A Smarter Way To Monitor the Cloud Before and After Migrations

And Ensure Quality User Experience

ENTERPRISE

TABLE OF CONTENTS

Executive Summary 3
Trends in Cloud Migrations 4
What are End-Users Expecting? 4
Cloud Migration Trends and Challenges 5
Assuring User Experience Before and After Cloud Migrations 6
End-Through-End Visibility and Control 6
Client Edge 7
Network Edge 7
Data Center / Cloud Service Edge 7
How To Optimize Visibility and Control 8
Packet-Based Monitoring 8
Synthetic Testing and Monitoring 8
Superior Visibility Offered with Combination of Packet and Synthetic Monitoring 9
USE CASE / VISIBILITY BEST PRACTICES 10
Conclusion 10
About NETSCOUT Omnis Smart Edge Monitoring Solution 11

Executive Summary

Organizations are striving to implement digital transformations at an unprecedented rate, which accelerated recently in response to changing business conditions brought on by the COVID-19 pandemic. According to IDC reports, during the January – March 2020 quarter, while non-cloud-based spending was down more than 16%, cloud infrastructure spending increased 2.2%. A survey by MariaDB revealed that 40% of enterprises reportedly were fast-tracking cloud-based initiatives in response to the pandemic. In order to realize the benefits of these innovations, organizations need to ensure application performance and user experience with these services are optimized – before and after migrations occur.

This paper will explore the challenges facing enterprise organizations engaged in high-profile, business- impacting, and end-user affecting migrations. It will provide recommendations on a smarter approach to gain end-thru-end visibility to monitor, measure, and remediate user experience and performance issues across complex, multi-cloud environments.

Trends in Cloud Migrations

With 90% of companies using some cloud and 58% of companies currently using a hybrid cloud strategy, it is no surprise that the global cloud computing market is expected to eclipse $330 billion in 2020 and grow to over $620 billion by 2023. In fact, according to hostingtribunal.com, approximately 33% of IT budgets are directed at cloud services. As of 2019, 60% of workloads were running in hosted cloud, which was up from 45% in 2018. It is further projected that as much as 94% of workloads will be processed in cloud data centers by 2021.

These staggering numbers, in such a high-growth market, highlights the significant activity in application workload migrations now and into the future. Operational efficiencies like automation, cost controls associated with the elasticity of cloud resources, and greater agility offered by cloud deployments are primary drivers for moving to the cloud. And, given the shift in so many employees still working from home, cost savings from only paying for what you use is also a driver.

It is no wonder that application migrations, either lift-and-shift or modernization / refactoring, are among the most common cloud initiatives. What hasn't changed is the expectation that the quality of end-user experience with those applications is anything less than equivalent or even better than its previous deployment model.

What are End-Users Expecting?

Businesses move applications to the cloud, according to recent trends, now more than ever. IT organizations, corporate lines of business, and senior executives all understand that one measure of the success of a migration will be their end-users' experience. So, what matters to those users, the customers, employees, partners, or even B2B systems? As they access any number of different applications – Web or e-commerce application, an internal process application, or a contact center – they want that experience to be, simply, flawless! Poor user experience may mean failed transactions for employees or a painfully poor voice call with customers. This impacts a company's reputation and/or their employee's productivity and, in some cases, the company's revenue and customers' loyalty.

Four key characteristics are associated with defining user experience. Availability, Reliability, Responsiveness, and Quality (Figure 1).

- **Availability** – Users must be able to connect to and use the application whenever they want. Availability is an assurance that remote employees and customers can access what they need and when they need it, regardless of where in the hybrid cloud it resides.
- **Reliability** – Once connected, the user can stay connected and perform tasks without interruption.
- **Responsiveness** – The application is not seen as slow or unresponsive to the user. Speed of replies when using applications and services determines how responsive it is to the user.
- **Quality** – Users should enjoy a high-quality experience, particularly when using real-time collaboration services. Collaboration platforms are essential to today's workforce and business, often being the "face" of a company. Proper visibility into key UC systems is required to understand UC experience issues.

Users expect the same quality performance in their cloud services regardless of whether they are accessing them from headquarters, branch locations, or home offices.

Figure 1: User experience depends on excellence in four key areas – Availability, Reliability, Responsiveness, and Quality.

All four characteristics MUST be tracked and measured to ensure end-user experience targets are achieved. This means consistent, constant monitoring of actual user activity as well as transaction testing even when the user is not using the service. Availability, reliability, responsiveness and quality all need to be analyzed and validated 7x24 to stay ahead of issues before end-users are impacted. Quick diagnose of the cause or source of an emerging issue enables swift restoration of services to optimal performance. Anything less will mean an organization risks missing revenue, injuring customer loyalty, damaging their reputation, failing to meet regulatory compliance, and / or lowering employee productivity in that area.

Cloud Migration Trends and Challenges

The cloud adoption cycle has moved from early-adopter to main-stream phase as the industry has matured over recent years. Continuous innovations in cloud technology, coupled with multiple vendors offering practical and affordable hybrid cloud solutions, has led many IT organizations to embark on major workload migration projects in the cloud. Enterprises are moving services to the cloud and infrastructure to Co-Los, innovating their private data centers, and adopting new Software-as-a-Service (SaaS) applications at record rates.

In fact, as cloud has matured, so too have the ways enterprises have migrated their services to the cloud. Initially, lift and shift migrations by organizations was a favored approach, for agility in the cloud and to take advantage of capex and opex savings. Cloud based tools were used to help evaluate and troubleshoot issues and tended to be very cloud dependent. Cloud migrations that involved re-factoring followed next and had enabled organizations to optimize time-to-value and improve their competitive advantages. Code-based and synthetic instrumentation provided focused analysis on the client and workload end-points to use when problems would emerge – end-to-end troubleshooting that was dependent on the application in use at the time.

Today, the maturity of the cloud and workload migrations has entered yet a third phase marked by operationalization of the services and enables organizations to maximize the value of their cloud investments. It is here that the performance problems have become complex and challenging, chief among them is the lack of visibility and control.

Assuring User Experience Before and After Cloud Migrations

Hybrid, multi-cloud environments involve multiple vendors, each delivering different pieces of a critical service, making researching problems somewhere in that communications path complicated and time consuming. Assuring end-user experience, with any application, through any multi-cloud environment is critical and operationalizing troubleshooting that is both cloud agnostic and application agnostic is essential.

Cloud migration challenges used to be infrastructure buildout, technical readiness, and affordability, but have evolved to be more about preserving user experience and quality of service, during and after the workload move. This is difficult as parts of the underlying infrastructure is now owned and operated by multiple IT domains and external third parties.

Troubleshooting issues impacting cloud migrations are complicated. War rooms of the past, where there was one owner, one General if you will, no longer exist. The responsibility is shared now among several owners and operators and these multiple forces can lose valuable time in finger pointing during triage and troubleshooting. There can be an infinite number of possible causes of a problem across these environments and a variety of ways the issues are reported or revealed. It is fair to ask if the right party will even be invited to the war room to evaluate the situation considering all the different possible participants. Meanwhile the issue and poor user experience persists.

Instrumenting the communication end points, that is, the client end point with synthetic agents and the application workload endpoint with APM agents, is one approach. However, it provides an incomplete picture focused only on the end point performance. There is a lot of activity occurring between these two points involving WANs, MPLS, SD-WAN, to name a few, where the issue can lie. End-point views are insufficient to solve today's complex problems and show the lack of visibility and control necessary to gain what is really needed. End-Through-End views, that reveal all the domains potentially impacting end-user experience, are required to show what is happening across the path and how that looks and feels to the user. And, ultimately this will help pinpoint where the issue affecting that user's experience lies.

Troubleshooting issues impacting cloud migrations are complicated by the virtually infinite number of potential causes throughout the multi-vendor environment.

End-Through-End Visibility and Control

There are three general areas for visibility (location types) in today's complex networks – Client, Network and Data Center / Cloud Service edges (Figure 2). Each of the three edges described share a common characteristic in that traffic is altered as it crosses domains between them. That means, as traffic moves between each IT domain, it is altered for the next domain. Examples include traffic from wireless to wired connections, or LAN to WAN, or ISP to Co-Lo, Co-Lo to cloud, cloud to server workload, etc. Each of these points of change is an edge that presents a potential gap in visibility.

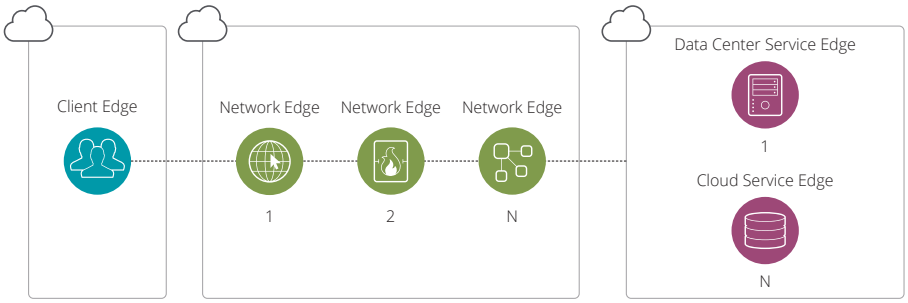

Figure 2: The three locations that exist across today's complex, multi-cloud environments are the Client, Network, and Data Center / Cloud Service Edges.

Client Edge

The Client Edge is where the user lives, also known as the client or group of clients, and it is also where user experience is felt. This is where they connect to the network using shared infrastructure to access particular services. Some examples include:

- Remote branch offices
- Home offices
- Contact centers
- Office floors at a company headquarters
- Manufacturing factories and plants
- Hospitals and medical buildings
- Warehouses and Distribution Centers

Network Edge

The Network Edge is defined as those boundaries between different IT domains or an interface between two network segments and could be thought of as the transport edge. Here, the packets may be modified in some way as they move through to the next service edge. These different domain points may or may not be owned or operated by different organizations.

Examples include:

- The connection of a remote site to the Internet, WAN, or Software-Defined WAN (SD-WAN)
- Peering locations in co-los, e.g., with public cloud, Internet
- VPN concentrators and/or VDI load-balancers for remote access
- Firewalls, DMZs, load-balancers for security
- Centralized SD-WAN concentrators at data centers, co-los (e.g. Equinix), and/or public cloud (e.g. Amazon AWS, Microsoft Azure, Google Cloud, Oracle Cloud)
- Private data center edges

Data Center / Cloud Service Edge

The Data Center / Cloud Service Edge is the first server that traffic from the client hits, wherever that application stack resides. This includes data centers, private or public cloud, and where communications between the different application tiers or microservices over east-west connections occurs.

Examples include:

- Data Center / Cloud Service Edge 1 will be the first server that traffic from the client will hit in the application stack.
- Additional servers that are between the different application tiers or microservices, including application servers, database servers.

With the number and variety of potential edges involved in any given application service, ffrom client to network to data center/cloud service edge, it is essential to have vendor independent visibility to quickly troubleshoot issues as they emerge to reduce mean-time-to-knowledge (MTTK) and mean-time-to-restore (MTTR). Using Network Edge 1 or 2 as a vantage point for monitoring wire data, for instance, will quickly show IT if the problem they are investigating is to the left on the client side or to the right towards the Network and Data Center / Cloud Service Edge. This information is powerful as it helps get critical evidence in the hands of the right IT group or third-party vendor to aid in rectifying the disruptive issue.

Cloud migration challenges used to be infrastructure buildout, technical readiness, and affordability, but have evolved to be more about preserving user experience and quality of service, during and after the workload move.

How To Optimize Visibility and Control

The Edges identified here offer unique visibility vantage points for ongoing monitoring and troubleshooting end-user experience and application performance. Strategic instrumentation of some, or all, of these vantage points returns visibility and control to IT Operations, enabling the realization of the intended benefits of digital transformations by improving its performance and meeting or accelerating its implementation.

The combined use of both packet data and synthetic test technologies for monitoring at these different edges is an imperative.

Packet-Based Monitoring

Passive monitoring leverages visibility into the actual packets that traverse the enterprise environments, seen as wire-data on north-south physical Ethernet segments, as well as across virtualized environments over east-west traffic paths. Monitoring actual user packets, in a passive, non-intrusive manner provides consistent, real-time visibility that is essential to gain precise metrics on actual activity of users as they interact with application services in their data centers, co-los, public cloud, and SaaS providers.

With sophisticated analysis, packet-based monitoring enables vendor-independent metrics and trending analysis throughout the path for use in capacity and bandwidth planning, identification and troubleshooting of emerging and existing problems, and details on users and activity enterprise-wide so IT organizations can make informed, evidence-based decisions. Key to the effectiveness of passive packet-based monitoring is:

- Scalability across all types of infrastructures, technologies, and speeds, regardless of network size or global scope to provide vendor and technology agnostic visibility.
- Fidelity to provide real-time, in-depth metrics related to utilization, performance, errors, availability, users, etc., across virtually any application in the environment – voice, video, and business data applications and protocols – for analysis across complex application stacks.
- Multi-purpose to offer broadest support for multiple enterprise use cases including triage, troubleshooting, capacity analysis, migration planning, user and application trending, and overall health, delivering a full range of baseline metrics, performance analysis, deviation alarming, and configurable reporting, to name a few.

Synthetic Testing and Monitoring

Synthetic or active monitoring simulates end-user actions with applications across an enterprise environment by providing scheduled, consistent tests, typically from a client perspective, that delivers detailed performance and availability analysis of targeted applications and/or network areas. Tests that are run automatically from remote user locations, at set time intervals, even when users are not active, provide a rich set of monitoring metrics and trended information that can reveal subtle changes in many areas. These could include SaaS applications, Unified Communications-as-a-Service (UCaaS) solutions, and even the impact of Wi-Fi or local Ethernet networks at remote locations that may have on overall impact on transaction performance.

Key to the usefulness of synthetic testing and monitoring is:

- Flexibility in evaluating individual experiences under a variety of different conditions including locations and technologies such as home offices, branch locations, or manufacturing plants, over Wi-Fi or Ethernet, as well as with private, public cloud, or SaaS environments.
- Customizable tests designed to analyze and trend a range of user actions including navigation to a URL, such as, company.com, login to a URL or site, more detailed steps or actions, and logout.
- Multi-purpose to cover testing and analysis of both business applications as well as unified communications services using the same synthetic testing technology which greatly increases its value and usefulness in understanding end-user experiences.

The combination of real-time passive packet monitoring with continuous active testing provides the most comprehensive visibility for high quality application performance and user experience.

Superior Visibility Offered with Combination of Packet and Synthetic Monitoring

Ultimately, the combination of passive, packet-based monitoring and active, synthetic testing is a more comprehensive way to optimize user experience as well as assure network and application performance in today's complex, multi-cloud environments. As Figure 3 illustrates, when problems are experienced and reported by users, complexity will impact both ownership and speed of resolution. Time could possibly be lost as the different teams either blame each other as the source of the problem or as they use point tools in an attempt to prove it is not their area of responsibility causing the issue.

Synthetic agents from the remote perspective (Figure 3), testing the workflows are very effective at executing the same transaction repeatedly, downloading a web page, logging into Office 365, or voice call test, for instance, and will detect and alert if there is a change. Not unlike a smoke alarm, it will let you know a problem is emerging. Knowing a slowdown may be occurring, enables IT to turn to their packet monitoring solution to evaluate where in the path the slowdown is developing. Evidence collected with this analysis can be shared among the IT team, co-lo, WAN provider, and/or cloud vendor for resolution.

The process of using both packets and synthetic tests is highly effective, avoiding and reducing overall downtime experienced by customers and users. By cutting MTTK and reducing MTTR, user impact is minimized and user experience optimized, while simultaneously improving productivity of the IT staff and overall employee community.

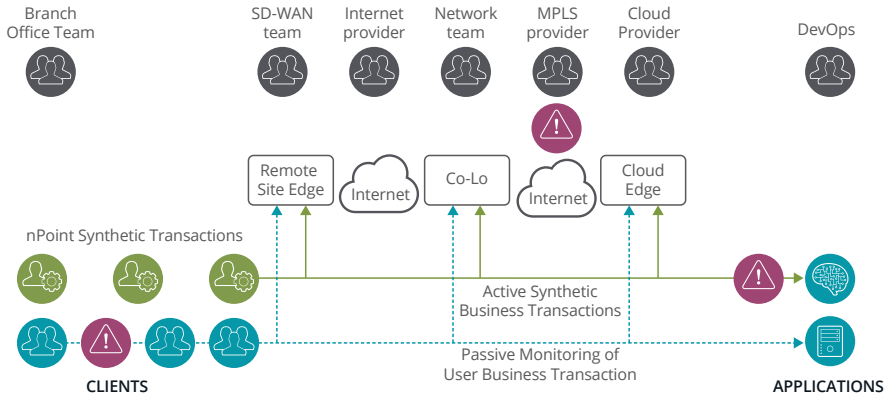

Figure 3: When employees in a branch office encounter poor user experience with a corporate application, rapid triage and isolation are challenging for IT without comprehensive visibility. The diagram shows the complexity involved from the client edge at the branch office with SD-WAN, through the four network edges (Internet, Co-Lo, MPLS, and Cloud provider), to the point where it meets the data center / cloud service edge of the application servers. Synthetic tests and packet-based monitoring capabilities help reduce MTTK and MTTR.

USE CASE / VISIBILITY BEST PRACTICES

Mitigating Application Migration Challenges

BUSINESS PROBLEM
- Employees are unable to process orders efficiently in applucation recently migrated to cloud
- Businesses loses customers to competitors
- Reputational damage to businesses

WITHOUT VISIBILITY
- Comlex path with multipleowners
- IT does not know where to start looking
- Lost time with finger pointing between teams/vendors
- Problems remain unresolveed

PASSIVE & ACTOVE MONITORING
- Visibility isolates cause of high latency
- Active, synthetic testing confirms that business processes are slow to execute
- Passive, packet analysis shows appication design is unsuitable for high-latency links

OUTCOME
- Application temporarily repatriated to origial location
- Business process efficiency restored rapidly
- Application redesigned so it can be safely moved back to public cloud

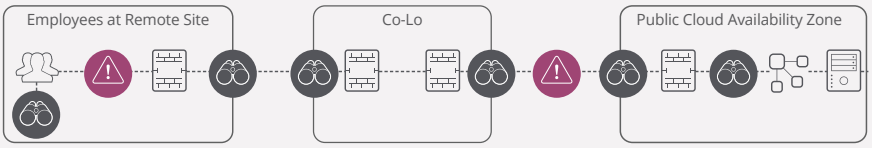

A group of employees at a remote site was experiencing slowdowns processing orders. Without proper visibility and control as this application was migrated to the cloud, this problem could have perpetuated for a protracted period of time, putting customer service, revenue, and reputation at significant risk. However, this enterprise deployed synthetic testing for their recently migrated order processing application from the branch office client edge.

They also implemented packet-based monitoring at the network edges and at the services edge of the public cloud availability zone for visibility into the application performance network-wide. The combination of the analysis from the synthetic testing alerted to latency issues and the packet-based analysis showed that this application design was unsuitable for high latency WAN links and would need to be redesigned.

IT repatriated the application temporarily to its original host and performance returned to normal It was the combination of the two visibility technologies – synthetic testing and packet monitoring that truly isolated the source of the problem with this migrated application, and allowed them to develop an action plan that would meet their business goals to move to the cloud and maintain high quality for end users with the order processing application.

Conclusion

In a world of complexity, this could not be simpler. The combination of comprehensive, scalable, passive packet monitoring with scheduled, continuous, active synthetic testing is necessary to provide the most comprehensive visibility to ensure availability, reliability, responsiveness, and quality for excellence in end-user experience across the client, network, and data center / cloud service edges in today's intricate multi-cloud environments.

Adoption of such broad visibility protects critical business drivers including reliability of revenue-impacting services, availability of customer-facing applications and portals, responsiveness of business applications for employees job responsibilities, quality voice/video communications with corporate representatives and contact centers, all of which protects corporate responsibility for regulatory compliance and reputational standing. With such critical business requirements dependent on performance excellence of application services, implementing a combined solution for passive and active monitoring is simply good business practice!

About NETSCOUT Omnis Smart Edge Monitoring Solution

Leveraging its 30 years of market and technology leadership in patented smart data technology and smart analytics, NETSCOUT® is revolutionizing the industry with the introduction of the Omnis™ Smart Edge Monitoring solution that provides comprehensive, borderless monitoring and visibility unparalleled in the industry today. NETSCOUT's Omnis Smart Edge Monitoring uniquely combines and extends the benefits of our market leading monitoring solution, nGeniusONE® Service Assurance platform, with our award-winning synthetic testing offering, nGenius®PULSE, to provide critical insights into end-user experience.

The anchor of the Omnis Smart Edge Monitoring solution is the NETSCOUT InfiniStreamNG® (ISNG) appliance with Omnis™ Cloud Adaptor which uses NETSCOUT's patented, advanced Adaptive Service Intelligence® (ASI) technology to combine passive, packet-based monitoring and synthetic testing intelligence in a single data source. Building upon our Visibility Without Borders capability, NETSCOUT provides 7x24 testing and monitoring of critical business services – even if no one is actually using them. nGeniusPULSE nPoint synthetic test technology, including business transaction tests (BTT), continuously performs customer designed tests on applications and services. Our ASI technology combines the test traffic with packet monitoring traffic to provide an unmatched perspective of the service delivery and user experience as analyzed and viewed in nGeniusONE.

NETSCOUT is elevating the way IT organizations have been approaching application service delivery, performance, service edge monitoring, and end-user experience assurance in their ever-evolving, multi-cloud environments. Troubleshooting performance issues is more proactive and efficient as Omnis Smart Edge Monitoring provides insight into the entire path between the client and the application, thereby enabling the source of a problem to be rapidly identified. For problems that originate with third-party providers, Omnis Smart Edge Monitoring also allows evidence to be collected which can help the third-party providers confirm the nature of the issue and resolve it faster.

NETSCOUT's transformative approach dramatically changes the way organizations plan, deliver, integrate, test, and deploy services and applications in the cloud, accelerating their digital transformations. Application migrations, as described in the example of this white paper, will be more timely, successful, and seamless with the NETSCOUT Omnis Smart Edge Monitoring solution in action. With detailed evidence, Omnis Smart Edge provides the pinpoint accuracy to evaluate performance at the client, network, and data center / cloud service edge, as well as to measure user experience. This allows you to respond and triage the issues impacting the success of your cloud migration like never before.

Today, with your teams and your applications on the move, the NETSCOUT Omnis Smart Edge Monitoring solution removes the gaps in visibility for service edge monitoring and end-through-end insight into application service responsiveness, availability and reliability for high quality user experience, bringing control back to your IT operations teams.

NETSCOUT

Corporate Headquarters
NETSCOUT Systems, Inc.
Westford, MA 01886-4105
Phone: +1 978-614-4000
www.netscout.com

Sales Information
Toll Free US: 800-309-4804
(International numbers below)

Product Support
Toll Free US: 888-357-7667
(International numbers below)

NETSCOUT offers sales, support, and services in over 32 countries. Global addresses, and international numbers are listed on the NETSCOUT website at: www.netscout.com/company/contact-us

EWP_027_EN-2101 05/2021