spider web of networking," a systems engineer at FIS noted. "Where communications used to be contained within a data center, the user experience continues to depend on networks and environments outside of the data center. To be able to continue to deliver 24/7 availability to clients, we have to be able to successfully monitor those domains as well."

Through NetOps by Broadcom, FIS gained unified visibility into legacy networks, SD-WAN, and user experience. FIS was able to reduce SLA breaches and penalties. In the past, teams were coming close and sometimes missing SLAs.

"With the Broadcom solution, our team is better equipped to meet or beat our SLAs," a systems engineer explained. "This means we can avoid the financial penalties and poor customer experiences associated with SLA breaches."

In addition, the team was able to accelerate triage by up to 95% and improve customer satisfaction.

WHY BROADCOM

NetOps by Broadcom combines active and passive monitoring technologies to provide unparalleled visibility across internally managed and externally managed networks. The solution delivers the insights needed to boost network management, monitoring, and observability.

By remaining vendor agnostic and focusing on data analysis, Broadcom can provide third-party validation to any network and at any scale. Proven in the largest environments, the Broadcom solution eliminates visibility gaps and delivers a comprehensive network observability and management solution designed with tomorrow in mind.

With the Broadcom solution, customers can optimize internal network operations, accelerate network transformations in the cloud, and enhance connected experiences.

Validate-Cloud-Connections-White-Paper_BC-NO_CE-4116_v2 February 1, 2024