

EBOOK

# IS YOUR NETWORK REALLY READY FOR AI?

10 Ways to Future  
Proof Today's Network  
Infrastructure for  
Tomorrow's AI Workloads

# The State of Network Operations

**77%**

report that network operations is significantly more complex today.

**82%**

report network visibility blind spots because of ISP, cloud and SaaS adoption.

**41%**


site a lack of skilled personnel as the #1 challenge for NetOps success.

Source: Cloud and Internet Usage Generates Network Observability Blind Spots - <https://academy.broadcom.com/network-observability/cloud-and-internet-usage-generates-blind-spots?interests=dx-netops,appneta>

# AI's Demands on Today's Networks


Characteristics	Traditional Network Traffic in Data Centers	Generative AI Network Traffic
<b>Data volume</b>	Relatively smaller data volumes	Massive datasets for training and inferencing
<b>Traffic pattern</b>	Client/server, north-south	Distributed, east-west traffic within clusters, GPU-to-GPU
<b>Bandwidth requirement</b>	Moderate to high bandwidth requirements	Extremely high bandwidth requirements for data transfer
<b>Latency sensitivity</b>	Moderate latency tolerance	Extremely low latency requirements for real-time inferencing
<b>Compute intensity</b>	Moderate compute demands	Highly compute intensive workloads (e.g., training)
<b>Network resilience</b>	Best-effort delivery	Lossless, reliable data transfer required
<b>Traffic burstiness</b>	Relatively steady traffic flows	Bursty traffic patterns during training/inferencing
<b>Network topology</b>	Hierarchical, tree-like topologies	Flat, high-radix topologies for low-diameter networks
<b>Data movement</b>	Storage to compute	Bidirectional GPU-to-GPU, GPU-to-Storage
<b>Network protocol</b>	TCP/IP, Ethernet	RDMA, InfiniBand, RoCE
<b>Traffic monitoring</b>	Basic monitoring	Advanced telemetry and analytics for AI workloads
<b>Network congestion</b>	TCP incast/outcast	Potential RDMA, congestion

# Our Networks Are Not Ready for AI



“We have restored wireless service to all our affected customers,” AT&T said on its [website](#) Thursday afternoon. “**We sincerely apologize to them ...** we are taking steps to ensure our customers do not experience this again in the future.”

Source: MSN



“**Network I/O** is key for recommendation workloads.”

Source: OCP keynote by Alexis Bjorlinat, 2022 OCP Global Summit





# 10 Ways to Future Proof Today's Network Infrastructure for Tomorrow's AI Workloads

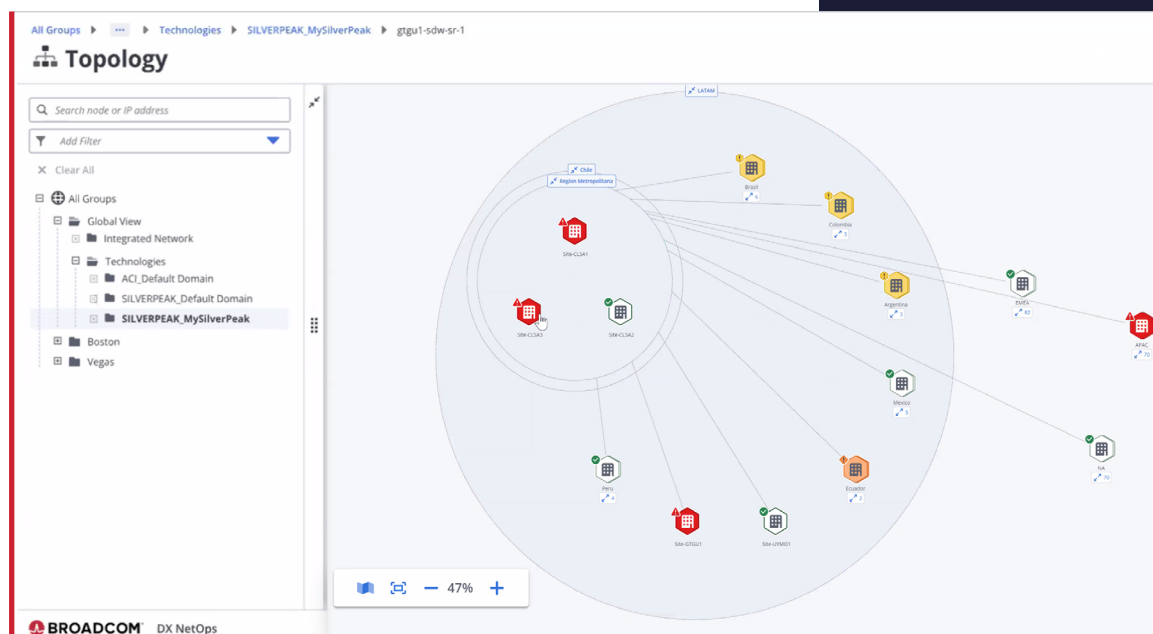
# Fault Isolation and Alert Suppression

The first rule of any network operations team is to reduce the troubleshooting noise generated by modern network technologies when managing user experiences.

Leading network observability solutions should offer the following proactive, multi-vendor visibility for success:

- Device/component status, relationships and dependencies
- Fault domain analysis and suppression
- Large-scale event correlation

Employing analytics that suppress the noise from downstream infrastructure events and alarms is a proven method to simplify troubleshooting and accelerate root cause analysis, particularly when AI workloads are disrupted.



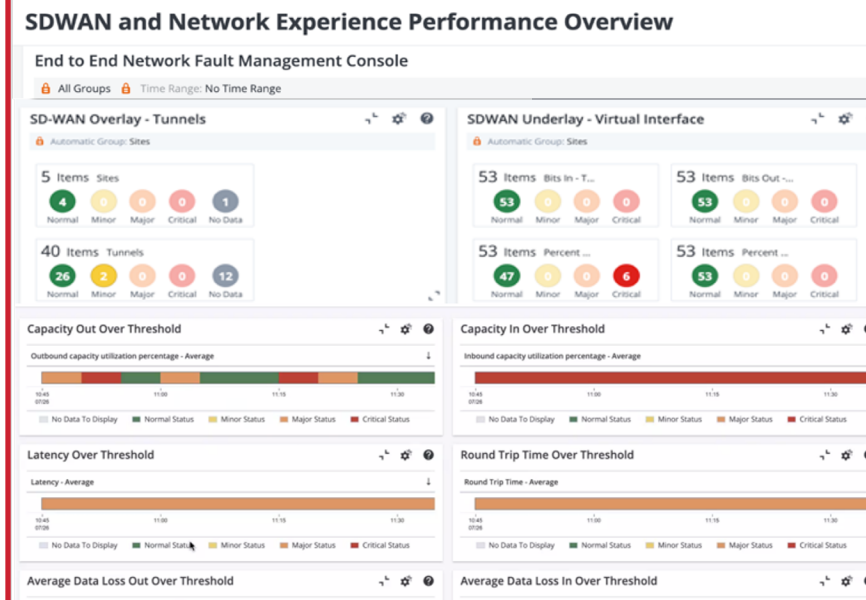
# Alarm Noise Reduction

Modern network technologies, such as Cisco Application Centric Infrastructure (ACI), can generate a high volume of unnecessary alarms and events, often complicating and delaying troubleshooting.

When selecting a network observability solution, look for expertise in the following areas:

- Proven analytical capabilities back by customer testimonials
- Tailored troubleshooting workflows following standard operating procedures
- A network-agnostic, flexible and extensible platform

When it comes to isolating the root cause of issues in modern networks, AI and advanced analytics undoubtedly hold promise—but this promise will only be realized when advanced capabilities are effectively tailored to the networking domain. When done successfully, these features can significantly enhance operational efficiency, speed up remediation, and support network team effectiveness.



**SDWAN and Network Experience Performance Overview**

End to End Network Fault Management Console

All Groups Time Range: No Time Range

no symptoms no maintenance Quick Filter Acknowledge Unacknowledge Clear Troubleshoot

Severity	Date/Time	Item Name	Alarm Title	Impact
Critical	Feb 20, 2024 16:58:...	oaimrb01.us-east1-...	DEVICE HAS STOPPED RESPONDING TO POLLS	1
Critical	Oct 10, 2023 16:30:0...	hq_rtr_01	DEVICE HAS STOPPED RESPONDING TO POLLS	140 3
Major	Jul 11, 2024 14:23:0...	Fa2/0 - FastEthernet...	%CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered on FastEthernet2/0 (not half duplex), with R1.mydomain.com FastEthernet0/1 (half duplex)	0 7564

Over 7k events, but only 1 alarm

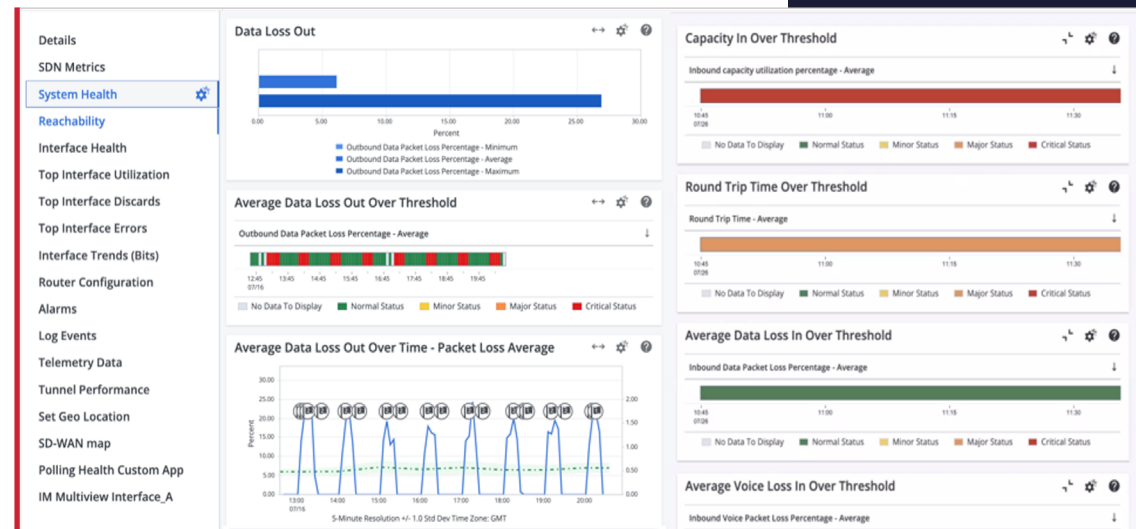
# Root Cause Analysis (Powered by Contextual Data)

Building on the previous steps, applying analytics to the right data means you can expect intelligent answers.

Data collected become correlated troubleshooting “breadcrumbs”. Look for solutions that offer granular visibility into:

- Faults, performance, flows, configs, logs
- Capacity, packet loss, latency, jitter
- SD-WAN underlay/overlay correlation
- Telemetry data
- Application experience impact analysis

When implemented correctly, expect minimal troubleshooting efforts or even a ‘one-click alarm to root cause diagnosis’.

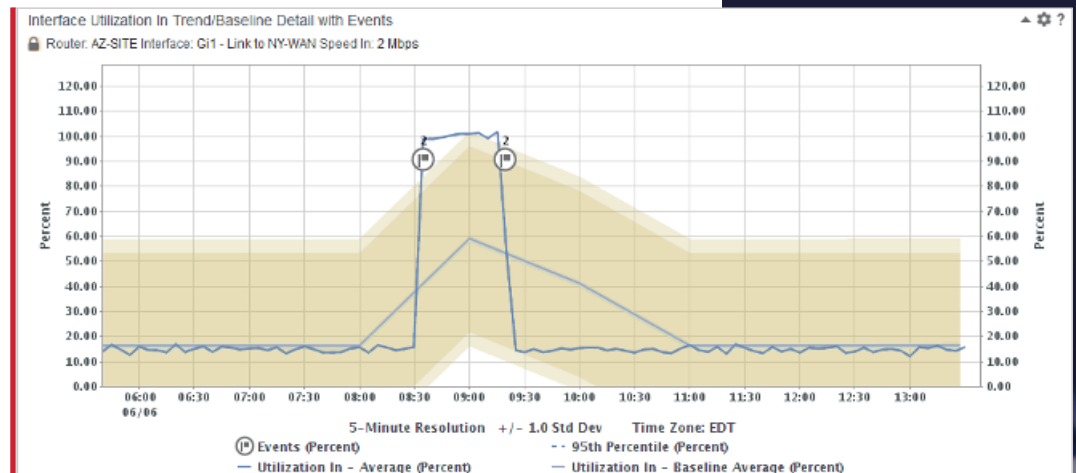




# Baseline Thresholding → Predictive Intelligence

Calculating future network performance based on historical metric data is essential for any modern network or AI initiative. Metric projections are valuable for capacity planning, such as ensuring that the interface bandwidth on a device—whether owned by your organization or an ISP—will be sufficient for a new application launch or network technology deployment. These projections can also help cut costs, protect revenue, and safeguard the company's brand.

When planning AI projects, these proactive measures should not be overlooked, especially since many enterprises implementing artificial intelligence for enhanced customer experiences may anticipate seasonal workload peaks.



BW Util Out - Interface - Utilization Out - 95 Percentile

User Group: ForwardInc

Timeframe: Last 30 Days Metric Calculate Level: by Component Resolution: Daily Roll-up

Critical Status (80) Major Status (60) Minor Status (40) Normal Status

Quick Filter

Group/Item Name	Device Name	Sep 13, 2024	Overall Average (%)	Projection: Dec 12, 2024	Projection: Jan 11, 2025	Projection: Feb 10, 2025
ForwardInc	ForwardInc	0.42%	0.42%	0.42%	0.42%	0.42%
Fa0/0 - Interface to PE	CE2.mydomain.com	96.26%	96.35%	96.5%	96.57%	96.64%
Fa0/1 - Interface to Loca	CE2.mydomain.com	0.42%	0.42%	0.42%	0.42%	0.42%
Fa0/1.100 - FastEthernet	CE2.mydomain.com	0.42%	0.42%	0.42%	0.42%	0.42%
Fa0/0 - Interface to PE	CE1.mydomain.com	0.27%	0.27%	0.27%	0.27%	0.27%
Ethernet1 - Ethernet1	Switch2.mydomain.com	0.12%	0.12%	0.12%	0.12%	0.12%
Ethernet2 - Ethernet2	Switch3.mydomain.com	0.12%	0.12%	0.12%	0.12%	0.12%
Fa0/1 - Interface to Loca	CE1.mydomain.com	0.03%	0.04%	0.04%	0.04%	0.04%

# #5

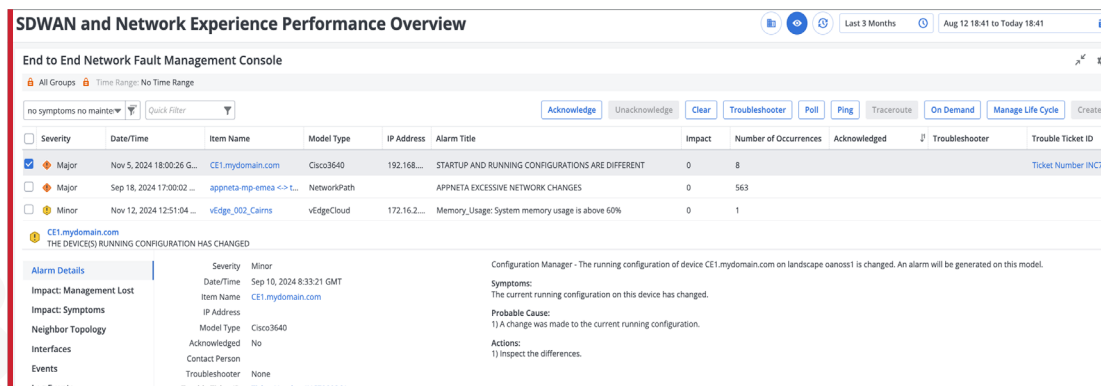
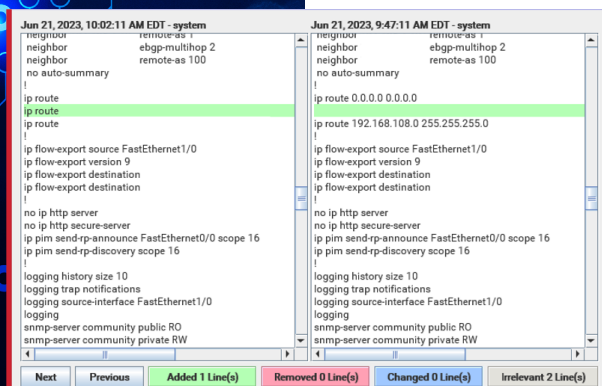
## Configuration Compliance and Automated Roll Back

Unplanned or unauthorized network device configuration changes are among the leading causes of network outages. Although most enterprises have compliance policies in place to mitigate this risk, many still face this longstanding challenge.

Solutions that offer a centralized portal for configuration analysis of any network device can expect the following benefits:

- Improved network security
- Better network stability
- Regulatory compliance
- Automated roll back
- Audit trails and reports
- Visibility, accountability

By correlating outages with configuration changes, you can expect to minimize the impact of unexpected downtime on operations, customers, or AI projects, and even prevent issues before they occur.



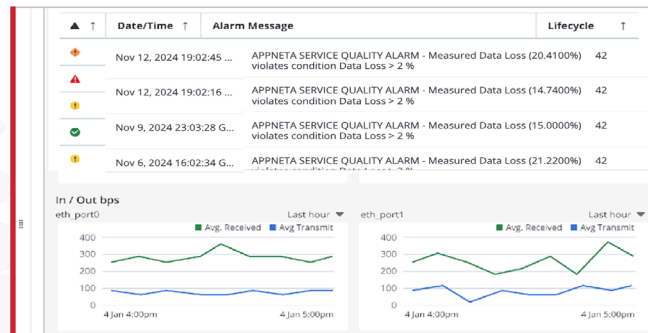
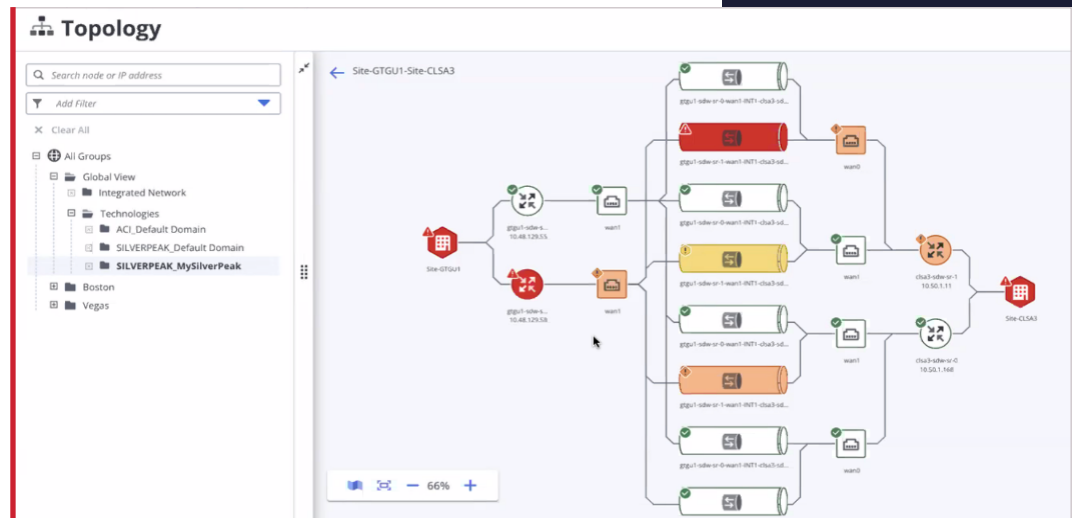
# Topology Analytics

Visualizing the entire global network landscape is always beneficial to operations teams; even when it's a mix of legacy and modern multi-vendor technologies, ISP and cloud network paths, and hybrid Wi-Fi setups.

With advanced topology analytics, you should expect:

- Traditional/SDx/cloud network relationships
- End-to-end, multi-hop views and triage
- Global views of multi-vendor network health
- Telco-scale network support
- Correlated device metrics like performance, flow, logs and alarms for pinpoint triage

With 'one-click' visibility into the health and performance of your entire global network, you can confidently identify the root cause of managed or unmanaged network issues and assess their impact on branch sites, remote user productivity, or even the reliability of Gen-AI-driven prompt responses.

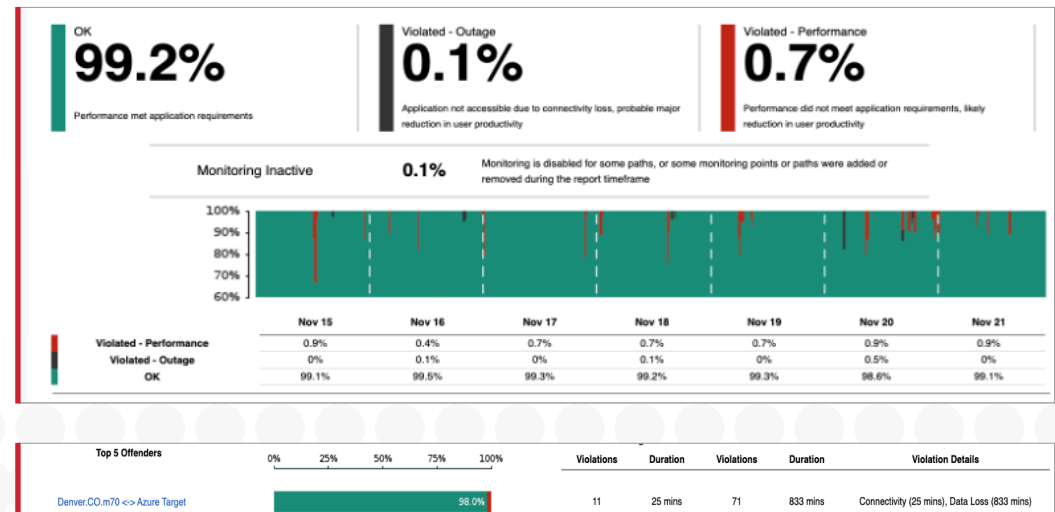
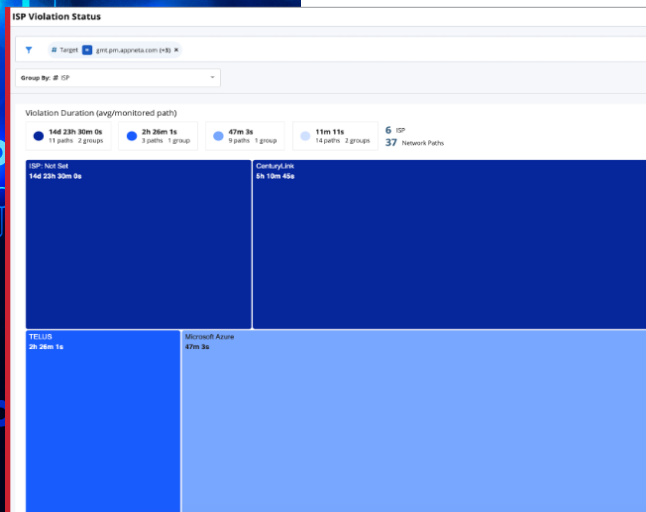


# #7

## Site-to-Site Network Delivery SLA Visibility

Internet service providers play a pivotal role in supporting business-critical apps and services. Additionally, with hybrid work, the requirements on IT to troubleshoot residential, branch office, and last-mile ISP network issues is growing, even as inherent visibility into these networks remains lacking.

Network Observability should enable active monitoring of the complete end-to-end network path. This includes visibility into transit and last-mile ISP performance to help IT quickly understand the root cause of performance issues. So, if that root cause is with the ISP, you can easily accelerate your mean time to insights and prove your innocence with granular performance metrics from inside the ISP network.





# Real-Time, End-to-End, Hop-by-Hop Network Path Analysis

Network teams have lost the operational visibility required by AI needed to ensure quality end-to-end, hop-by-hop network delivery.

To meet this requirement, look for network observability solutions that include:

- ISP/cloud network path visibility to find team innocence, while holding providers accountable for performance.
- Web/URL testing powered by synthetic transaction monitoring to identify SaaS and web app issues caused by unexpected network performance.
- Deep packet inspection for raw data analysis, when and where you need it.

With end-to-end views into site-to-site, site-to-cloud, and cloud-to-cloud network paths for business-critical services, teams can improve problem isolation and remediation times, and confidently network delivery performance by leveraging intelligent insights into unmanaged network devices.



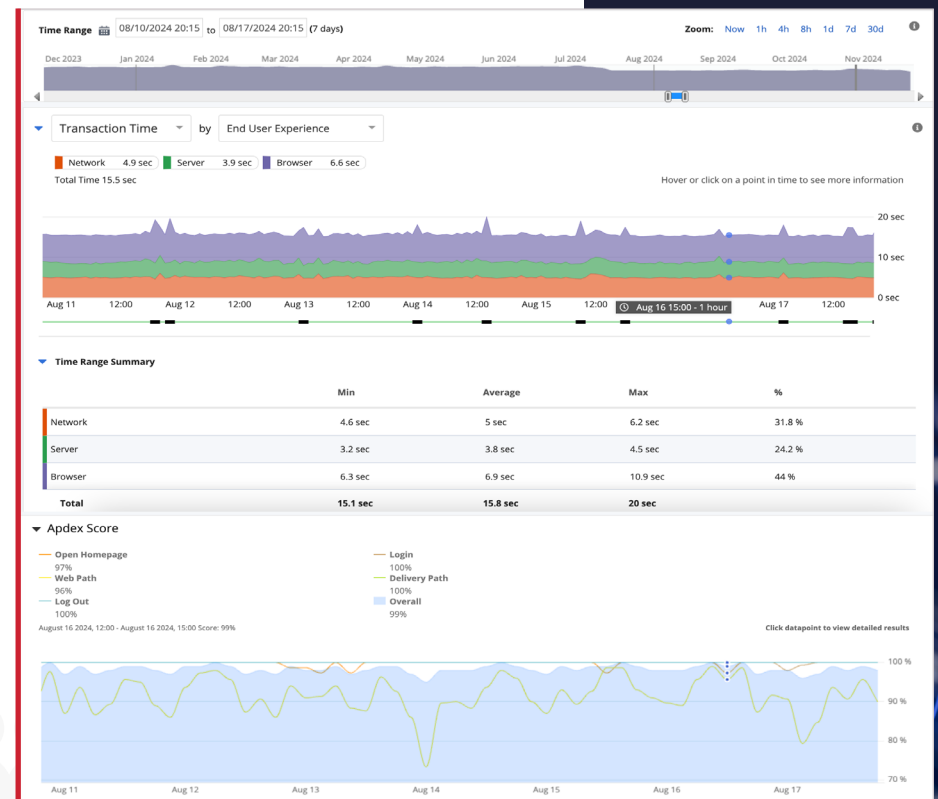
# App-Aware Network Synthetic Testing

While apps rarely go down these days, “slow” is the new down for users distributed across the globe. A deeper understanding of network delivery’s impact on application performance is required now. Synthetic transaction monitoring provides details for any application by sending tests via scripts to mimic user interaction.

Collecting the following metrics across network server and browsers will always help to understand the negative network performance impact on user productivity:

- Transaction times
- Page load times
- Redirects
- DNS lookups
- TCP/SSL connections
- HTTP server response
- Download times
- Apdex scores

App-aware network synthetic testing can help teams quickly answer the question “is it the app or is it the network?”

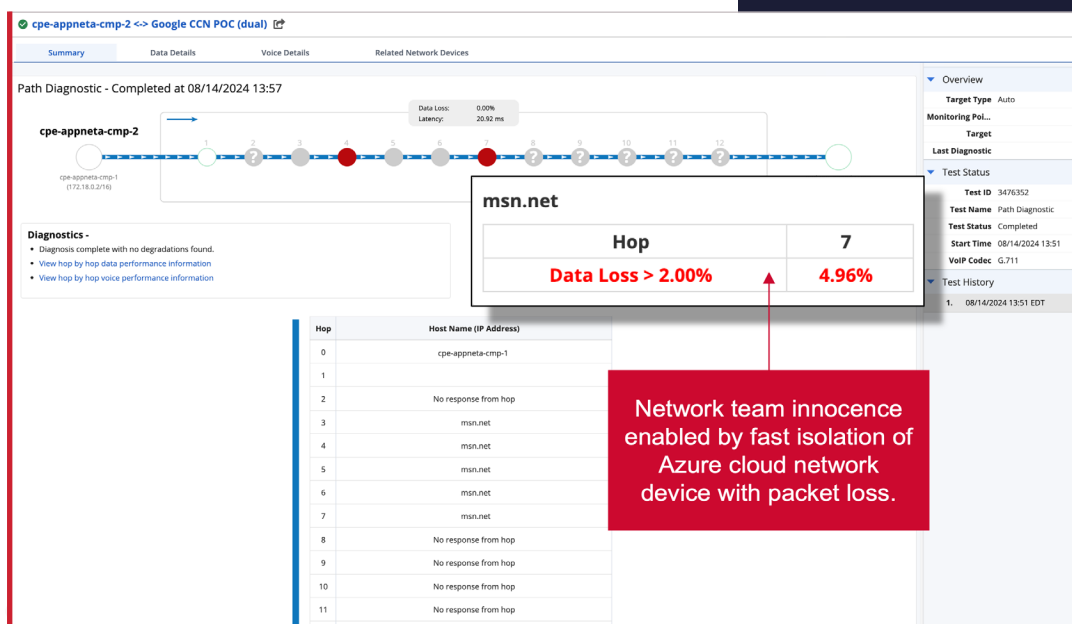


# Active Performance Validation of ISP/Cloud Networks

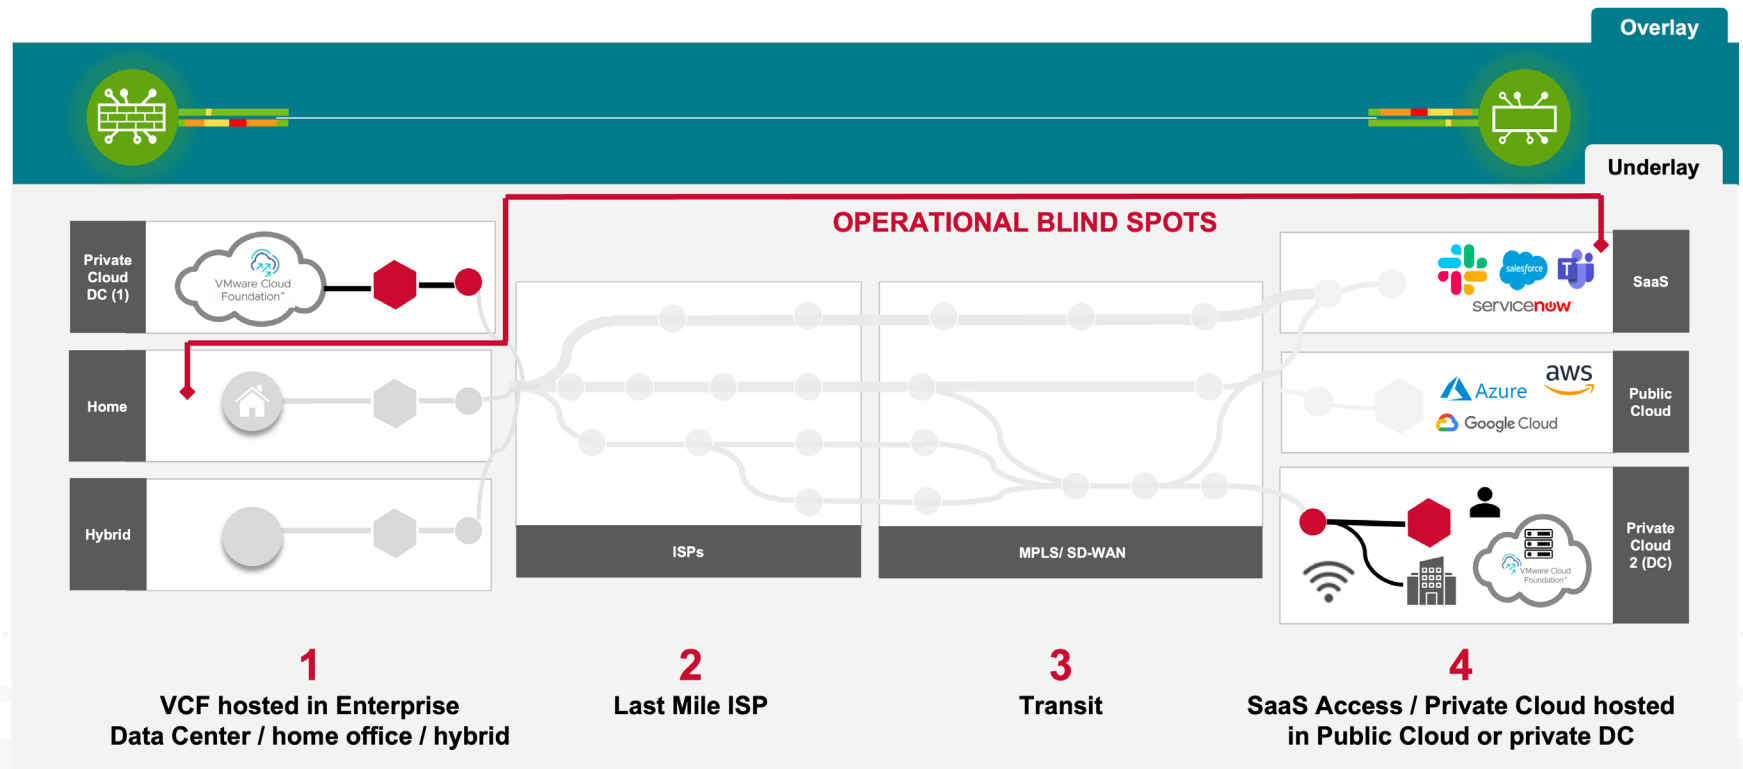
With the widespread use of public networks now, network teams should adopt a “trust but verify” approach but unlike in private data centers, we typically don’t have full administrative access to public ISPs or cloud infrastructure. To combat this, look for observability solutions that offer the following capabilities:

- Proprietary TruPath technology based on packet train dispersion to isolate hop-by-hop performance, regardless of who owns the network.
- Visibility into hybrid work devices, including host metrics, top processes, and connectivity type (such as VPN, wired, or wireless).
- Robust JSON-based API that enables push and pull data integration into existing IT operations solutions and workflows.

AI workloads will undoubtedly be disturbed for some use cases. It will always be better to visualize, analyze and validate every hop in your AI workload network path to ensure success.




# Embrace the New AI Workload Network Path





# And Pinpoint Issues Affecting Your AI Workloads



# You Trust Us to Build Your AI Networks, Trust Us to Assure Them



"Broadcom brings user experience metrics into the NOC for a better understanding of the managed and unmanaged network delivery performance impact on applications and users. By correlating network path metrics with network device performance, root cause and end-to-end network path health are surfaced, enabling the operations teams to get a better perspective of user experience impact."

**Andrew Green, Technology Analyst,**  
GigaOm Radar for Network Observability, 2024



"Broadcom offers a strong feature set with very powerful scalability. It also offers a formidable combination of observability for corporate networks, the internet, and the cloud."

**Shamus McGillicuddy, Vice President at Enterprise Management Associates (EMA)**  
EMA Radar Report for Network Performance Management, 2024



# IS YOUR NETWORK REALLY READY FOR AI?

LEARN MORE TODAY.

[HTTPS://NETWORKOBSERVABILITY.  
BROADCOM.COM/](https://networkobservability.broadcom.com/)



## About Broadcom

Broadcom Inc. (NASDAQ: AVGO) is a global technology leader that designs, develops, and supplies a broad range of semiconductor, enterprise software, and security solutions. Broadcom's category-leading product portfolio serves critical markets including cloud, data center, networking, broadband, wireless, storage, industrial, and enterprise software. Our solutions include service provider and enterprise networking and storage, mobile device and broadband connectivity, mainframe, cybersecurity, and private and hybrid cloud infrastructure. For more information, go to [www.broadcom.com](http://www.broadcom.com).

Copyright © 2024 Broadcom. All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.  
December 5, 2024