DATASHEET

BROADCOM

Cross-Cloud Networking

Discover the best practices for network observability when using Google Cross-Cloud Networking.

For Google Cloud customers that have network or app infrastructure across multiple cloud environments, Google Cross-Cloud Networking is essential. Regardless of app ownership, network operations teams are responsible for troubleshooting network connections to business-critical services. With multicloud environments there is often less visibility into the network delivery paths between end users and applications.

To best assure multicloud services, it's crucial to baseline application and network performance in order to proactively validate the network paths between cloud providers. The use case described below is common for multicloud and hybrid cloud deployments, with active layers 3, 4, and 7 monitoring.

Observability Focus

For multicloud visibility the first step is to baseline and proactively validate the end-to-end network connection between the cloud endpoints. Network performance visibility is critical in order to isolate where the root cause of issues are between one cloud, the transit backbone and the Google cloud edge as well as the final application environment. Using AppNeta's continuous performance analysis to determine if the selected Google Cloud Standard or Premium tiers are appropriate for the customer applications and user base.

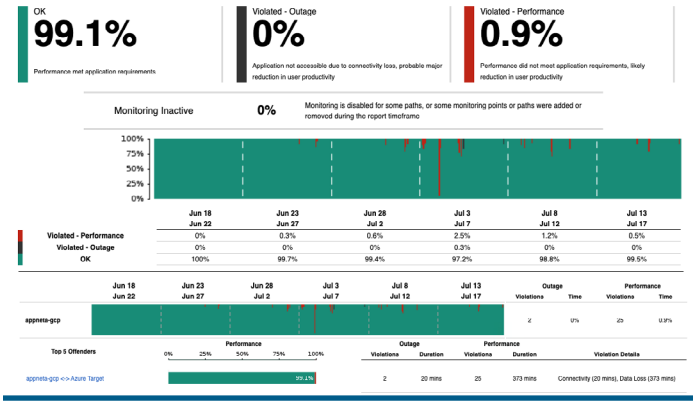

SLA Validation

For multicloud environments, visibility into the network mesh between Google Cloud, and other Tier 1 Cloud Providers like AWS, and Microsoft Azure is a common requirement.

This dashboard allows network operations to identify common or independent events between multiple cloud service providers. In this case, the SLA between Google Cloud and a target within Azure are within compliance.

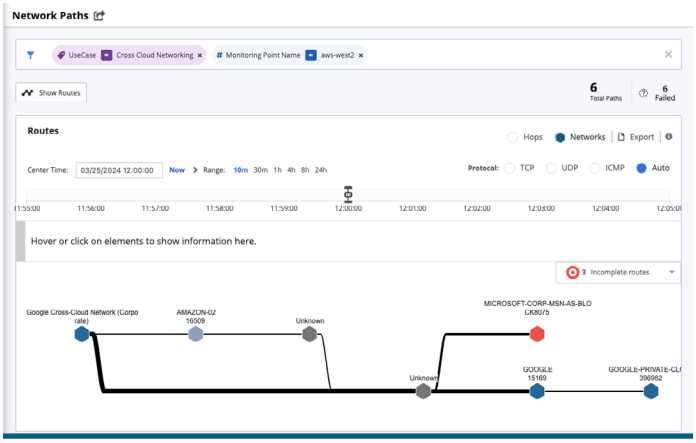

Multi-Route View to Cloud Apps

Visibility into the network delivery paths between cloud service providers and interconnect routes helps network operations identify potential outages like the Microsoft outage (red circle on the opposed diagram).

Through active visibility and continuous testing teams are able to identify the specific hop where the outage occurred and therefore assign resources or contact support depending on whether or not the issue is internal.

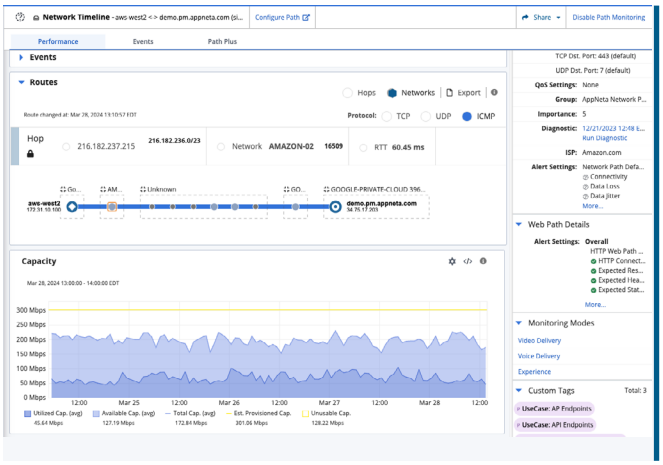

Continuous Performance Data

In this example on March 25th there was an event with high Data Jitter which correlated with Data Loss, high Latency, and high Round Trip Time (RTT). This event exceeded the acceptable level of Loss, which automatically started a Diagnostic Test.

Through the diagnostics network operations teams can get hop-by-hop network performance metrics leading to a discovery of Data Loss in the AWS domain of an outbound path to a Google-hosted application.

User Experience for Hosted Applications

To identify if there are any application performance impacts, network operations teams can continue on to the Layer 7 Experience monitoring. A significant spike in the maximum response time can be identified at the same time as the Data Loss event along with high Latency and RTT.

Teams need to be able to identify if the network event impacted the overall user experience of the application. By looking at both network and application metrics it's possible to do this quickly and efficiently.

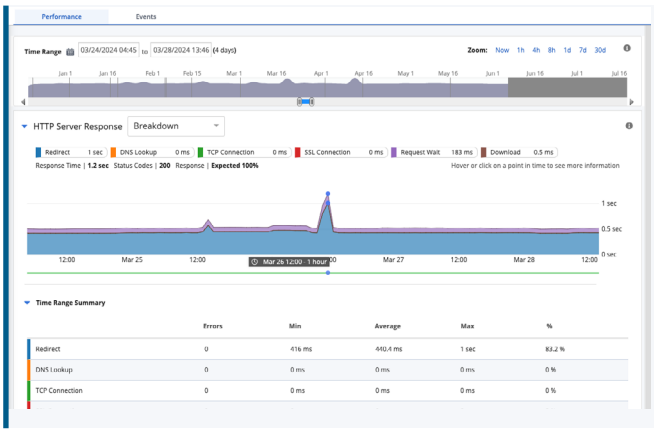

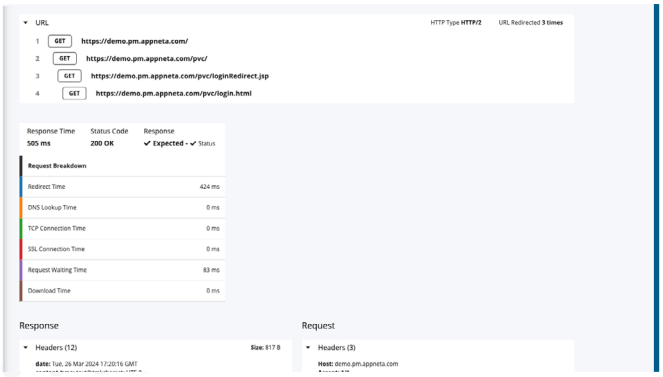

Application Drilldown

If the application is at fault, identifying what is wrong is simple by just selecting a timeframe to review and then viewing the Response Time, Response Headers, and page Body contained during the event.

With complete visibility via Web Synthetic Transaction Monitoring, teams don't have to guess what happened hours or days in the past because it's all captured at the time of the event.

Root Cause

In the case depicted, a network hop within the AWS domain was forwarding traffic to the application hosted in Google Cloud when the path experienced Data Loss. With active testing, network operations teams are able to pinpoint the root cause of the issue and who is responsible for fixing it. This is a very common use case for hybrid and multicloud environments as many enterprises use multiple cloud service providers to deliver business-critical applications to users across the globe.

BROADCOM
connecting everything®

For more information, visit our website at: www.broadcom.com

Cross-Cloud-Networking-Data-Sheet_BC-NO-2024_CE-4544_v2 July 22, 2024