White Paper

Close the Network Visibility Gaps That Legacy and App-Centric Monitoring Miss

A Technical Perspective for Unified Network Observability

SNMP UPTIME DOESN'T PROVE PERFORMANCE. NETWORK TEAMS MUST MOVE BEYOND STATIC POLLING TO CONTINUOUS, ACTIVE PATH TESTING THAT DELIVERS HOP-BY-HOP VISIBILITY ACROSS ON-PREMISES, CLOUD, AND PROVIDER NETWORKS. THIS EXPOSES PACKET LOSS OR CONFIGURATION DRIFT INSTANTLY, ENDING THE *PROVE IT'S NOT THE NETWORK* CYCLE AND PROTECTING CRITICAL SERVICES AND AI WORKLOADS.

Introduction

Today's network operations team is trapped in a multi-screen, swivel-chair nightmare.

Imagine a network analyst halfway through a grueling 12-hour shift, racing against a ticking 30-minute SLA window to isolate a Sev-1 blast radius. Under a tight SLA, a tired analyst parses a flapping SD-WAN BGP peer on Monitor 1, unable to distinguish unmanaged ISP brownouts from configuration drift. Monitor 2 triages a legacy SNMP alert, blind to real-time packet loss and latency spikes due to slow polling cycles. On Monitor 3, an APM dashboard falsely blames network latency for what are actually TCP socket timeouts on an overloaded container host.

This *good enough* approach fosters a culture dominated by mean time to innocence (MTTI). Engineers waste hours on manual scavenger hunts—SSH'ing into CLIs, running traceroutes, and pulling interface stats—just to prove the underlay is innocent. This reactive fragmented monitoring fails under demanding AI workloads, where even minor, intermittent packet loss or network congestion can derail high-throughput GPU training.

The Network Observability Maturity Model

The network observability maturity model benchmarks organizational capability against industry standards, but they shouldn't judge frontline NetOps skill. They should measure the operational cognitive load imposed by the toolset. In early maturity stages, under-tooled teams are unfairly forced to act as the human correlation engine. Real progress across the following dimensions of network observability—coverage, architecture, workflow, and analytics—liberates engineering talent. Replacing fragmented dashboards with automated correlation shifts teams from constant firefighting to proactive optimization.

Advancing along the maturity curve eliminates outage-driven administrative toil. High-maturity observability automates triage, removing the need to manually build timelines or copy-paste logs across disconnected systems. Trusted, correlated data is essential for safe automation; without it, triggering SOPs or agentic AI remediation becomes dangerous guesswork. Integrated platforms instantly map dependencies and enrich tickets with forensic context, sharply reducing cognitive fatigue and allowing engineers to focus on solving complex issues rather than wrestling with fragmented tools.

How does maturing network observability improve the NOC's ability to detect, triage, and resolve network issues?

ALERT FATIGUE IS AN ARCHITECTURE PROBLEM, NOT AN OPERATOR PROBLEM. A MODERN PLATFORM MUST AUTOMATICALLY CORRELATE THOUSANDS OF DOWNSTREAM SYMPTOMS INTO ONE ACTIONABLE ROOT-CAUSE INCIDENT. THIS SHIELDS FRONT-LINE TEAMS FROM ALERT NOISE SO THEY CAN FOCUS ON THE REAL ISSUE AND DRIVE FASTER RESOLUTION.

Stabilize the Internal Baseline

Before network teams can confidently predict future performance drops, it must first gain a high-fidelity view of the present and stabilize internal operations. This foundational visibility is established during Maturity Levels 1 and 2. The primary challenge at this stage is normalizing massive performance datasets derived from complex multi-vendor hardware environments at scale.

AI WORKLOADS DEMAND MINIMAL TO ZERO PACKET LOSS, AND STATIC BASELINES CAN'T GUARANTEE THAT. CORRELATING REAL-TIME PERFORMANCE WITH L2/L3 TOPOLOGY AND CONFIGURATION DATA REVEALS THE TRUE IMPACT OF EVERY CHANGE. THAT CONTEXT SUPPRESSES FALSE ALARMS AND PROTECTS THE LOW-LATENCY, HIGH-THROUGHPUT BEHAVIOR YOUR GPU CLUSTERS RELY ON.

Stabilizing an enterprise network at Levels 1 and 2 isn't an abstract data exercise—it's a daily battle against configuration drift. Brittle, home-grown cron jobs and disjointed tracking scripts frequently fail at scale, leaving L1 teams completely blind when unauthorized changes occur. True stabilization happens when real-time performance baselines are natively integrated with network configuration states. When an uncoordinated change occurs at 3:00 a.m., the NOC shouldn't have to wait for morning user validations or manually run diffs against a text file; the performance engine should instantly flag the configuration drift before the alert storm hits the service desk.

Extend Visibility Outward: End-to-End Network Path Insights

TRANSITION EXTERNAL PATH TRACKING TO CONTINUOUS ACTIVE SYNTHETIC TESTING TO EXPOSE CLOUD AND SAAS BLIND SPOTS, AND ADD DUAL-ENDED PATH MONITORING TO MEASURE CAPACITY, LATENCY, JITTER, AND PACKET LOSS AT EVERY HOP.

Once internal infrastructure is stabilized, the next step forces a transition to a proactive operational posture (Maturity Level 3). Traditional configuration and passive monitoring tools fall short here because they run into a blind spot beyond the corporate perimeter—specifically across public cloud environments, SaaS providers, and third-party ISPs.

Shifting to a proactive posture at Maturity Level 3 requires illuminating the blind spot beyond the corporate edge. The moment traffic leaves an SD-WAN gateway and enters a public cloud VPC or a third-party ISP path, traditional passive monitoring completely breaks down. When remote users hit latency bottlenecks accessing critical SaaS fabrics, the internal corporate interfaces report completely green status. Shifting outward means augmenting passive traps with dual-ended active synthetics that continuously map external transit. This path observability measures capacity, jitter, and packet loss hop-by-hop, allowing the NOC to instantly pinpoint whether an application issue sits with a cloud transit gateway or a localized provider outage before the ticket queue explodes.

How can network teams achieve true data center-to-cloud visibility to speed up issue detection and response?

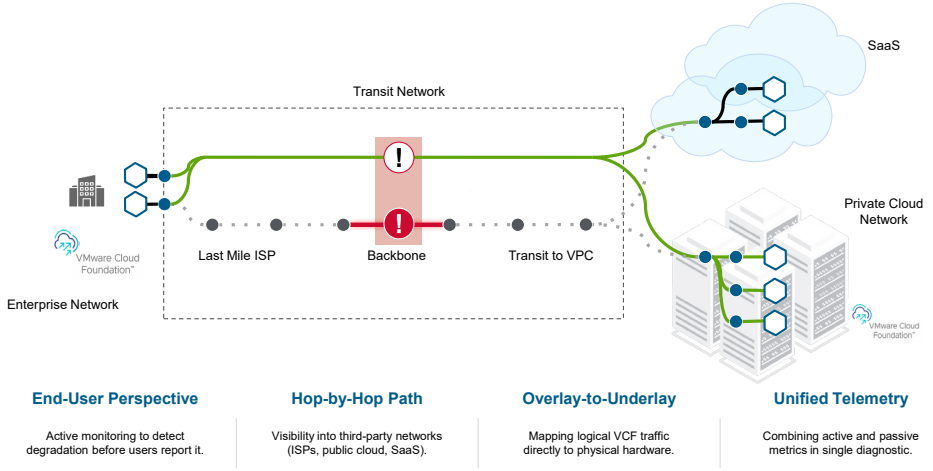

The Critical Flaws of App-Centric Models

SEPARATE APPLICATION-LEVEL DIAGNOSTICS FROM DEDICATED, MULTI-VENDOR NETWORK OBSERVABILITY LAYERS, AND DEPLOY FULL-FIDELITY, UNSAMPLED FLOW ENGINES TO PRESERVE DEEP FORENSIC PATH TRACING ACROSS EVERY HARDWARE DOMAIN.

Some cloud-native, application-centric monitoring tools such as DataDog, Dynatrace and others, attempt to stretch their footprints into network territory by relying on agents or Extended Berkeley Packet Filter (eBPF) collection. While valuable for code profiling, this can create a network blind spot. eBPF demands specific host kernels, making it irrelevant for headless network appliances, core routers, hardware firewalls, or third-party SaaS fabrics.

When an app-centric agent flags an application slowdown, it leaves the infrastructure team blind to backend transport realities. You cannot install an agent on a specialized headless switch or a raw RDMA over Converged Ethernet (RoCE) fabric. If network congestion or asymmetric path mapping throttles internode GPU communications, an APM dashboard sees a generic timeout, but it cannot distinguish between a slow database query and network path degradation. Complex network operations requires an unsampled flow engine capable of correlating application logical traffic directly with the physical underlay across all multi-vendor hardware domains.

APP-CENTRIC TOOLS ARE VITAL FOR DEVELOPERS, BUT THEY WEREN'T PURPOSE-BUILT TO HANDLE MULTI-VENDOR NETWORK COMPLEXITY. MODERN NETOPS REQUIRES NORMALIZING ALL TELEMETRY—SNMP, FLOW, AND CLOUD APIS—INTO A UNIFIED DATA MODEL. THAT CREATES A SINGLE SOURCE OF NETWORK TRUTH, GIVING YOUR TEAM END-TO-END VISIBILITY WITHOUT CONSTANT TOOL-SWITCHING.

Unify Network Insight, Not Just Managing Tools

Insisting on a single, aggregated monitoring agent for all application and infrastructure domains usually stems from a mandate to simplify compliance and governance. Enterprise leadership fears that maintaining distributed domain tools will lead to operational silos and hidden blind spots. However, forcing an app-centric or hardware-centric tool to act as an all-in-one network engine can introduce critical bottlenecks.

Enterprise IT can elegantly resolve this friction by centralizing telemetry data rather than forcing a single control tool. By normalizing disparate telemetry sources—including SNMP, REST API, IPFIX, and active synthetic results—into a unified operational data model, specialized teams can use optimized, domain-specific pipelines. An open platform avoids data silos by delivering authoritative network truth via standards-based integrations, including Kafka streaming and OpenTelemetry export, enriching data lakes and third-party systems. This strategic shift guarantees full visibility across the application delivery chain without sacrificing the deep forensic capability engineers require.

How can Broadcom help network teams go beyond *good enough* monitoring to achieve true network observability?

FAULT | Analytics

Reduce MTTR by **rapidly isolating error domains using root cause analysis, topology, and configuration management** across multi-vendor, multi-technology networks.

PERFORMANCE | Analytics

Ensure service availability by leveraging **algorithmic anomaly detection, automated baselining, trend analysis, and capacity planning** across diverse data streams (SNMP/API/Telemetry).

FLOW | Analytics

Improve network efficiency with end-to-end **traffic visibility by understanding bandwidth use, capacity, configuration, and app performance,** leading to optimized resource allocation.

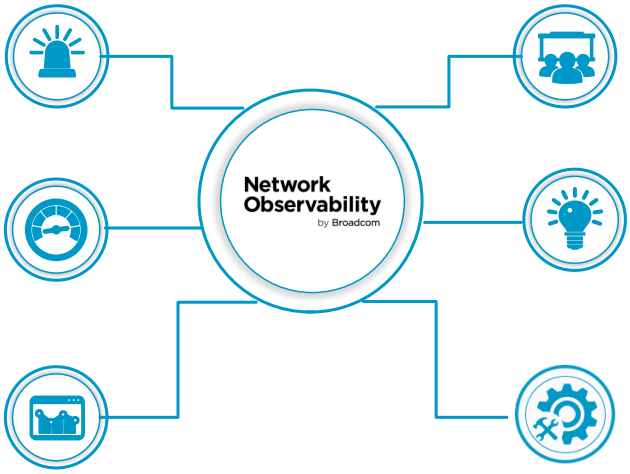

USER EXPERIENCE | Monitoring

Enhance **user experiences** and quickly identify potential issues with **proactive visibility into network and application performance** across all domains.

MODERN | Network Monitoring

Optimize SDx performance with **comprehensive visibility into SD-WAN, SDDC, and API** controllers via event, topology, and performance mapping.

NETWORK CONFIGURATION | Configuration

Optimize performance with **the ability to monitor, change and update configurations of devices and systems**

Comprehensive, industry-leading end-to-end Network Observability solution that is technology-agnostic and scalable.

PREVENT CROSS-DOMAIN FINGER-POINTING BY IMPLEMENTING A UNIFIED OPERATIONAL DATA LAKE ACROSS TECHNOLOGY SILOS, AND CORRELATE PHYSICAL UNDERLAY HEALTH WITH VIRTUAL OVERLAYS TO INSTANTLY ISOLATE ROOT-CAUSE FAULTS ACROSS HYBRID APPLICATION PATHS.

Real-World Application: Move From Finger Pointing to Fast Collaboration

Allowing separate application and infrastructure domains to operate without a correlated data lake introduces a painful side effect: an entrenched mean time to innocence (MTTI) culture. A global financial services firm managing over 300 remote branch offices running an SD-WAN overlay uncovered this visibility gap during a major cross-domain disruption. Their legacy APM tools flagged intermittent user experience degradation on a core web application, but could not pinpoint the cause. By aggregating multi-vendor state logs, unsampled IPFIX flow data, and active synthetic path tests into a single data lake, the NOC instantly isolated the root cause: an asymmetric routing loop between their multicloud VPC edge and an external SaaS provider network, bypassing local firewall rule states. Fault isolation dropped from four hours of siloed finger-pointing down to minutes of collaborative remediation.

Ready to Dismantle Your MTTI Culture?

Good enough monitoring isn't enough. To move from reactive firefighting to proactive network operations, unified end-to-end visibility is essential—and the Broadcom® Network Observability solution delivers that today while enabling network operations for the demands of AI-driven workloads.

We're advancing the human-to-platform interface through ongoing investment in AI, enabling teams to query the full depth of observability data in natural language and receive immediate, actionable answers. Looking ahead, our vision extends further: unified APIs that empower Agentic AI to evaluate network state, diagnose faults, and trigger remediation. This progression leads to our North Star—Zero-Dashboard Operations—where the platform proactively identifies issues, determines root cause, and trigger needed remediation workflows.

Learn more about the Broadcom Network Observability solution:
networkobservability.broadcom.com

BROADCOM®

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

Unified-Network-Observability-WP100 June 18, 2026