




The First Network AI
for the Modern Enterprise

Managed Service Provider Case Study

A Global MSP Augments their Enterprise SD-WAN Service with Augtera Network AI to Solve Application Performance Challenge

Augtera networks enabled an MSP to improve their competitive positioning and the user experience of their customers by detecting circuit performance issues, without telemetry from the circuit provider.





Challenge

A large Managed Service Provider (MSP) needed a way to detect performance issues on SD-WAN underlay circuits when they were not the circuit provider and only had access to the SD-WAN telemetry.

Solution

Augtera Networks AI platform: Real-time AI/ML automatic anomaly detection and correlation of the SD-WAN probe telemetry (e.g., RTT, loss, jitter) per site and path using native Versa logs.

Results

MSP has now achieved a competitive advantage by enabling their enterprise customers to detect and be notified of performance degradations on non-MSP managed circuits.

Problem Statement

“Enterprises often need to use multiple sources and types of broadband circuits in our managed SD-WAN service. They wanted us to be able to detect when a circuit we didn’t manage was having performance issues. We had challenges doing that with our current SD-WAN tools” said MSP’s Director of Customer Service.

Customers of MSP requested a capability to detect and be notified of issues on all broadband connections used in the managed SD-WAN service. Solving this problem would improve the user application performance and give MSP a competitive advantage. Many of MSP’s enterprise customers augment MSP circuits, by using other 3rd party circuit providers. Having a 3rd party ISP as part of the underlay, with no ability to monitor the circuit directly, in unison with a SD-WAN overlay, translated into a series of circuit performance detection challenges, ultimately impacting application performance when issues continued undetected.

Company Profile

This MSP provides worldwide network and voice connectivity to enterprises, across 30 countries and 200 cities. MSP’s enterprise market segments include financial, media, and cloud services. MSP has built a large SD-WAN practice for secure hybrid-cloud, using Versa Networks.



UNDERLAY PROBLEM

Enterprises often have limited choices when choosing broadband circuits, because of cost and/or location. Enterprises were unable to detect performance degradation, especially when their edge devices are serviced by multiple circuit providers and service types (e.g., LTE, MPLS, Internet). MSP was previously only able to detect issues on circuits they provided using their internal monitoring platform. MSP and their enterprise customers both had limited telemetry from 3rd party circuit providers. A solution was needed that could leverage telemetry from the overlay network, to detect an underlay issue with a specific circuit.

OVERLAY PROBLEM

MSP's Versa SD-WAN, along with other SD-WAN vendors, monitor the performance of the SD-WAN overlay using synthetic probe data such as RTT, packet loss, and jitter for each path and tunnel. This telemetry is primarily used for application performance monitoring and is the key to finding underlay issues. The challenge is analyzing the very large volume of data generated, and knowing when a particular

metric is abnormal for a given path and tunnel. Typically fault detection would use static threshold settings for each path. However, this manual method was untenable due to the large number of paths and because each path has unique performance characteristics. As an example, a SD-WAN VPN with 100 sites connected in full mesh could equate to 237,600* graphs¹.

MSP REQUIREMENTS

The ideal solution starts with a consolidated view of each enterprise customer's SD-WAN network, available to both the MSP and the enterprise, for visualizing the underlay circuit performance. The solution needs to learn what is normal for each path, using all available metrics, and detect when there is an abnormal performance degradation and identify the specific faulty circuit(s). In addition, the solution should be able to send a notification to the enterprise's operations tools. This enables a ticket to be opened with the provider for the faulty circuit.

¹ 100 locations * 99 full meshed paths * 2 circuits per location * 4 Queues * 3 metrics (RTT, Loss, Jitter)



Solution

The Augtera Machine Learning offered an ideal solution for MSP, by natively providing three attributes:

1. A platform that can learn and detect abnormal changes in metric patterns, in real-time, and at scale.
2. An ability to correlate path/tunnel anomalies, with the circuit name/ID in order to isolate the issue to a specific circuit and/or provider.
3. Ability to leverage the available native SD-WAN telemetry. Augtera consumes data non-intrusively. Versa provides the data in the form of logs, and Augtera passively monitors those logs.



Results

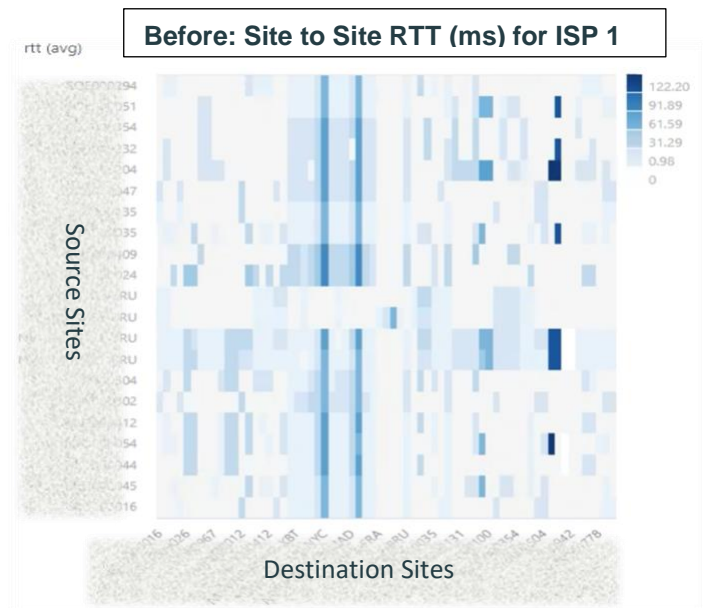
MSP was able to solve their Enterprise customers' challenge, and that has improved the enterprise customers' user experience. It also has improved MSP's competitive advantage, in a highly competitive provider market.

Results after deploying Augtera Network AI were:

1. MSP's enterprise customers can now observe all overlay performance degradations from any site to any other site, for any given underlay circuit. Augtera identifies ML-based anomalies, by removing noise in the data. Figure 1 below shows the "BEFORE" noisy raw fault data.

2. Enterprise customers can see anomaly patterns on a heatmap. For example, the heatmap can show if many or all paths have been affected simultaneously using same underlay circuit, thus pointing to a specific faulty ISP. Figure 2 shows the "AFTER" AI insight data produced by Augtera Machine Learning.
3. A notification can be generated to the Enterprise's operations tools, such as Slack or ServiceNow, when an underlay circuit anomaly is detected

Figure 1: BEFORE – The X and Y axis shows the raw RTT data between each site. It is very noisy and the variance in RTT between sites makes it extremely difficult to use thresholds as it is impossible to know if RTT is in a good or bad state at any given time.



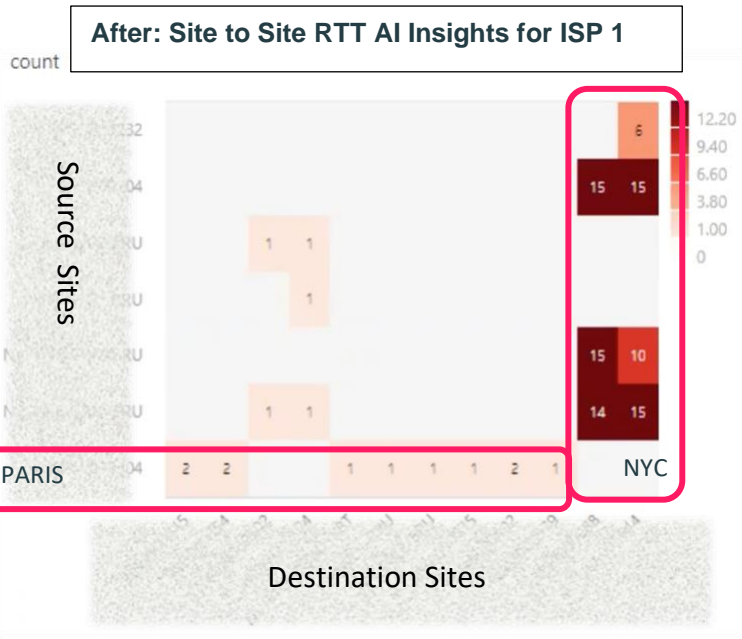
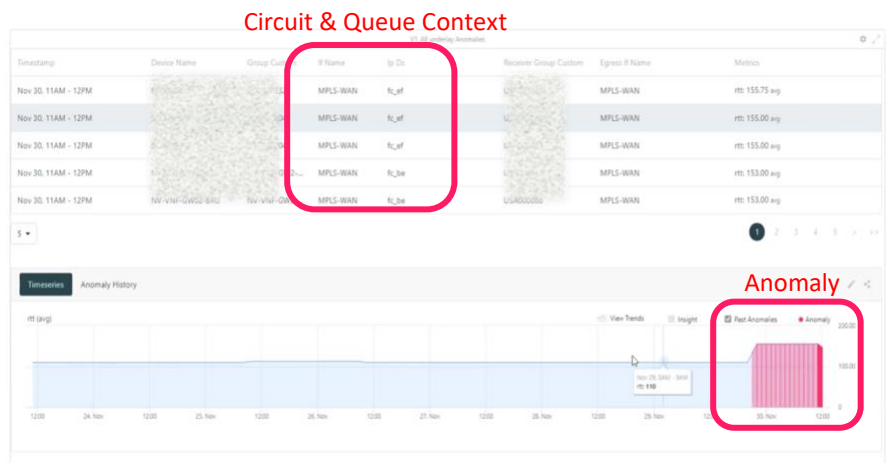


Figure 2: AFTER – AI insights produced by Augtera AI after processing the same RTT data from Figure 1. It shows the volume of RTT anomalies in the last 24 hours. Each number indicates a period of time that the RTT pattern was materially degraded. Paris had RTT anomalies across all other sites it connects to, but had a shorter count, illustrating a short duration. NYC on the other hand, had sustained RTT anomalies, indicating prolonged issues.

Actionable Insights: Augtera Performance Anomalies

Figure 3: The Augtera AI anomalies that are represented in Figure 2. The anomaly context shows the source device, receiving device and the circuit type along with the queue info. In this example, multiple anomalies on the MPLS paths indicate the MPLS circuit has sustained RTT degradation. These anomalies are sent directly to the enterprise operations tools.





Benefits of Augtera

- Detects underlay performance issues per circuit, regardless of circuit provider.
- Improves customer experience and competitive positioning for MSP.
- Eliminates manual configuration rules. Augtera ML continually learns anomaly patterns at scale, and dramatically improves operational productivity.
- Provides real-time and customizable heatmaps, with analytics of the SD-WAN environment.

Augtera Networks is changing the way organizations operate their networks. Please see our online [demos](#) and [contact us](#) to see what Augtera can do to improve your network operations.