

# Arista AVA: The Power of AI-Driven Networking

# Table of contents

|  |    |
|--|----|
| <b>Introduction</b>  | 3  |
| <b>AVA Architecture</b>  | 4  |
| <b>The AVA Advantage</b>   | 5  |
| <i>A Better Approach: Learning What's "Normal"</i>                     | 5  |
| <i>The Power of the Knowledge Graph</i>                                | 5  |
| <b>AVA Delivers Solutions</b>  | 6  |
| Case Study 1 – Proactive Network Operations                            | 6  |
| Case Study 2 – AIOps for Network Access Control                        | 7  |
| Case Study 3 - Rapid Issue Inference and Mitigation                    | 8  |
| Case Study 4 - Accelerating Network Troubleshooting with Generative AI | 9  |
| Case Study 5 – Autonomous Network Detection and Response               | 10 |
| Case Study 6 – Quality of Experience (QoE)                             | 11 |
| Case Study 7 – IoT Observability and Security                          | 12 |
| Case Study 8 – Enriching the Ecosystem with AVA Insights               | 13 |
| <b>Summary: Data-Driven Networking Made Possible</b>                   | 14 |

## Introduction

Artificial intelligence (AI) has been all over the news especially as large language models like GPT-4 capture the imagination of even the average internet user, let alone the technophiles. These solutions have the potential to change the very definition of the enterprise, what it means to work, and lower the cost of doing business. Of course, in parallel, the world is producing data at a blistering pace. A World Economic Forum study estimates that by 2025, we will be creating 463 exabytes of data every day. For context, that's 463, followed by 18 zeros! Or as the authors put it, "Forty times more bytes than there are stars in the observable universe." Of course, what isn't likely to happen in that same time frame is that humans will suddenly evolve to consider all of that available data, factor probabilities, and make the optimal operational choices.. That is not to suggest that humans are becoming expendable. Far from it, we have highly tuned abilities to recognize patterns<sup>2</sup>, understand abstract relationships and generalize. For instance, we don't need hundreds or thousands of training samples to know that a user being targeted by a phishing attempt is the CFO of the organization. We recognize the name right away and our instincts take over to drive remediation next steps. The task at hand then for the AI tools, is to surface just the right information that enables operators to make decisions which ultimately lead to the most favorable business outcomes.

In the world of **IT and security**, the network serves as a fundamental source for the data that fuels decision-making. However, the rapid expansion of networks has led to an overwhelming volume of data, making it difficult to find the right balance of actionable information. So how do you get to the goldilocks balance of "just right" information? It comes down to two key capabilities:

- **Efficient and Real-time Data Collection:** A system that can not only efficiently collect raw, or "ground truth," data from the network but also do so in real-time. Such a robust data foundation can be achieved with the following three key attributes:
  - » **Data Diversity:** Collecting from multiple network domains to give a comprehensive view.
  - » **Data Richness:** Capturing an extensive number of features to provide detailed context.
  - » **Data Fidelity:** Using real-time streaming to ensure the highest level of accuracy.
  - » **Intelligent Data Processing:** Having the Intelligence to extract valuable information and context from this raw data and present it in a clear, digestible format.

Arista is uniquely positioned to deliver on both of these capabilities. Arista EOS® based on our network data lake (NetDL™)<sup>3</sup>, provides a multi-domain, multi-modal, multi-tenant capable data lake that offers real-time network telemetry to other Arista solutions as well as those from our partners. Arista Autonomous Virtual Assist (AVA) uses an AI-driven approach to anticipate operator questions, extract answers from NetDL and deliver the insights necessary for effective human decision-making. With this combination, Arista is providing networking for the data-driven enterprise. In this paper, we will share how AVA uses cutting-edge AI models to solve real-world challenges for network and security operators and present specific case studies that illustrate the business value the approach can generate.

---

<sup>1</sup> <https://www.weforum.org/agenda/2019/04/how-much-data-is-generated-each-day-cf4bddf29f/>

<sup>2</sup> <https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf>

<sup>3</sup> <https://www.arista.com/en/solutions/cloud-networking>

AVA Architecture

AVA is an AI-enabled decision support system that provides an unprecedented level of visibility and responsiveness for network and security operations. It combines cloud scalability with the expertise of real-world professionals.

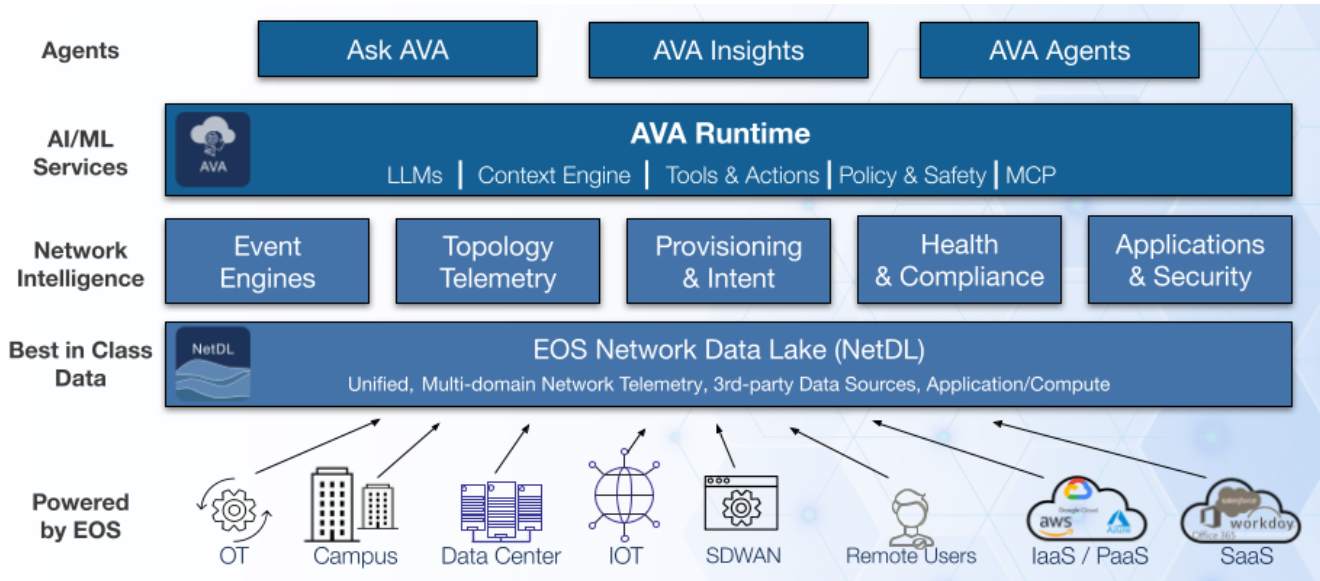


Figure 1: The Arista CloudVision's AI for Networking Stack

• **Foundational Network Intelligence**

At the core of AVA's AIOps capabilities is the EOS Network Data Lake (NetDL), which collects real-time, multi-domain telemetry data from a wide range of devices. This creates a unified and historically rich view of the entire network, allowing operators to manage it holistically.

AVA also surfaces early, weak signals of network issues along with corroborating evidence, while eliminating signals that can't be corroborated to avoid wasting valuable human cycles. This puts the operations team in a better position to act decisively. In fact, field results show that AVA often finds more incident-related activity than a senior human operator analyzing the same activity.

Furthermore, building on this foundation, AVA Events uses machine learning algorithms to deliver high-signal alerts for network problems. It proactively identifies issues like unusual connectivity jitter, failing optics, or a lack of disk space, reducing operational noise so teams can focus on the most impactful issues.

• **The Power of Agentic AI**

A new agentic AI layer builds on this foundation, beginning with Ask AVA. This chat assistant is integrated directly into CloudVision and uses generative AI to help operators troubleshoot and manage their network with simple, conversational answers. Its capabilities also expand to include automating provisioning changes and running network audits.

With AVA Insights, the vision extends beyond a chat assistant and moves from a reactive model to a proactive one, autonomously identifying and investigating issues as they arise, alerting users and providing answers before an operator even asks a question.

Looking ahead, a long list of other agentic features enabled by AVA are envisioned. We will have agents that analyze change controls to ensure deployments align with best practices, automate provisioning new networks, and autonomously monitor routing protocols to ensure changes match an administrator's intent. Each of these AVA agents will simplify network management, allowing human operators to focus on high-level strategy and innovation.

## The AVA Advantage

Legacy data science approaches face significant challenges with the scale and diversity of modern networks. AI models are only as effective as the labeled samples used to train them, but the sheer volume of network data and the types of conditions on even the simplest of today's networks make effective labels scarce. This problem is compounded by the fact that issues and threats evolve rapidly, making labels obsolete, and requiring constant retraining and high operational costs. Even if these challenges are overcome, the resulting models are often massive, monolithic, slow and lack explainability. In other words, a human analyst seeing the result often has no idea why something is flagged as a network or security issue. This leaves human analysts without confidence in the information or clear next steps. (Figure 2).

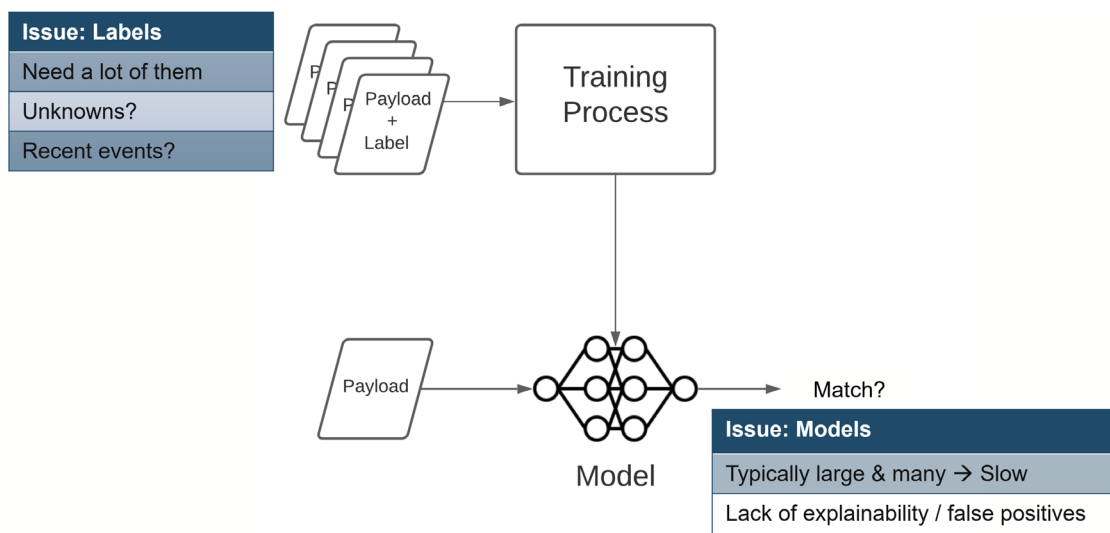


Figure 2: Challenges with legacy artificial intelligence approaches applied to network data

## A Better Approach: Learning What's "Normal"

Unlike legacy solutions (as described in Figure 2) that try to learn what "bad traffic" looks like, AVA takes a different approach. It starts by learning what is normal and routine on the network. This includes common behaviors like email, patching, and everyday application usage. This approach works because there is an abundance of labeled data for these mundane network behaviors. By eliminating this "hay" from the network data haystack, AVA leaves a much smaller set of data to analyze for potential issues.

## The Power of the Knowledge Graph

A key to AVA's success with processing the voluminous data is its use of a **knowledge graph**. Here's how it works:

- The voluminous data from the NetDL is first processed to discover entities, such as users, devices, VLANs and applications, and the relationships between them. This forms the core of the knowledge graph.
- This knowledge graph can be enhanced with insights from domain experts in the form of heuristics and ongoing feedback from human operators.
- The small size of the knowledge graph, compared to the raw data, allows for better, iterative AI analysis.
- Unlike traditional models, AVA captures algorithmic outputs as human-readable entities and properties. This delivers **explainable AI** with clearly defined next steps for the analyst.

AVA Delivers Solutions

The data-driven architecture, coupled with artificial intelligence, enables several network and security operations use cases. This section will provide a few examples of how AVA is optimizing workflows, speeding mean time to resolution, and improving security outcomes.

Case Study 1 – Proactive Network Operations

AVA enables network reachability modeling to deliver proactive notifications when a specific network service or application is experiencing reachability issues. Using dynamic anomaly detection, AVA identifies anomalies based on deviations from a learned reachability/latency baseline. The historical bounds and anomaly scores adapt to normal variations as time goes on. Access to NetDL’s historical data supports both real-time and forensic troubleshooting of any issues identified.

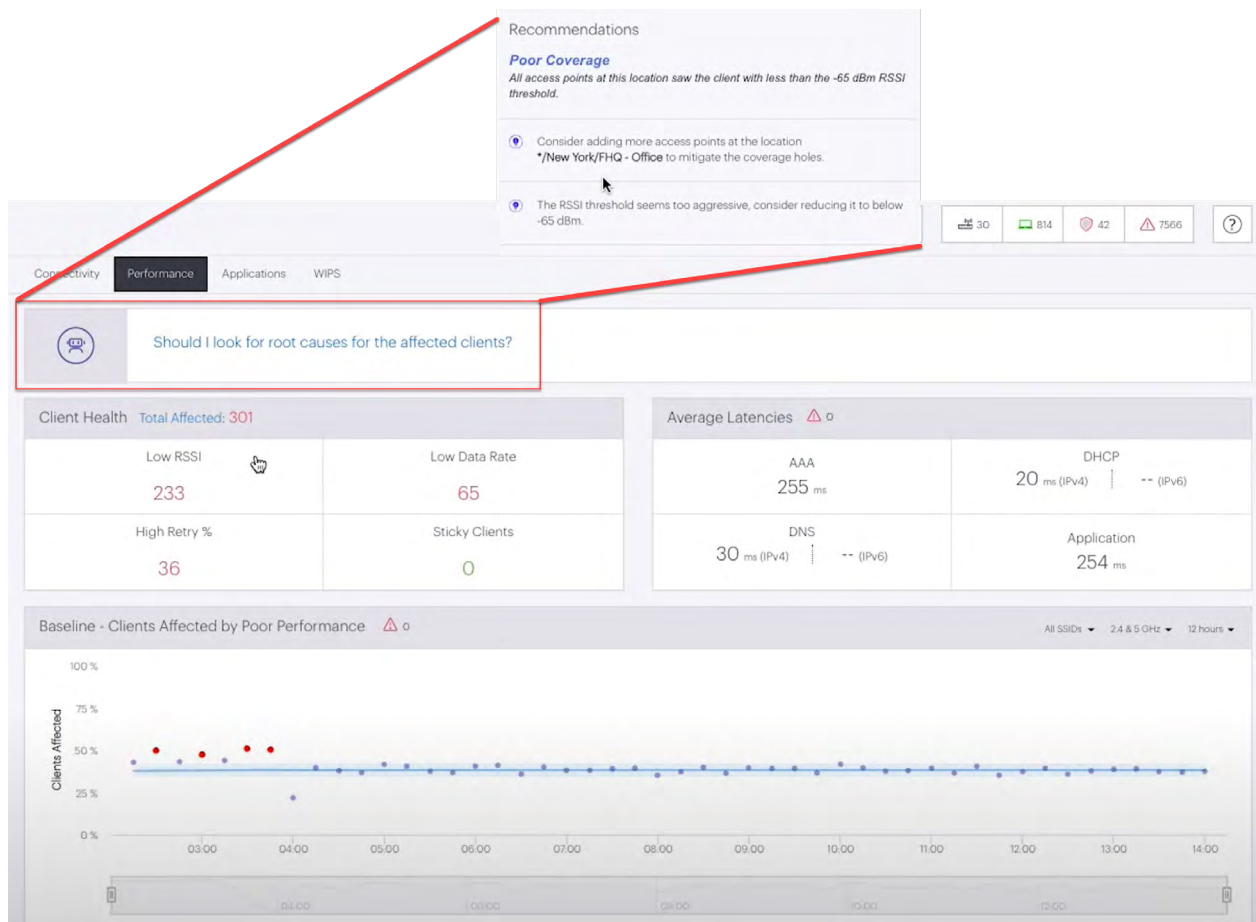


Figure 3: AVA proactively identifies network and application reachability issues

AVA also assists with another common network operation challenge: switch table overflow, resulting in a network outage. With the myriad chipset implementations and associated hardware table capacities, it is difficult for operators to monitor, track and extrapolate utilization trends manually. Instead, AVA models hardware resource utilization growth trends, enabling predictive assessments and notification ahead of exceeding capacity. Customers use these notifications to take preventive measures and avert a crisis.

← **Anomaly in Connectivity Monitor latency metric** on **bri516**

Lasted 34s — Started Jun 14, 2021 11:27:15 (17m ago)

Acknowledge

Configure Event Generation

Event Description

Detected anomaly in Connectivity Monitor latency metric

Cloudtracer latency values were detected outside historical bounds. The historical bounds are indicated by the envelope around the historical latency mean in the graph below.

CloudTracer Latency and Anomaly Graph



### Case Study 2 – AIOps for Network Access Control

Deploying and managing network access control (NAC) solutions have historically been cumbersome and error-prone, requiring expensive and hard-to-find human experts. Arista CloudVision AGNI innovates in this area by automating and optimizing key workflows, including providing a conversational AI capability called Ask AVA for configurations, troubleshooting, and simulations. Ask AVA uses various AI techniques including a generative pre-trained transformer (GPT)-based approach for training purposes.

AGNI has two sets of models: one based on classic machine learning uses support vector machines while the other uses ChatGPT’s large language models via OpenAI APIs. Ask AVA uses a chat-like interface and natural language processing (NLP) to aid even a junior analyst with tasks such as configuring, troubleshooting, and analyzing policy configurations. From the NLP queries, AVA autonomously identifies the analyst’s intent and assists in configuration, provides contextual output to troubleshoot problems, and analyzes the correctness of network and security policies. The net effect is a simpler and more secure NAC deployment.

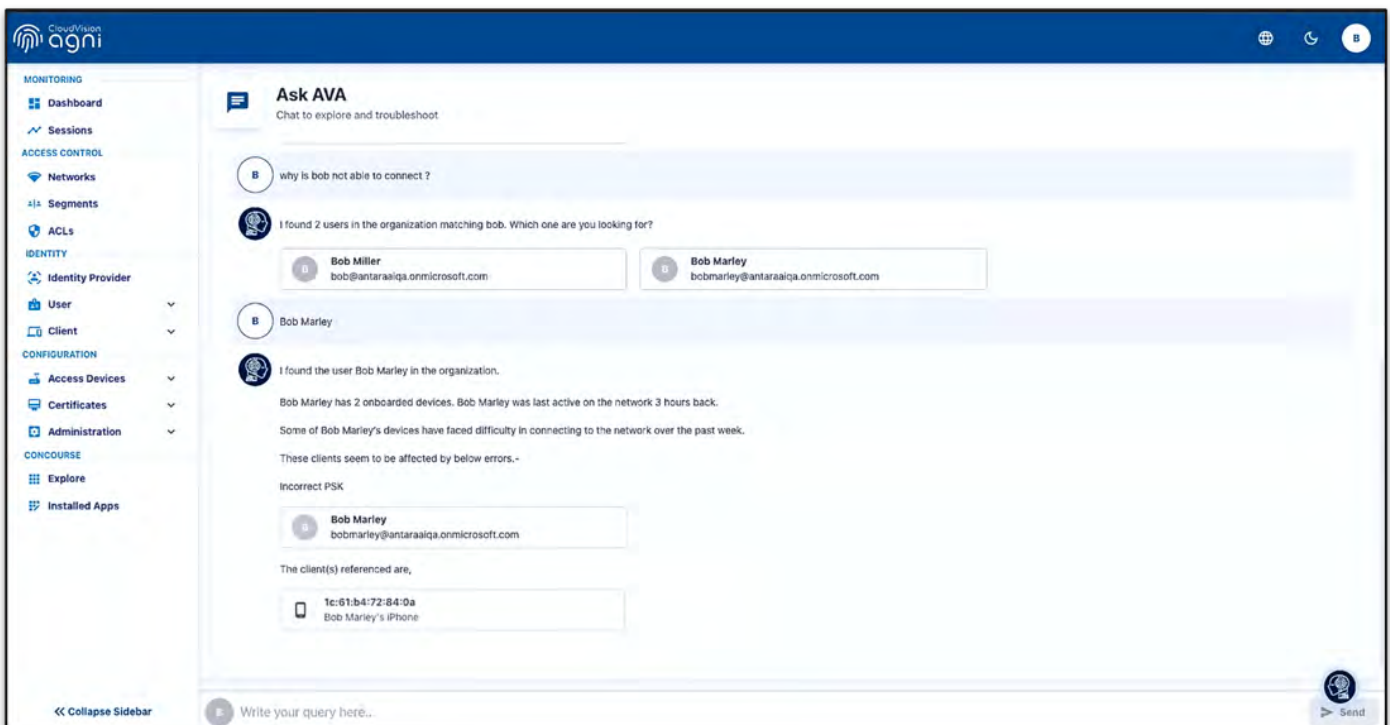


Figure 4: AVA streamlines operations and reduces errors in network access control workflows using conversational AI

### Case Study 3 - Rapid Issue Inference and Mitigation

With increasing network complexity, NetOps and DevOps teams struggle to pinpoint the root causes of performance issues. Questions like, “Is the issue the application or the network?” or “Where in the network is the problem?” lead to outdated and labor-intensive troubleshooting methods. CloudVision Universal Network Observability (CV UNO) uses AVA-based analysis to address these challenges. By leveraging composite application and network performance data stored in NetDL, CV UNO infers topology-aware correlations across events, changes, and anomalies. This proactively detects issues, accelerates problem-solving, and expedites resolution.



Figure 5: Using traffic analysis, AVA can provide rapid issue inference for network / application related issues

- Rapid Impact Analysis and Issue Inference

Modern network infrastructure generates an overwhelming volume of event data, making it difficult for operators to find the precise information they need to resolve issues. AVA cuts through this noise by surfacing only the early, weak signals of a network issue and corroborating evidence, eliminating unrelated signals avoiding wasted human cycles. This approach leaves a significantly smaller set of data to analyze, empowering the operations team to act decisively. Overlaying an end-to-end application flow through the network topology helps to accurately determine the root cause for network or even application/host related changes, eliminating the “finger-pointing” common with legacy approaches.

- Proactive Risk Analysis

Arista’s solutions are uniquely designed to move beyond reactive troubleshooting to a model of proactive prevention, helping you stop problems before they start. The Arista CI pipeline automates the entire network lifecycle, from design to deployment, to significantly reduce human error. As part of its powerful pre-flight checks, the pipeline integrates with CV UNO that leverages AVA’s analysis to identify and assess the potential impact of network changes on critical applications before they are ever deployed into production. This creates a proactive risk analysis workflow that prevents disruptive, application-impacting configurations from ever reaching your mission-critical networks. This is a unique offering that shifts the network troubleshooting operations from a reactive stance to a proactive one.

## Case Study 4 - Accelerating Network Troubleshooting with Generative AI

Network operations teams are under constant pressure to maintain uptime and ensure performance. Despite powerful monitoring tools, pinpointing the root cause of a specific issue in a complex, dynamic environment remains a time-consuming, manual process. Engineers must often sift through vast amounts of data, execute a series of ad-hoc commands, and manually correlate disparate information, leading to high Mean Time to Resolution (MTTR) and business disruption. Ask AVA is an agentic AI assistant integrated directly into CloudVision that transforms how network engineers troubleshoot. Using natural language, you can query, diagnose, and manage your network, empowering you to move from symptom to root cause in minutes.

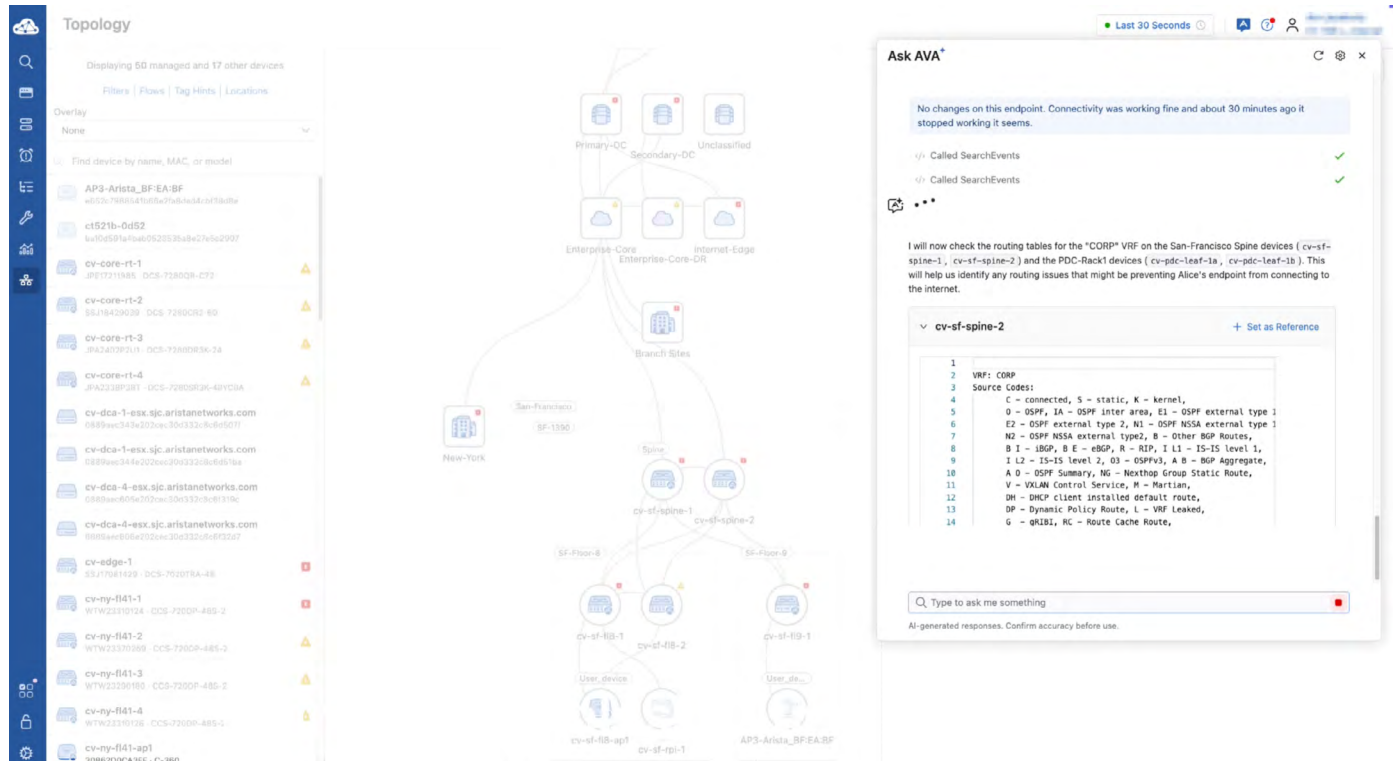


Figure 6: Using Generation AI, AVA transforms troubleshooting into a streamlined, guided experience

Take a case of a particular user not being able to reach the internet. Instead of a multi-hour manual investigation, a network engineer simply asks CloudVision the question as to why that user can't reach the internet. Ask AVA's AI engine takes over, seamlessly moving from a natural language query to a conclusive diagnosis:

- It automatically identifies the user's device and runs a series of diagnostic commands, like ping and traceroute.
- It intelligently analyzes the results and correlates data from across the platform, including ACLs and configuration change logs.
- In minutes, it pinpoints the exact line of code and change control responsible for the user's connectivity issue.

This process, which would have taken hours of manual work, is completed in minutes, transforming troubleshooting into a streamlined, guided experience. Ask AVA provides an intuitive, powerful way to leverage all of CloudVision's data and features. By intelligently aggregating and analyzing information, it empowers engineers to resolve complex issues with unprecedented speed and accuracy, freeing up valuable time to focus on strategic initiatives.

## Case Study 5 – Autonomous Network Detection and Response

Like the network operations use cases above, AVA can also deliver day 0 value to the security operations team. For instance, AVA uses unsupervised machine learning to identify and track users, devices, applications, etc., over time. AVA can also use this information to cluster similar entities. This presents a significant enhancement from legacy approaches that rely on unsupervised learning to spot anomalies from “normal” baselines for individual IP addresses rather than entities. Attributing behaviors to an IP address leads to high false positives and negatives, which translates to operational burdens on analysts.

Similarly, AVA uses supervised machine learning to identify patterns of activity that relate to attacker tactics, techniques, and procedures. For instance, AVA can classify remote access tools, reverse shells, unauthorized applications used for command and control, etc., without the need for decrypting the underlying data. This encrypted traffic analysis eliminates the privacy, policy, and technology challenges of decrypting data for analysis.

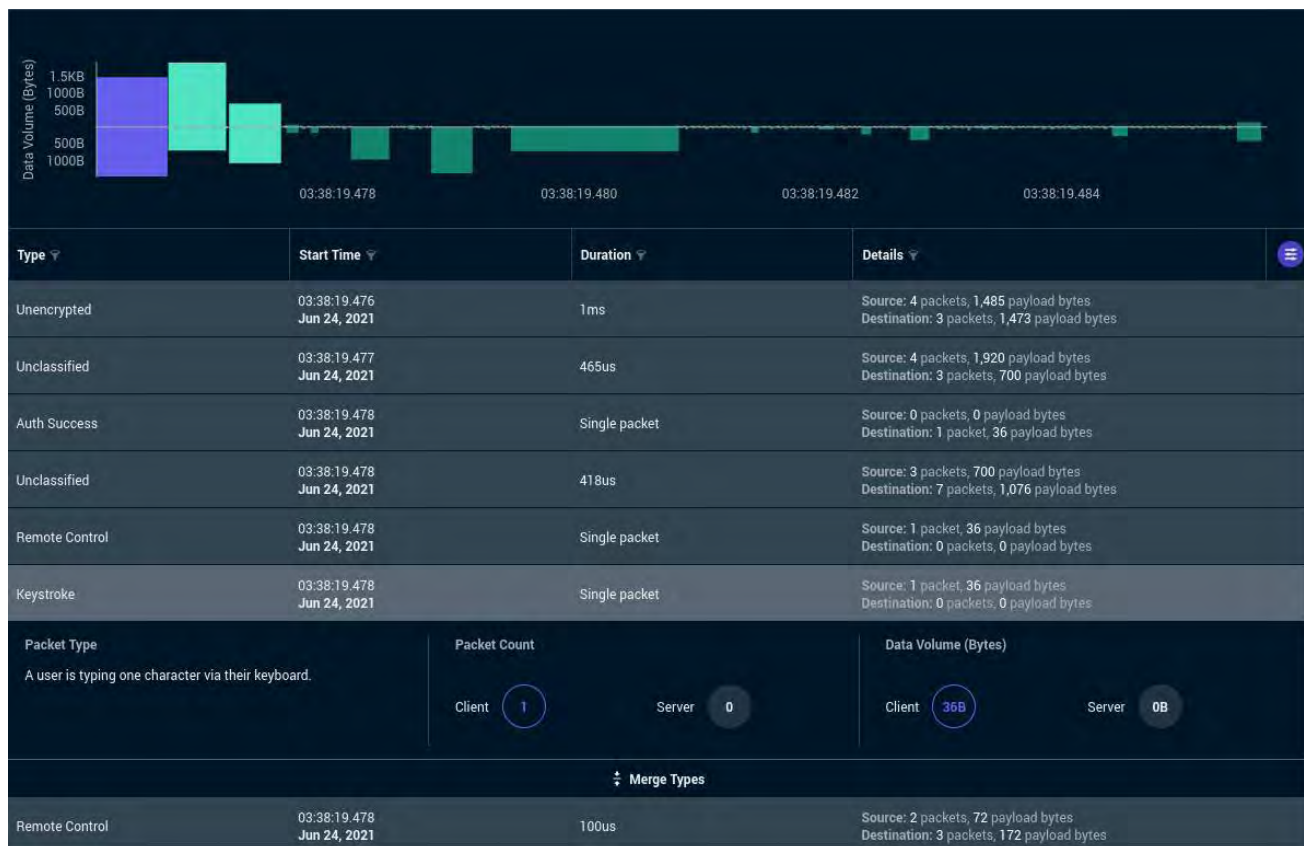


Figure 7: AVA flushes out the entire scope of an attack, enabling a decisive and rapid response

Another instance of AVA driving value for the security team is the ability to autonomously pull open-source and threat intelligence and thus, contextualize a potential threat uncovered in the environment. For example, when confronted with a suspect domain or IP address, much like an experienced security expert would, AVA pre-computes answers to questions such as:

- What other domains or destinations first showed up on the network at approximately the same time as the initial suspect domain?
- Did other devices attempt to connect to any of the same domains?
- Was there any trace of lateral movement activity beyond the initial victim?

Answering questions like these requires analyses of both internal data sources and external information via search engines, threat, and vulnerability databases, etc. AVA analyzes these results using natural language processing techniques such as entity extraction and topic modeling. For the operator, the benefit is that AVA helps uncover all potential victims within the organization and different parts of the attacker infrastructure, e.g., multiple command and control domains and IPs, all on a single screen.

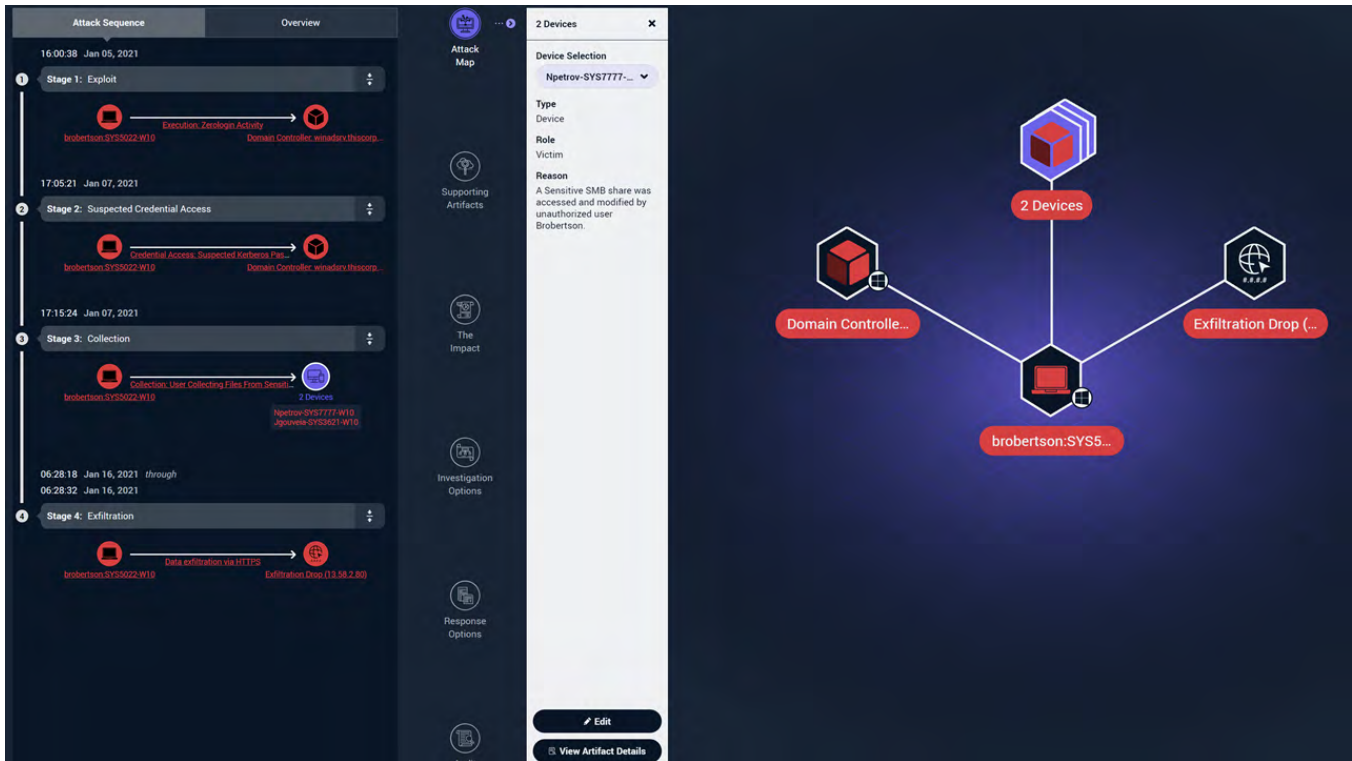


Figure 8: AVA flushes out the entire scope of an attack, enabling decisive and rapid response

**Case Study 6 – Quality of Experience (QoE)**

AVA helps provide the network operator with a clear view of the root causes of poor user experience and what remedial actions can be taken to improve that experience. This day 0 capability analyzes real-time data with the benefit of lab-trained models that understand the causes of network QoE issues, their interaction, and their effects. As a specific example, AVA employs a support vector machine (SVM) classifier to determine the performance of voice and video collaboration applications. The model is trained using a vast library of voice and video call flows labeled as a good or bad experience.

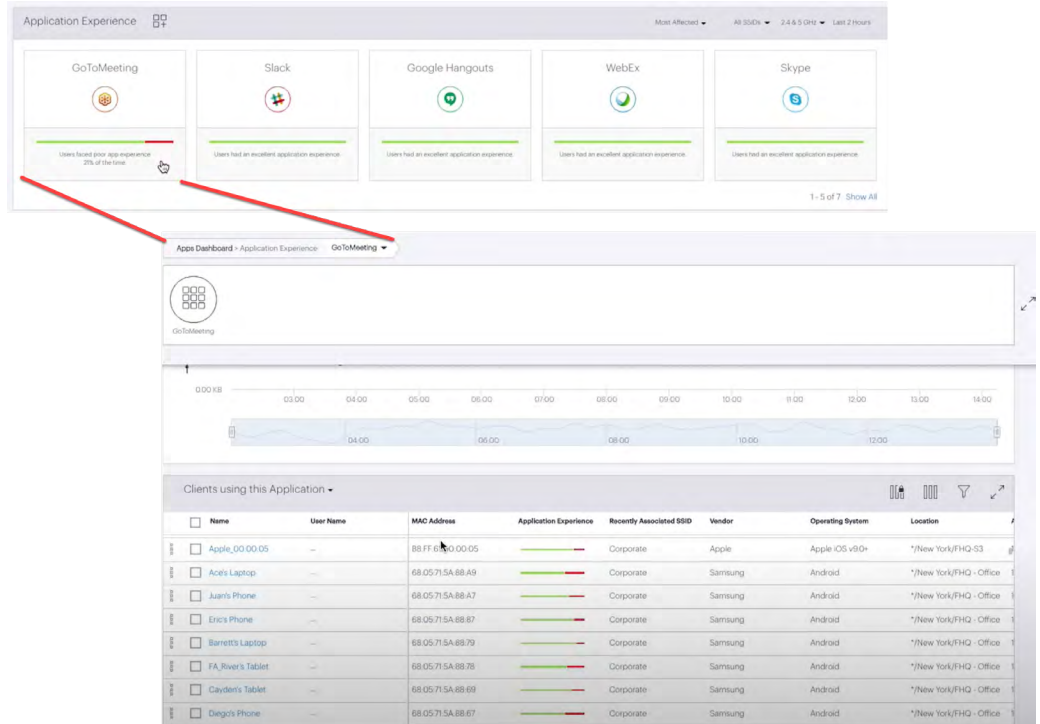


Figure 9: AVA proactively surfaces application quality of experience issues and recommends fixes

**Case Study 7 – IoT Observability and Security**

Detecting unknown IoT devices on the network offers a great example of how AVA automates human expertise.

A human expert intuitively describes an IoT device as one that most often doesn't have a browser, doesn't use enterprise protocols like SMB and Kerberos, and typically communicates with a small set of destinations. Using the information in NetDL, AVA can infer and index properties like these for the devices on the network, thereby easily highlighting the IoT devices. AVA then goes further by using recommendation systems algorithms to tag other IoT devices that are not captured by the originally encoded human intuition. This iterative approach is both quick and comprehensive at the task of eliminating an important blind spot for customers dealing with an explosion of devices on the network.

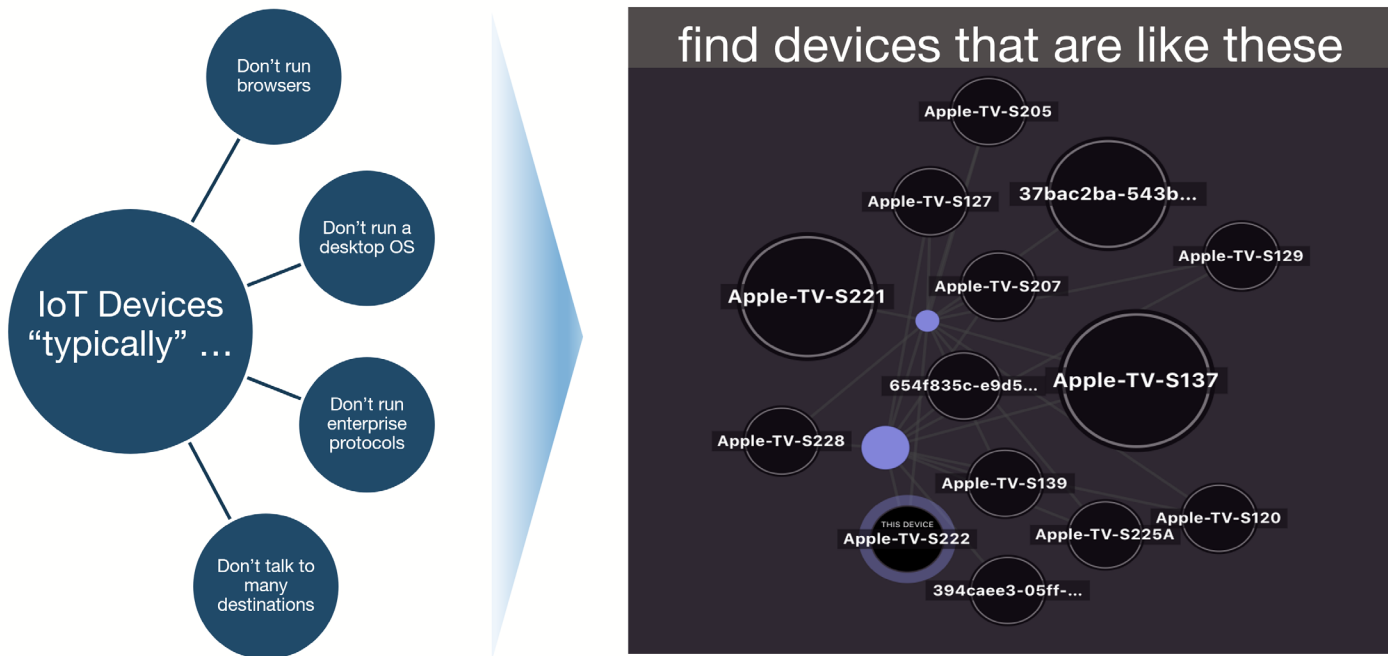


Figure 10: AVA encodes human intuition to discover IoT devices on the network

Consider the following real-world customer example which illustrates the benefits of this IoT observability. An IoT device plugged into a critical device's USB port was being used to intercept keystrokes between the keyboard and the computer. Detecting this threat first required identifying the shadow IoT device. As Figure 9 shows, the device had been sending encrypted email and a proprietary UDP stream to locations in Germany and Malaysia. These "weak signals" added further conviction, allowing AVA to trigger a device quarantine via CloudVision AGNI, Arista's NAC solution.

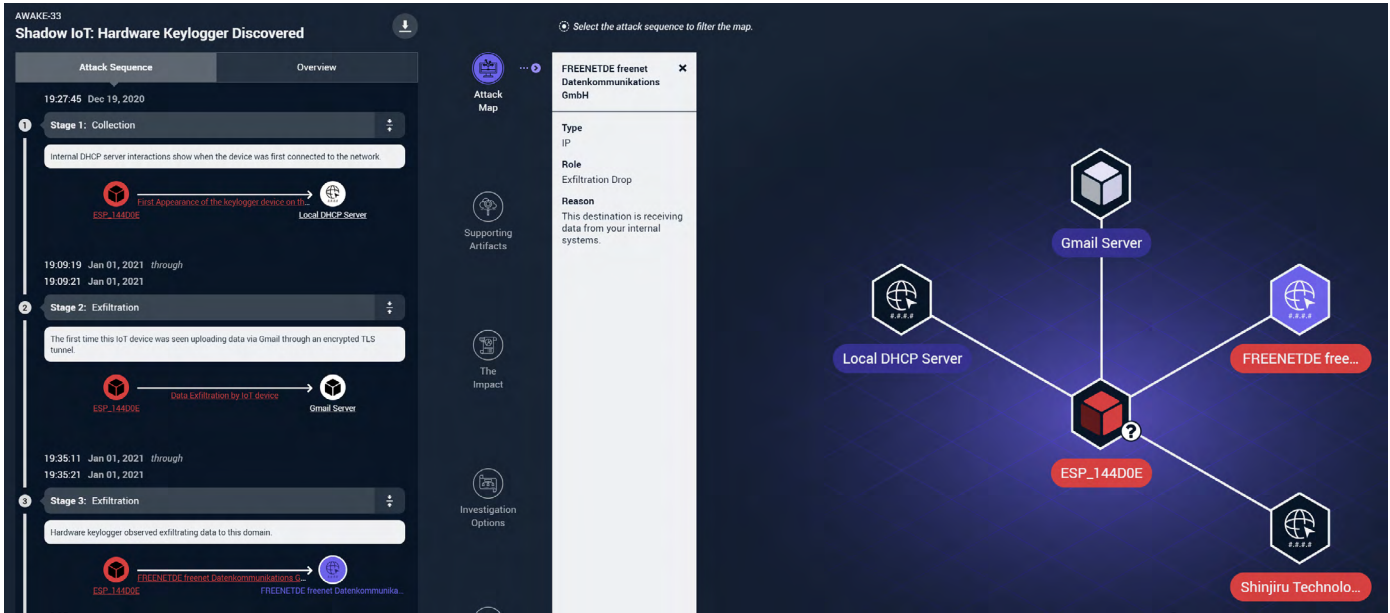


Figure 11: Discovering and remediating rogue IoT devices on the network

## Case Study 8 – Enriching the Ecosystem with AVA Insights

AVA insights can also be used to enrich other parts of the customers’ IT and security ecosystem. For example, with one click, analysts using log aggregation and SIEM tools such as Splunk or Azure Sentinel can pivot from a meaningless IP address in those tools to an AVA-enriched profile that includes the name of the device, its primary user, applications running on it and other similar devices, as well as a forensic timeline of device activities. In addition, the analyst has a detailed timeline of that device’s activities and can therefore make appropriate risk management decisions.

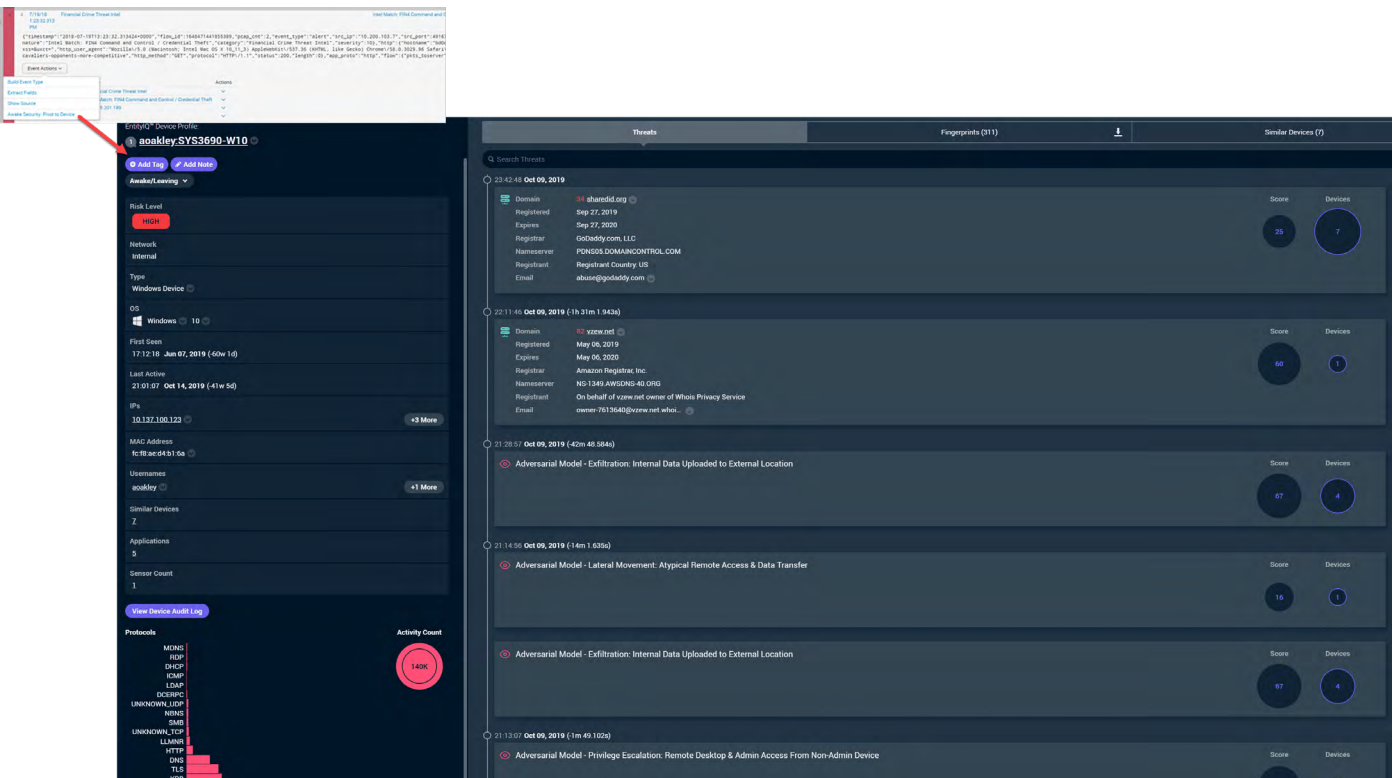


Figure 12: Enriching Splunk with AVA context

## Summary: Data-Driven Networking Made Possible

The combination of Arista AVA and EOS NetDL provides predictive and prescriptive intelligence for data-driven networks. AI/ML enrichment and analytics in combination with a broad ecosystem of vendors/partners, deliver market and customer-specific security, application, and network performance analysis, feeding continuous awareness and assurance. This provides a single source of truth and a decision-support architecture for Arista customers that ultimately delivers better business outcomes.

### **Santa Clara—Corporate Headquarters**

5453 Great America Parkway,  
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: [info@arista.com](mailto:info@arista.com)

### **Ireland—International Headquarters**

3130 Atlantic Avenue  
Westpark Business Campus  
Shannon, Co. Clare  
Ireland

### **Vancouver—R&D Office**

9200 Glenlyon Pkwy, Unit 300  
Burnaby, British Columbia  
Canada V5J 5J8

### **San Francisco—R&D and Sales Office 1390**

Market Street, Suite 800  
San Francisco, CA 94102

### **India—R&D Office**

Global Tech Park, Tower A, 11th Floor  
Marathahalli Outer Ring Road  
Devarabeesanahalli Village, Varthur Hobli  
Bangalore, India 560103

### **Singapore—APAC Administrative Office**

9 Temasek Boulevard  
#29-01, Suntec Tower Two  
Singapore 038989

### **Nashua—R&D Office**

10 Tara Boulevard  
Nashua, NH 03062



Copyright © 2025 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. December 10, 2025 02-0114-03