



How To Remove IT Visibility Gaps and Ensure Successful Public Cloud Migrations

TABLE OF CONTENTS

Executive Summary	3
Introduction	4
Use Cases, Deployment Strategies, and Value	5
Gap Analysis: Self-Audit	7
Case Study	8
Conclusion	8

Executive Summary

Migrating workloads to the public cloud has grown to be a key driver behind digital transformation initiatives and an essential building block to amplify business processes. Yet initiating a cloud migration without sufficient monitoring to ensure application performance and availability before, during, and after can create visibility gaps that cause IT organizations to lose control of their service environment performance, especially if workloads are spread over multiple on-premises or additional public cloud domains. To succeed in competitive markets and maximize their public cloud investments, organizations also need comprehensive visibility into their network environments to enable faster, more efficient remediation of issues and, ultimately, ensure the success of existing deployments and ongoing migrations.

Introduction

Public cloud usage has experienced almost a decade of more than 30 percent annual growth, according to Forrester. And it is accelerating. The public cloud market is expected to reach \$1 trillion worldwide by 2026 even amid budget-tightening concerns. Migration to public cloud services (such as AWS, Azure, and Google Cloud) enables organizations to leverage quick, scalable models to deploy compute and services for mission-critical applications. But lack of visibility and control in public cloud environments are often listed among the top challenges to successful cloud migration.

Visibility means understanding everything that is happening on the network and the traffic moving through it in detail to identify and manage performance degradations that are otherwise hard to detect. Public cloud environments are, by design, constantly changing and difficult to monitor, due in large part to workloads often being splintered across data centers, colocation facilities, and physical and virtualized infrastructures.

In this whitepaper, we explore challenges associated with public cloud migration and IT visibility gaps caused by dependencies before (Day 0), during (Day 1), and after (Day 2) so that organizations can maintain application performance and service availability. We also highlight the key features and capabilities of packet-based solutions from NETSCOUT® to greatly reduce mean time to repair (MTTR) by accurately identifying, triaging, and isolating problems affecting application performance so companies can overcome unexpected visibility challenges.

Challenges

As public cloud services are typically deployed in 'Regions' and 'Availability Zones' close to where users are located, they can have the effect of expanding IT's range of responsibility beyond what was once a small number of centralized data centers. With deployment now spread across multiple worldwide availability zones for resiliency/redundancy, IT must ensure cloud customers have a stable connection to a cloud service in the geographic area that is closest to them.

Users either access resources in the public cloud over the internet or through a colocation data center facility. Colos provide direct peering connections with public cloud providers, for example via AWS Direct Connect or Azure ExpressRoute, rather than through the internet, to guarantee low latency and predictable bandwidth capacity. As these are critical, high-traffic transition locations in the network, they risk becoming potential points of failure without essential visibility.

The typical deployments NETSCOUT sees are not single cloud only, but a mix of data center, colo, and often one or more public clouds—which can complicate migration plans and visibility strategies and create service interdependencies. Not every application is cloud-ready or built to be completely cloud-native. Without sufficient multi-domain monitoring, management of distributed applications can cause unanticipated IT visibility gaps that may affect service performance. Visibility becomes a priority throughout the three critical phases of migration in the following ways:

- **Day 0 - Before the migration:** Visibility of the performance and the dependencies of applications while they are still on-premises or hosted on a public cloud provider's server to establish a performance baseline and to clarify what needs to move so nothing is left behind.
- **Day 1 - During the migration:** Instrumentation in the cloud can identify and help to resolve performance and availability problems to ensure a rapid and successful migration.
- **Day 2 - After the migration:** Continuously monitoring and optimizing the performance of applications, especially since they frequently span multiple on-premises and cloud domains.

The need for visibility does not stop once an application has been successfully migrated. Future changes in the environments can create additional IT visibility gaps and unexpected disruptions. One of the biggest challenges facing enterprises should not be “why,” but rather “how to,” give evolving virtualized environments packet-level visibility to quickly identify and reduce the mean time to repair (MTTR) of emerging service issues or failures that affect mission-critical applications and the users and customers who depend on them.

nGeniusONE®, the core of NETSCOUT’s nGenius Enterprise Performance Management solutions, provides a unified, real-time view of application performance data from different traffic sources in a single pane of glass, so teams can quickly troubleshoot and resolve migration issues, maintain operational efficiency, and safeguard reputation and public cloud services, regardless of where workloads reside (See Figure 1 and Figure 2).

Visibility to minimize MTTR of mission critical application issues

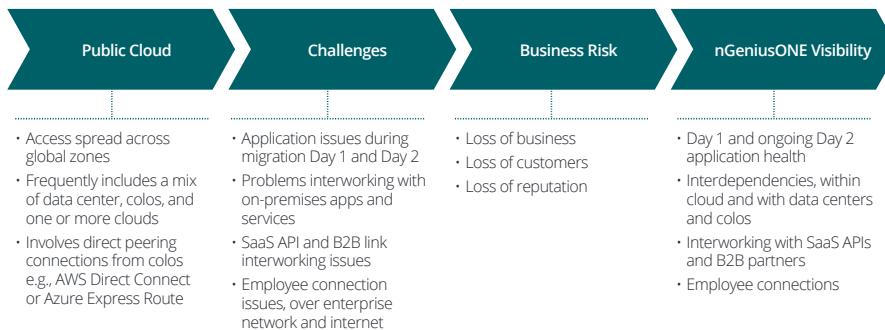


Figure 1: Instrumenting in the Public Cloud Overview

Use Cases, Deployment Strategies, and Value

Public cloud can reduce the operational burden of running and maintaining hardware and software. But virtualized environments present unique visibility challenges, often requiring robust metadata presented in an easy-to-view format, such as nGeniusONE, fed by highly accurate solutions that gather reliable packet-level data from virtualized and cloud infrastructures, such as NETSCOUT’s vSTREAM® virtual appliance.

The instrumentation strategy for public cloud often follows an “outside-in” approach, whereby locations carrying aggregate traffic are instrumented first. If enterprise IT teams are leveraging connections from a colocation data center, for example, it is worth considering what visibility can be obtained by instrumenting these connections in the colo.

Instrumentation in the public cloud may start with aggregation locations, such as inspection zones that include traffic flowing to/from the internet and between virtual private clouds (VPCs) in the public cloud availability zones. East-west traffic flowing between individual workloads can be monitored using cloud-native port mirror capabilities (if available) or a vSTREAM agent (see Table 1 and Figure 2 for instrumentation examples).

NETSCOUT leverages different packet acquisition strategies to ensure mission-critical services are performing according to baseline standards depending on the public cloud in use. For examples:

- **Amazon Web Services and Google Cloud Platform:** These vendors provide native mirroring capability of workloads (VPC traffic mirroring in AWS, packet mirroring in GCP). The mirrored traffic is sent to a target or set of targets (in our case, directly to vSTREAM monitoring interfaces or through a network load balancer front-ending a set of vSTREAM interfaces).
- **Azure:** NETSCOUT has developed the nGenius Cloud vTAP product, which provides a resilient inline tapping capability that allows east-west traffic between subnets to be mirrored to vSTREAM interfaces.
- **Where native packet-level visibility is not provided by the cloud provider:** Organizations may use the vSTREAM agent, which is installed into the workloads themselves.

Cloud Provider	Traffic Challenge	nGeniusONE Visibility Gap Analysis
AWS	Internet connected users, external services, connected colo users to/from services hosted in the public cloud	Traffic is rerouted via ingress routing to a Gateway Load Balancer (GWLB) to virtual appliances (such as 3rd party firewalls/IDS) which can be mirrored with virtual private network (VPN) traffic mirroring (1)
AWS	Workload to workload traffic, or transit connections between VPCs	Traffic is rerouted via AWS Transit Gateway (TGW) to GWLB to virtual appliances (such as 3rd party firewalls/IDS) which can be mirrored with VPC traffic mirroring (1)
AWS	Traffic to/from individual workloads	Individual workloads can be mirrored (3)
GCP	Traffic to/from groups of workloads	Subnet mirroring (2)
GCP	Traffic to/from individual workloads	Individual workloads can be mirrored by IP address (3)
Azure	Traffic to/from individual or groups of workloads	Traffic between Azure virtual network (VNet) subnets can be rerouted using user-defined routes to an Azure load balancer hosting NETSCOUT nGenius Cloud vTAPS, which will mirror traffic to vSTREAM (2)
AWS/GCP/Azure	Traffic to individual or groups of workloads	Specific workloads can be instrumented by installing a vSTREAM agent (4)

Table 1: Public Cloud Visibility Gaps and Instrumentation Locations.

Table 1 lists the IT visibility gaps that may be present in public cloud deployments along with the instrumentation locations according to the type of challenge and nGeniusONE detection. The numbers in brackets represent placement locations for selected nGeniusONE collection devices and are shown in Figure 2.

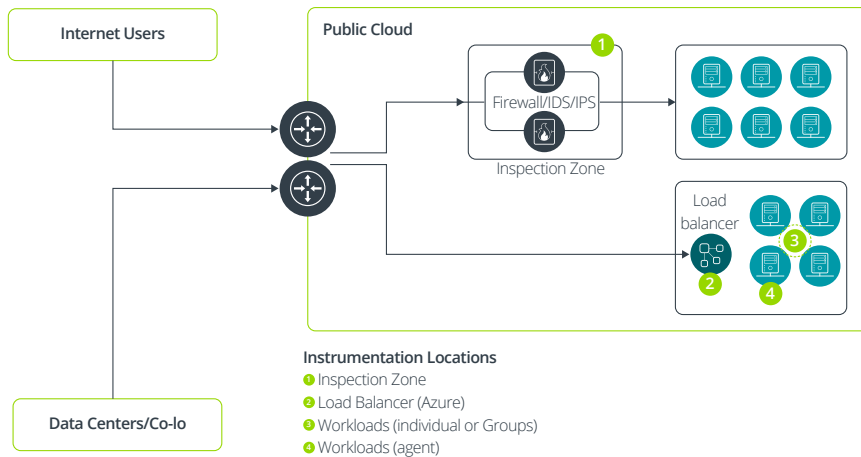


Figure 2: Instrumentation Locations.

vSTREAM is ideal for monitoring service-critical traffic running within virtualized or cloud infrastructures. Together with nGeniusONE, the two solutions provide deep insight into packet traffic and application workloads giving IT organizations end-through-end monitoring and control of IT visibility gaps in public cloud infrastructure regardless of whether applications are hosted on-premises, with a third-party provider, or if services are split between the two enterprise IT environments.

Decisions relating to the number of vSTREAMs will depend on factors such as the quantity of traffic, whether the number of monitoring locations are expected to increase over time, and cloud provider packet acquisition costs, which may vary depending on whether mirrored traffic traverses VPCs, Availability Zones, and other routing technologies.

Our solutions have been independently proven to reduce the MTTR of incidents by 80%.

Gap Analysis: Self-Audit

Every public cloud migration is different and guaranteed to be complex. Companies pursuing a migration to harness public cloud benefits—such as new revenue streams, remote work flexibility, data and infrastructure services, offering on-demand applications, improving customer and user experiences, etc—should conduct an audit of the business benefits and weigh those against critical needs to determine best practices. Here are some common questions to consider:

1. Which applications need to be migrated first?
2. What is the best migration approach for other applications?
3. Are applications performing according to the baseline recorded in Day 0?
4. Have IT visibility gaps emerged as applications moved to public cloud infrastructure since Day 1?
5. Has the migration affected your ability to accurately pinpoint the root cause of service degradations after Day 2?
6. What affect has the migration had on your team's MTTR?
7. Can you see everything that is happening in your on-premises, cloud, hybrid, and multicloud settings with your existing monitoring solution? Does it combine end-user experience, application, network, and cloud performance into a single view?
8. Is your team able to accurately map application interdependencies to the underlying infrastructure?
9. Would it be more cost-effective to slightly increase instrumentation investments as opposed to adding specialized team members?

Case Study

NETSCOUT has built its legacy by helping the world's leading organizations overcome their visibility challenges, improve uptime and operational efficiency, and fuel innovation and growth. In this example, a multi-billion-dollar business process outsourcing (BPO) company and trusted consultant to many businesses needed help managing a critical migration involving Microsoft Azure, SD-WAN, and colos. Their IT service organization realized these transformations exposed IT visibility gaps that required proactive enterprise-wide visibility to ensure performance management and help them meet service-level agreements (SLAs).

Read the Case Study to learn how NETSCOUT [ensured the BPO company's business-critical services did not affect global internal customers.](#)

Conclusion

Public cloud migrations require organizations to deal with different variables, including likely scenarios and unexpected challenges. Regardless of where you are on your cloud migration journey, NETSCOUT can help at every stage. By monitoring your entire infrastructure (both on and offsite), our packet-based solutions can identify application interdependencies, distributed workloads, IT visibility gaps, and other challenges before, during, and after migration. This can help you improve MTTR and gain a contextual understanding of application availability and performance to ensure your move to the cloud is a success.

If you are noticing visibility gaps before, during, or after your public cloud migration, let NETSCOUT help you cover them.



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