Arista Zero Trust Security for Cloud Networking

In 2020 the National Institute of Standards and Technology (NIST) defined the journey towards a Zero Trust Architecture (ZTA). Zero trust is not a single architecture but rather a framework that includes guiding principles for workflow, system design and operations, all designed to improve the security posture of a modern network. Arista is trusted by the world's largest data centers and cloud providers for the quality, security and performance of its products. The experience gained from working with thousands of customers has helped refine Arista's Zero Trust Security Framework across cloud networks. By leveraging a zero trust framework patterned off of the NIST and applying cloud networking principles, this architecture provides a solution to secure all digital assets across campus, data center, IoT and cloud from threats such as ransomware, data theft and supply chain compromises.

This white paper discusses the Arista Zero Trust Security architecture that delivers situational awareness, segmentation and enforcement and continuous diagnostic & monitoring that are key to effective defense against today's complex threats.

Table of contents

Introducing Arista Zero Trust Security 3

1. Situational Awareness and Network Visibility 3

1.1 Understanding Connected Endpoints with CloudVision®, Awake Security and Arista Partners 4

1.1.1 CloudVision Device Analyzer 4

1.1.2 CloudVision WiFi 4

1.1.3 Awake EntityIQ 5

1.1.4 Third-Party NAC 6

1.2 Network Switch Visibility with CloudVision 6

1.2.1 Network Compliance 7

Flow Analysis 7

2.0 Supporting Many forms of Segmentation 8

2.1 Arista Macro-Segmentation Service® (MSS®) 8

2.1.2 Multi-Domain MSS-Group Service 9

2.1.3 MSS Firewall Service 10

2.1.4 MSS Host Services 11

3. AI-Driven Continuous Detection & Monitoring 12

3.1 Awake Nucleus and Adversarial Modeling™ 12

3.2 Awake's Ava™ 12

3.3 Awake Third-Party Integration 13

3.4. Deploying a Monitoring Fabric with DANZ Monitoring Fabric™ 13

4. Incident Response Security Expertise 14

Conclusion 15

Introducing Arista Zero Trust Security

Arista's Zero Trust Security overcomes the concept of a trusted inside network that connects to an untrusted network through a traditional firewall. It eliminates the implicit trust associated with network location and instead places the onus on continuously monitoring all device and application access for mal-intent and then responding quickly. The Arista Zero Trust model is patterned on the NIST guidance laid out in 800-207 framework and leverages three fundamental pillars:

- Situational awareness to understand all the resources on the network
- Enforcement to implement zero trust access policies
- Continuous diagnostics and monitoring to uncover mal-intent and observe the ongoing network footprint of business processes

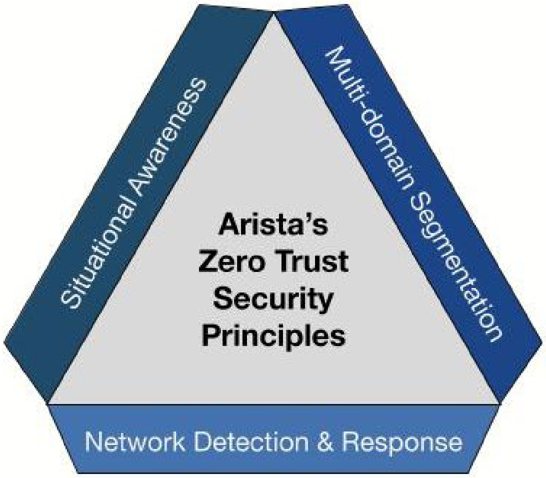
Situational Awareness
Multi-domain Segmentation
Arista's Zero Trust Security Principles
Network Detection & Response

Zero Trust Framework

- Situational Awareness → Visibility of all assets & workloads
- Enforcement → Restrict access to required connections
- Continuous Monitoring → Never trust continually verify

Zero Trust Assumes

- Successful authentication is not sufficient
- Device and User visibility beyond simple fingerprinting
- Client agents are not possible with many IoT devices

Adapting each pillar is dependent on the security administrator's specific requirements and leverages the integration with Arista partners, a portfolio of best of breed Arista switches, CloudVision network automation and telemetry, the Awake network detection and response (NDR) platform and the Arista and DANZ Monitoring Fabric, DMF. The Arista solution uses open standards and recognizes that all networks are composed of products from multiple vendors. Three building blocks define the Arista Zero Trust Security framework.

1. Situational Awareness and Network Visibility

Do you deeply understand all the connected assets in your network? A fundamental tenet of the NIST ZTA Architecture is the visibility of all resources and processes. Resources are defined broadly and include all data sources, computing services, users and IoT devices. Each resource may have state that could include things like software version, location, time/date, observed behavior, device analytics and more. Therefore the first step in migrating to zero trust model requires an organization to have detailed knowledge of its resources, privileges, and business processes. This knowledge is used for defining access privilege policy and enforcing those policies through segmentation or some other means.

Consider the example of segmentation. A group-based model places endpoints into behavioral groups and policies are defined that regulate communication both between groups and within groups. A limited number of groups such as "production", "preproduction" and "public" may be sufficient to secure many networks while other networks may require many groups for highly granular segmentation. Independent of the number of groups, segmentation begins with the need to understand connected endpoints and communication patterns.

Similarly, endpoints or network infrastructure that are susceptible to known defects such as those identified by known software defects or vulnerabilities such as PSIRTs (Product Security Incident Response Team) may also need to be identified and grouped into a high-risk group. A situational awareness strategy includes understanding the vulnerability of network infrastructure as well as the posture of the endpoint. Similar to device identification, the level of granularity required to analyze an endpoint's posture will vary based on specific customer requirements but could include, for example, OS version, connected peripherals such as removable storage devices, processes running, applications installed, memory utilization and more.

There are many components to situational awareness and not all components are needed for all customers. Arista provides a variety of visibility technologies for situational awareness, has established strategic partnerships with other vendors and pursues open standards to ensure interoperability.

1.1 Understanding Connected Endpoints with CloudVision, Awake Security and Arista Partners

Profiling or authenticating devices using network-based analytics is particularly important for IoT devices that do not leverage 802.1X or other agent-based technologies for device authentication. Users and other devices that support 802.1X leverage certificates and credentials to authenticate devices as part of the connection process. A combination of agent-based techniques and network-based techniques are required for device identification for most security architectures.

1.1.1 CloudVision Device Analyzer

Arista CloudVision Device Analyzer profiles connected endpoints using DHCP classification information. The Device Analyzer screenshot below shows 877 endpoints on the network. Devices are placed into various groups such as Android, phones, tablets, Amazon Echos etc along with information on their IP / MAC address and connected switch.

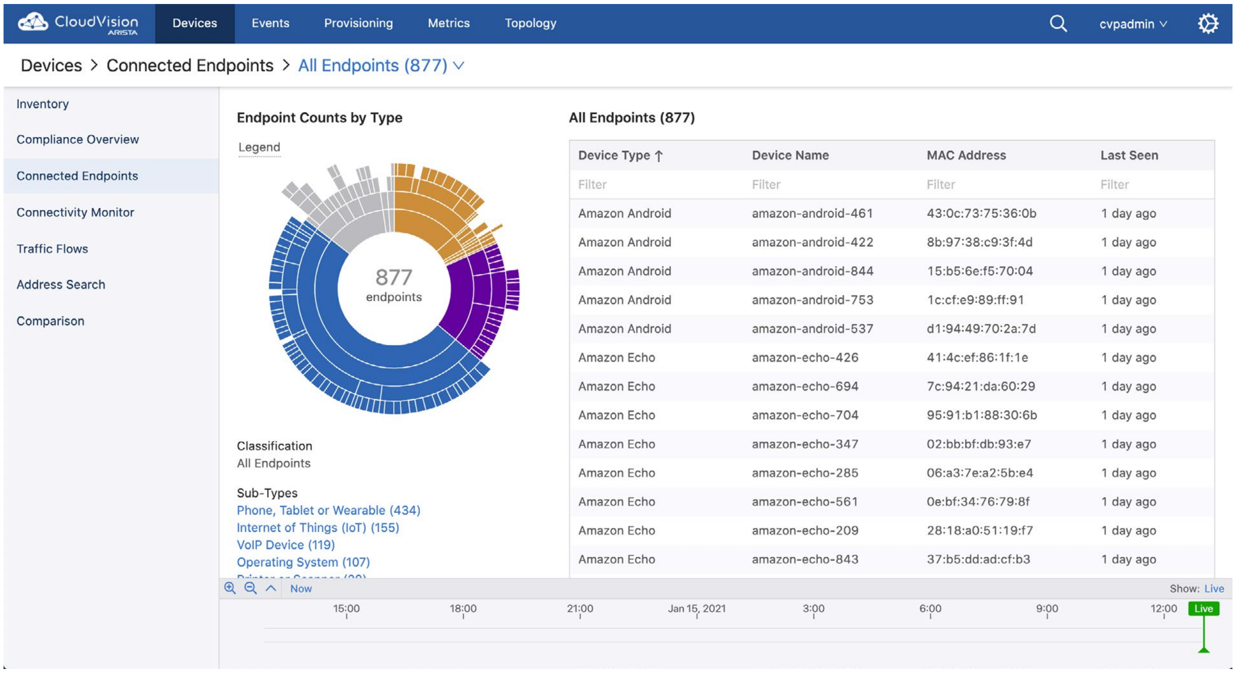

Classification
All Endpoints

Sub-Types
Phone, Tablet or Wearable (434)
Internet of Things (IoT) (155)
VoIP Device (119)
Operating System (107)

All Endpoints (877)

Device Type ↑	Device Name	MAC Address	Last Seen
Filter	Filter	Filter	Filter
Amazon Android	amazon-android-461	43:0c:73:75:36:0b	1 day ago
Amazon Android	amazon-android-422	8b:97:38:c9:3f:4d	1 day ago
Amazon Android	amazon-android-844	15:b5:6e:f5:70:04	1 day ago
Amazon Android	amazon-android-753	1c:cf:e9:89:ff:91	1 day ago
Amazon Android	amazon-android-537	d1:94:49:70:2a:7d	1 day ago
Amazon Echo	amazon-echo-426	41:4c:ef:86:1f:1e	1 day ago
Amazon Echo	amazon-echo-694	7c:94:21:da:60:29	1 day ago
Amazon Echo	amazon-echo-704	95:91:b1:88:30:6b	1 day ago
Amazon Echo	amazon-echo-347	02:bb:bf:db:93:e7	1 day ago
Amazon Echo	amazon-echo-285	06:a3:7e:a2:5b:e4	1 day ago
Amazon Echo	amazon-echo-561	0e:bf:34:76:79:8f	1 day ago
Amazon Echo	amazon-echo-209	28:18:a0:51:19:f7	1 day ago
Amazon Echo	amazon-echo-843	37:b5:dd:ad:cf:b3	1 day ago

1.1.2 CloudVision WiFi

CloudVision WiFi can also be used for wireless devices. CloudVision Wifi leverages packet information to determine not only the type of device but the applications in use as well. The CloudVision WiFi screenshot below shows various applications for example Amazon, Instagram, etc, the applications are classified into various categories such as web services, social networking, etc. With this kind of telemetry, CloudVision WiFi can immediately identify what endpoints are running a particular application as shown below.

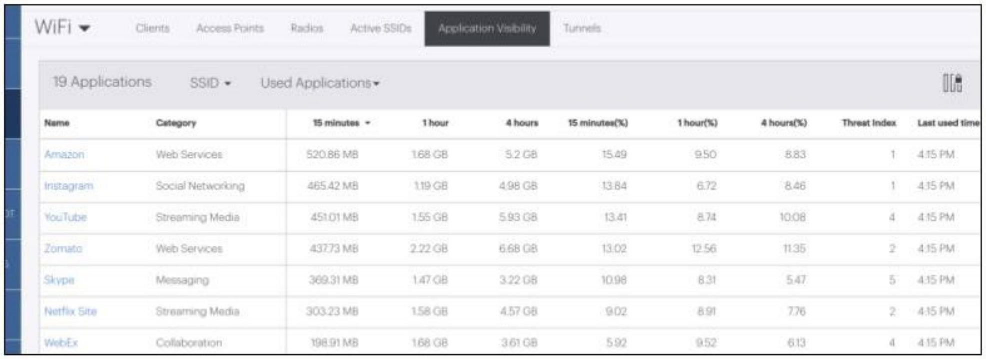

19 Applications SSID Used Applications

Name	Category	15 minutes	1 hour	4 hours	15 minutes(%)	1 hour(%)	4 hours(%)	Threat Index	Last used time
Amazon	Web Services	520.86 MB	1.68 GB	5.2 GB	15.49	9.50	8.83	1	4:15 PM
Instagram	Social Networking	465.42 MB	1.19 GB	4.98 GB	13.84	6.72	8.46	1	4:15 PM
YouTube	Streaming Media	451.01 MB	1.55 GB	5.93 GB	13.41	8.74	10.08	4	4:15 PM
Zomato	Web Services	437.73 MB	2.22 GB	6.68 GB	13.02	12.56	11.35	2	4:15 PM
Skype	Messaging	369.31 MB	1.47 GB	3.22 GB	10.98	8.31	5.47	5	4:15 PM
Netflix Site	Streaming Media	303.23 MB	1.58 GB	4.57 GB	9.02	8.91	7.76	2	4:15 PM
WebEx	Collaboration	198.91 MB	1.68 GB	3.61 GB	5.92	9.52	6.13	4	4:15 PM

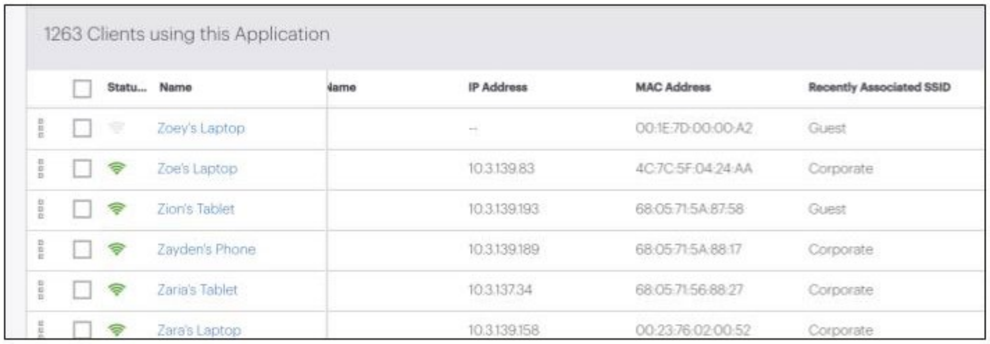

1263 Clients using this Application

☐	Statu...	Name	Name	IP Address	MAC Address	Recently Associated SSID
☐		Zoey's Laptop		--	00:1E:7D:00:00:A2	Guest
☐		Zoe's Laptop		10.3.139.83	4C:7C:5F:04:24:AA	Corporate
☐		Zion's Tablet		10.3.139.193	68:05:71:5A:87:58	Guest
☐		Zayden's Phone		10.3.139.189	68:05:71:5A:88:17	Corporate
☐		Zaria's Tablet		10.3.137.34	68:05:71:56:88:27	Corporate
☐		Zara's Laptop		10.3.139.158	00:23:76:02:00:52	Corporate

1.1.3 Arista Awake EntityIQ

Arista Awake EntityIQ provides behavioral device identification via an AI-based security knowledge graph that identifies, profiles and tracks devices, users and applications on an enterprise network with just a network connection and without the need for agents. Full packets are mirrored to the Awake Security Platform and analyzed along several axes beyond traditional IP address lookups. Devices are grouped into peer groups based on common behaviors and tracked as they move across the network and beyond, even as IP addresses change. For instance, TLS headers are analyzed for encrypted traffic; protocols like SMB and Kerberos are deeply parsed to identify devices and users; and DHCP / DNS transactions can be monitored to identify everything on the network- from IoT devices to operational technology. As shown in the screenshot below, EntityIQ is able to determine the device is a Windows 10 device, used primarily by aoakley and has had three different IP addresses in the recent past.

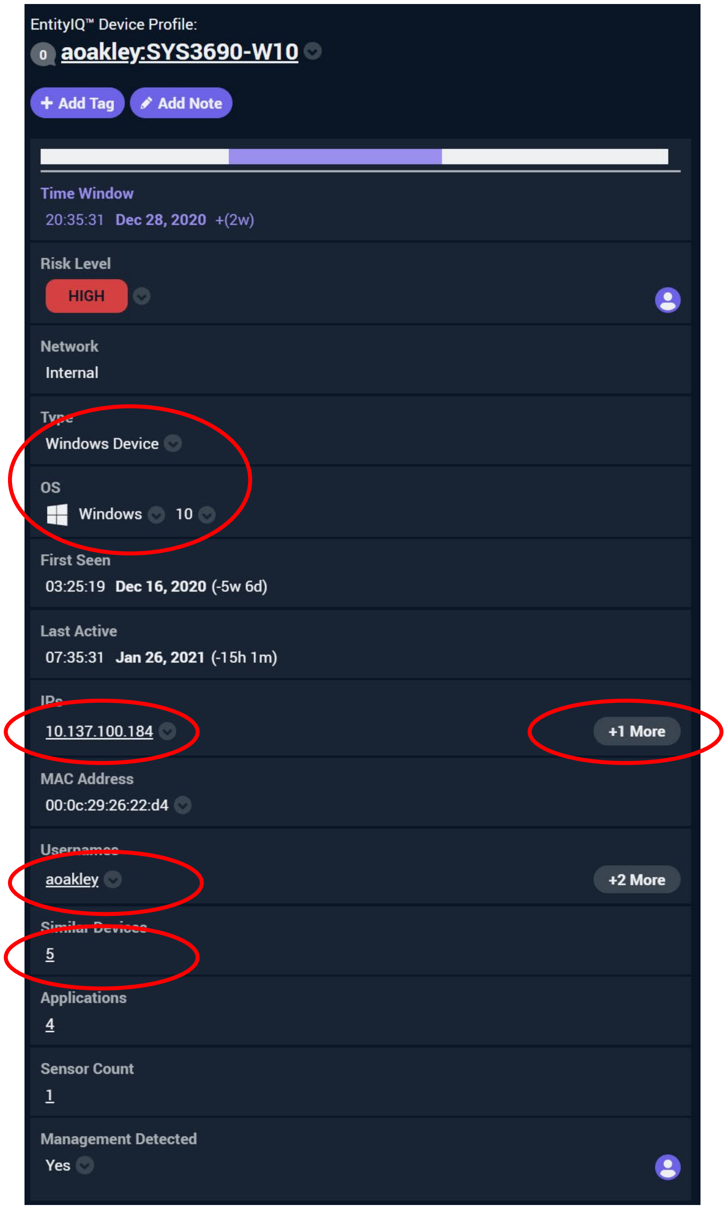

EntityIQ™ Device Profile:

0 aoakley:SYS3690-W10

Time Window
20:35:31 Dec 28, 2020 +(2w)

Risk Level
HIGH

Network
Internal

Type
Windows Device

OS
Windows 10

First Seen
03:25:19 Dec 16, 2020 (-5w 6d)

Last Active
07:35:31 Jan 26, 2021 (-15h 1m)

IPs
10.137.100.184

MAC Address
00:0c:29:26:22:d4

Usernames
aoakley

Similar Devices
5

Applications
4

Sensor Count
1

Management Detected
Yes

1.1.4 Third Party NAC

Traditional NAC products can also be used to identify and classify connected endpoints. Arista is interoperable with all the leading NAC providers including Aruba ClearPass, Cisco ISE and Forescout. A NAC product is also a RADIUS server that is responsible for authenticating devices through 802.1X or MAC Based Authentication (MBA). Authenticating IoT devices are frequently reliant on fingerprinting techniques such as DHCP, DNS, user-agent and SNMP.

1.2 Network Switch Visibility with CloudVision

In addition to profiling and classifying connected endpoints through Device Analyzer, CloudVision provides visibility into switch performance, network compliance and flow analytic processing.

1.2.1 Network Compliance

Ensuring that networking devices are not vulnerable to industry PSIRT advisories is critical for a secure network. This is a key component of the trust algorithm in a ZTA. CloudVision provides a simple compliance dashboard that reports observed PSIRT security advisories within the network. The dashboard also reports any known exposures to software defects that are relevant as well as out-of-band non-sanctioned configuration changes made to the switches under management.

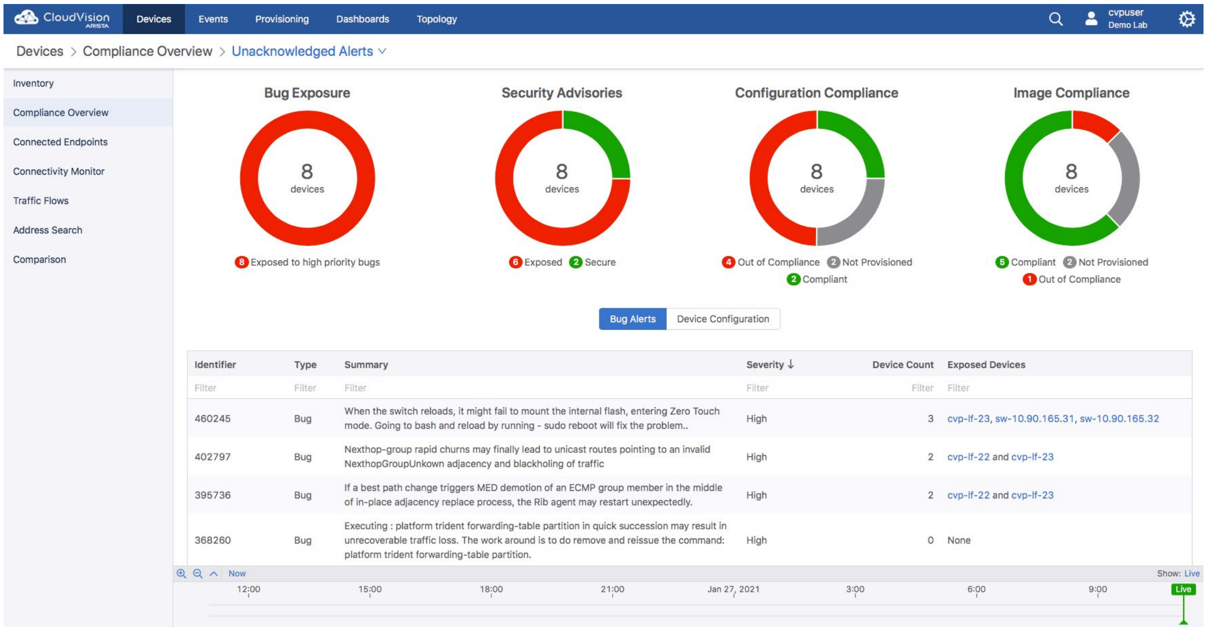

Identifier	Type	Summary	Severity ↓	Device Count	Exposed Devices
Filter	Filter	Filter	Filter	Filter	Filter
460245	Bug	When the switch reloads, it might fail to mount the internal flash, entering Zero Touch mode. Going to bash and reload by running - sudo reboot will fix the problem..	High	3	cvp-lf-23, sw-10.90.165.31, sw-10.90.165.32
402797	Bug	Nexthop-group rapid churns may finally lead to unicast routes pointing to an invalid NexthopGroupUnkown adjacency and blackholing of traffic	High	2	cvp-lf-22 and cvp-lf-23
395736	Bug	If a best path change triggers MED demotion of an ECMP group member in the middle of in-place adjacency replace process, the Rib agent may restart unexpectedly.	High	2	cvp-lf-22 and cvp-lf-23
368260	Bug	Executing : platform trident forwarding-table partition in quick succession may result in unrecoverable traffic loss. The work around is to do remove and reissue the command: platform trident forwarding-table partition.	High	0	None

Flow Analysis

CloudVision can analyze sampled or non sampled flow information received via SFLOW or IPFIX. CloudVision provides the ability to query conversations of who is talking to whom and flow traffic patterns. This information is useful to understand business processes while designing the ZTA as well as to monitor the network on an ongoing basis, something we cover in detail in the next section of this paper. The figure below is a simple query showing the top 20 hosts and port numbers used. Other queries include the amount of traffic for specific host-destination pairs.

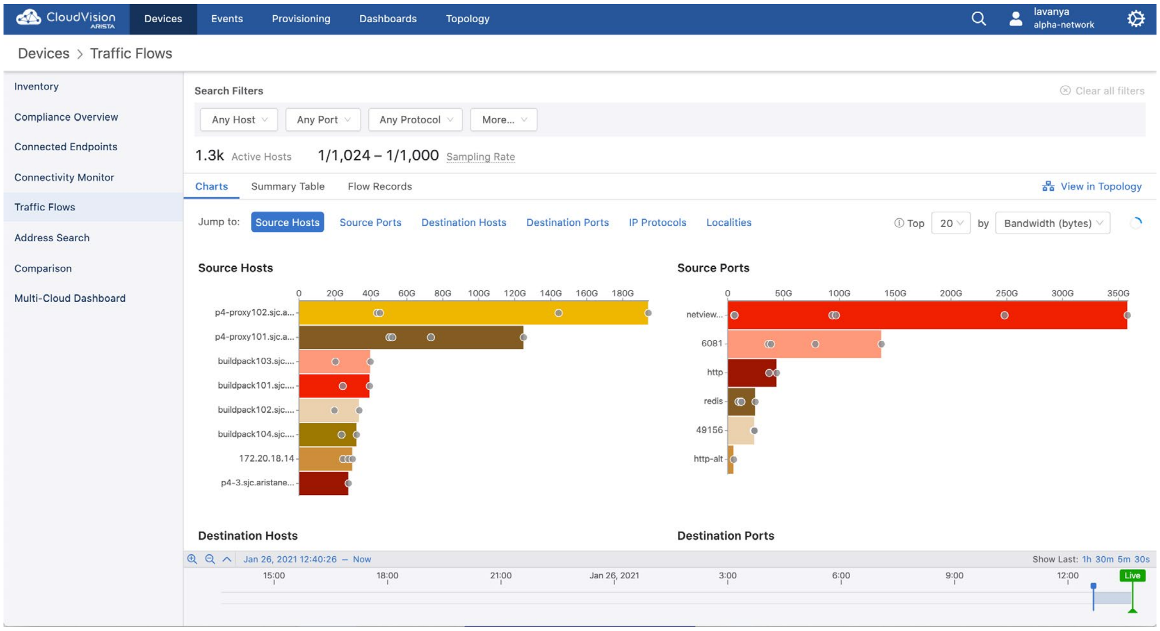

2.0 Supporting Many Forms of Segmentation

How do you control lateral movement and constrain malware proliferation? The second aspect of Zero Trust Security requires policy enforcement controls to ensure access is only provided based on runtime decisions made by the trust algorithm. Arista supports a variety of segmentation controls for multi domains across clients to DC to campus to cloud networks.

The historical and simplest form of segmentation are port-based Access Control Lists (PACLs), VLANs or VXLANs with Routed ACLs (RACLS) and VRFs to regulate client communications. Of course, firewalls are used at the network edge for protection from external connections as well as in the datacenter for protecting east-west connections. In other words, the concepts of segmentation are not new as administrators have been grouping clients and workloads into VLAN or VRF segments since the early 1990s.

Familiar segmentation methods such as VLANs and VRFs continue to be sufficient for many networks that only require a few segmentation zones. In Arista's Zero Trust model, administrators require more granular and dynamic segmentation groupings that are as close to the user or device as feasible. With more segments, it is operationally challenging to only associate security groups with IP addresses. For example in a VLAN or VRF architecture adding a new segment may require dividing an existing subnet into two smaller subnets and re-IPing the devices that were connected to the old subnet. Such network changes are cumbersome and disruptive. Standard PACLs have their own challenges as well, specifically with hardware resource scale in TCAM (Ternary Content-Addressable Memory).

Granular segmentation of campus and datacenter networks in a zero trust framework requires a different approach especially with the growth of IoT and OT. Fundamentally, security segmentation groups need to be defined independent of network IP constructs. For example, to protect the organization from the well publicized Mirai botnet, an administrator might want to define a group for security cameras and a different group for the networked digital video recorders (DVRs) and yet another one for the physical security administrators. A camera, per policy, should only be allowed to communicate with the DVR and security administrator. A camera should not be allowed to communicate with another camera even if it were on the same subnet. Because there may be multiple buildings, cameras could span multiple subnets in a classic L3 network design. Security and network administrators need to have the ability to easily define a segment and it's associated policy, that is independent of IP addressing and other network forwarding constructs.

2.1 Arista Macro-Segmentation Service (MSS)

The Arista Macro-Segmentation Service (MSS) solution set provides several leading edge segmentation options while continuing to support legacy models such as VRFs, VXLANs, and PACLs.

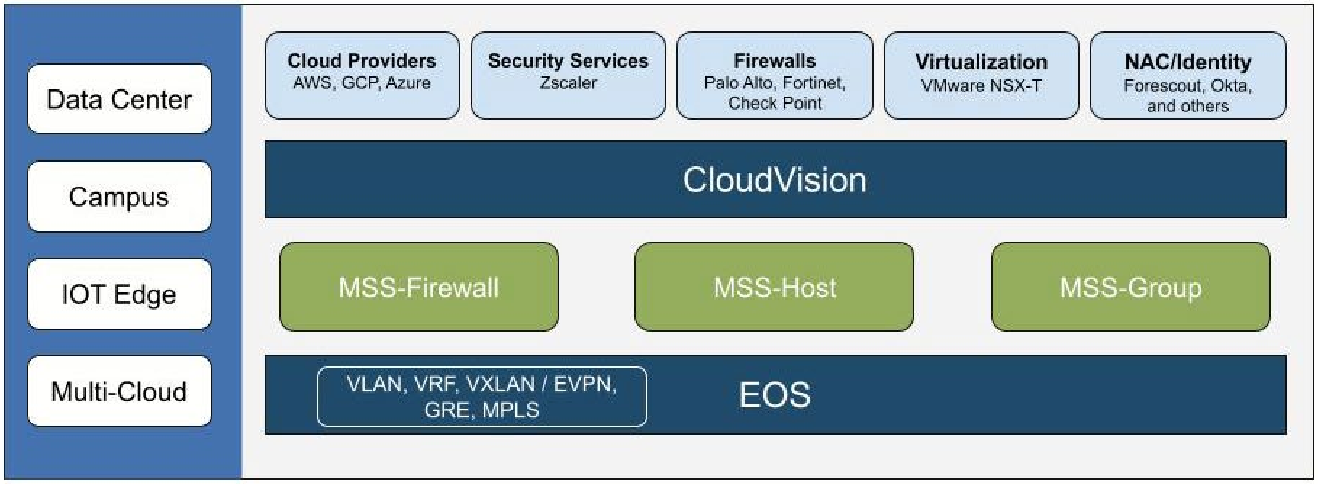

2.1.2 Multi-Domain MSS-Group Service

As part of the MSS solution set, Arista has introduced Multi-Domain MSS-Group Segmentation. MSS-Group applies authorization policies to security segment groups rather than interfaces, subnets or physical ports. IP addresses and / or IP subnets are placed into administratively defined security segment groups. Policies are applied to each group that define both inter and intra segment group communication. In the camera example mentioned above, MSS-Group enables an administrator to define a segment group called "camera" that per policy can talk to the defined segment group "DVR" but cannot communicate with anything else including each other. Because policies are defined on a per segment group basis, segmentation rules policies can be created independently from IP addressing.

The ability to create security segments and to enforce policies between segments is built into the Arista switch's hardware. The switch needs to be configured with what segment groups need to be created, with membership that defines what hosts or subnets belong to each segment and with policy that defines what other segments a given segment is allowed to communicate (including if it can communicate with other members on the same segment). The switch can be configured in a variety of ways including Arista standard CLI or EAPI which is sufficient for a few switches. For network-wide rollout, switches should be provisioned through an orchestration layer. The orchestration layer is a CloudVision function that pushes consistent configuration to all Arista switches performing MSS-Group enforcement. The orchestration layer receives group membership and policy information from a policy layer. The policy layer is a logical layer that may also be a CloudVision function. Within CloudVision static group segment policy and membership can be programmed by an administrator.

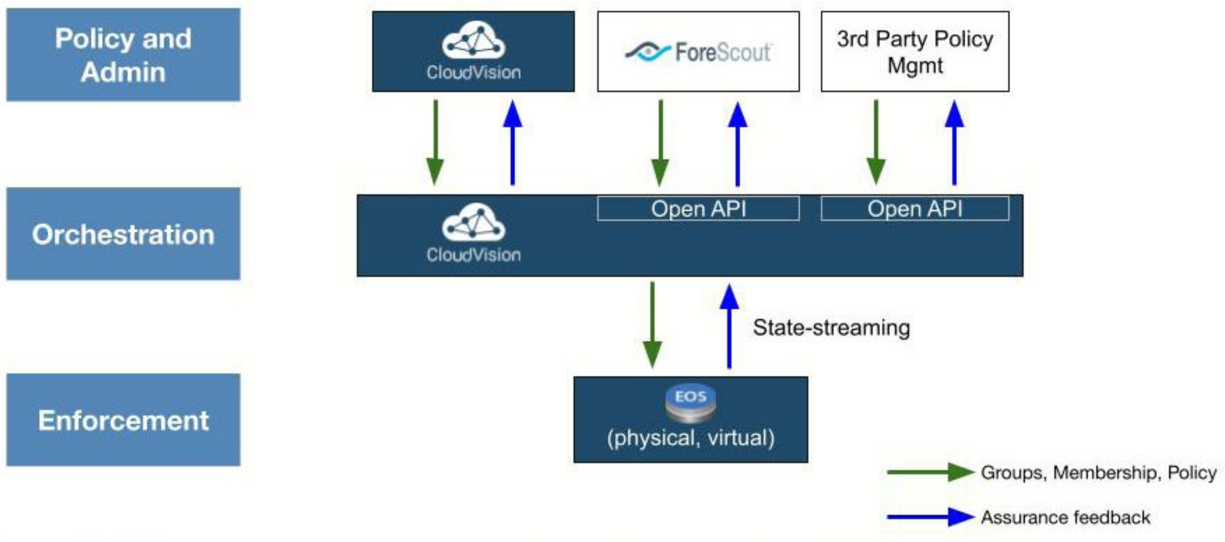

The MSS-Group solution is most powerful when CloudVision integrates with a dynamic identity layer. By leveraging APIs available in CloudVision, partners such as Forescout can ensure that different devices are put into logical groups based on device fingerprints, behavior, 802.1X authentication, and other mechanisms. The APIs allow Forescout to associate each device with its relevant security segmentation group and to apply the appropriate segmentation policies within CloudVision, which is then responsible for orchestrating the required policy to the various MSS-Group enabled switches. As new devices join the network or segmentation group membership changes, Forescout automatically updates CloudVision with these changes. In the reverse direction, CloudVision gathers hit count and segment drop information from the various switches. This information is used in CloudVision analytic reporting as well as forwarded to Forescout for Forescout reporting.

Unlike other solutions on the market, the MSS-Group segmentation architecture does not rely on proprietary ethernet tags or protocols. The upstream and downstream switches can be from any vendor. Arista MSS-Group capable switches can be deployed wherever enforcement is required. Enforcement policies can be created for any packets flowing through the switch. While it is most ideal to deploy MSS-Group at the access layer, the diagram below shows how a single switch configured with MSS-Group can be used to enforce both ends of a communication flow. Of course, additional enforcement points can be added as needed.

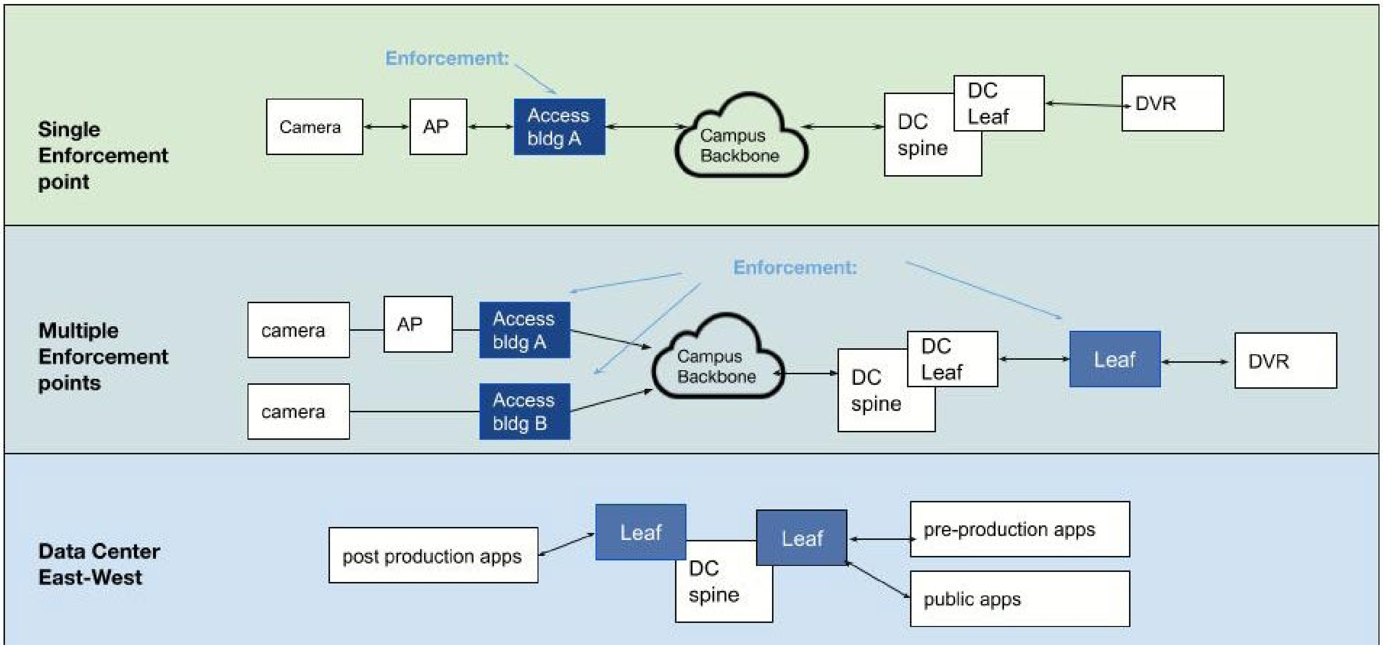

2.1.3 MSS Firewall Service

MSS Firewall is a MSS™ offering that enables an administrator to logically insert a Fortinet, Palo Alto, or Check Point firewall dynamically into the data path for traffic inspection. CloudVision connects with a supported firewall controller; an administrator defines the traffic inspection policy in the firewall controller; the firewall controller communicates the policy to CloudVision. CloudVision maintains a network-wide database of all state within the network called NetDB. NetDB is aware of where every workload is within the network, it learns in real time about new devices or workloads that are added, moved or removed from the network. Once CloudVision learns about the traffic inspection policy from the firewall controller, it leverages the switch location information in NetDB to configure the appropriate switch. The switch can be programmed by CloudVision to redirect specified traffic to the firewall or an ACL can be programmed to drop or forward selected traffic thus bypassing the firewall.

The diagram below shows how a single physical firewall can be used to inspect specific traffic from workloads that are anywhere in the network.

Large datacenters can centralize their firewalls in a service rack and insert them in the path between any workloads on-demand or based on a firewall policy. MSS Firewall uses standards-based forwarding to stitch service devices into the path of traffic, and it can fully function if the network is comprised of devices from multiple vendors.

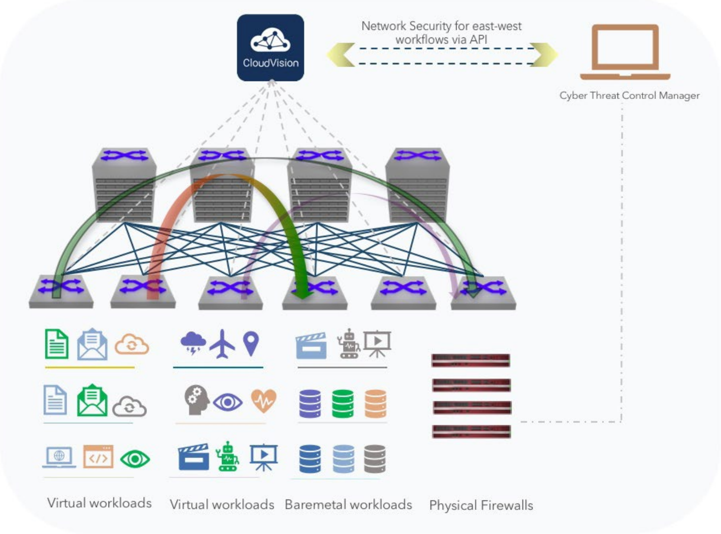

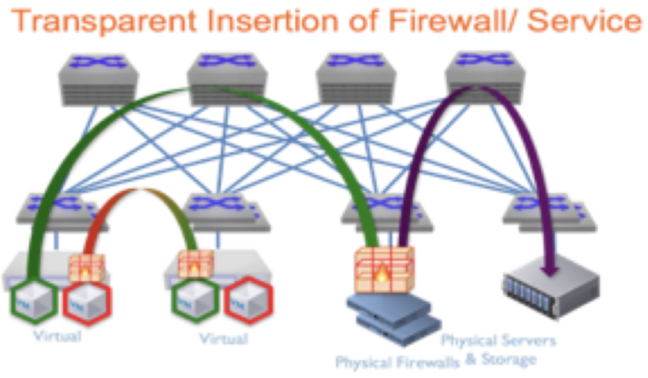
Transparent Insertion of Firewall/ Service
Virtual
Virtual
Physical Firewalls
Physical Servers & Storage

MSS Firewall was originally developed to segment traffic in the datacenter, however MSS Firewall is also being used to segment north-south traffic in the campus. In campus use cases, MSS Firewall is used to restrict traffic to secure applications and to guard against denial of service, DOS, attacks. All flows connecting to select applications are directed to the firewall service node for further inspections per firewall defined policy. Similarly, MSS Firewall can also be used to add additional inspection before a device is trusted or if the subject is deemed to be risky.

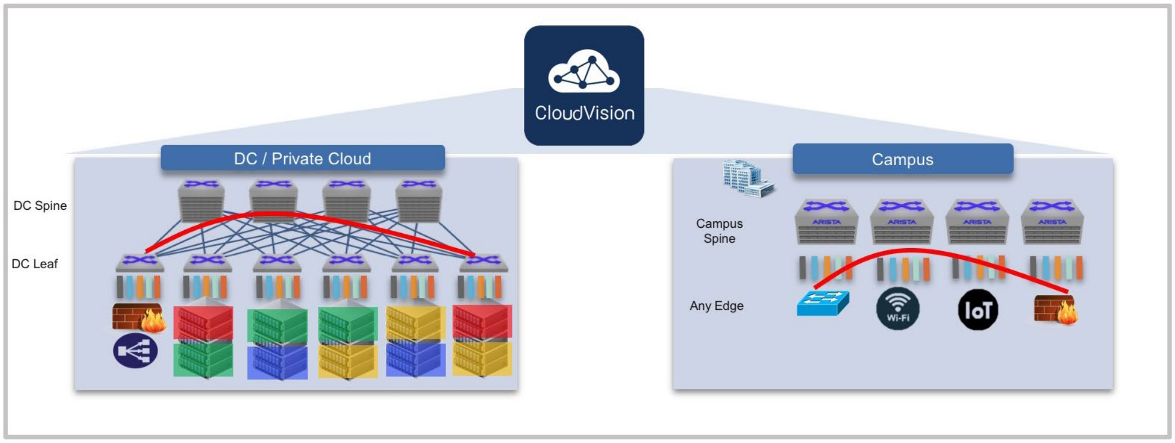

2.1.4 MSS Host Services

Arista and VMware have partnered to integrate VMware micro-segmentation technology with Arista MSS Host. The solution provides a single administrative domain to manage both VMs and physical workloads. Applying the security policy at the network edge for the physical workloads brings uniformity and consistency. In operation, Arista MSS Host will register with the VMware NSX controller and receive the policies. CloudVision will appropriately program the Arista switch or switch pairs to allow or deny conversation between the physical and virtual workloads. This allows for dynamic synchronization of security policy as new policies are created and existing policies are modified. The MSS Host and VMware solution allows enterprises to secure all assets with uniform policy implementation at scale, mitigating the overall risk and delivering agile services.

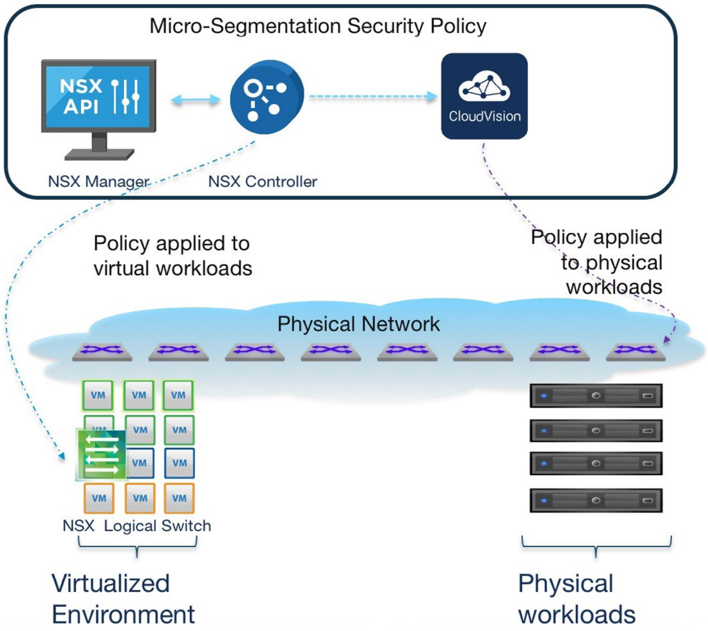

3. AI-Driven Continuous Detection & Monitoring

Your CEO is trustworthy but is his computer? Just because a device or user is on the network, does not imply it is to be trusted. The Arista zero trust architecture performs continuous monitoring to identify malicious intent originating from both outside the perimeter as well as from the inside. Threats and risk scores surfaced by the Awake Security Platform can then be used to make segmentation decisions by the enforcement controls we discussed above. For example, an at risk endpoint could be moved to the "high-risk" security group where it is given restricted access.

3.1 Awake Nucleus and Adversarial Modeling™

The Awake Security Platform is built on a foundation of deep network analysis across the campus, data center, IoT and cloud workload networks. Unlike other network detection and response (NDR) solutions, Awake parses over three thousand protocols and processes layer 2 through layer 7 data, including performing encrypted traffic analysis. As explained above, Arista Awake EntityIQ uses this information to autonomously profile entities such as devices, users and applications, while also preserving these communications for historical forensics. The Awake Nucleus then uses an ensemble of machine learning approaches to identify malicious intent hidden within.

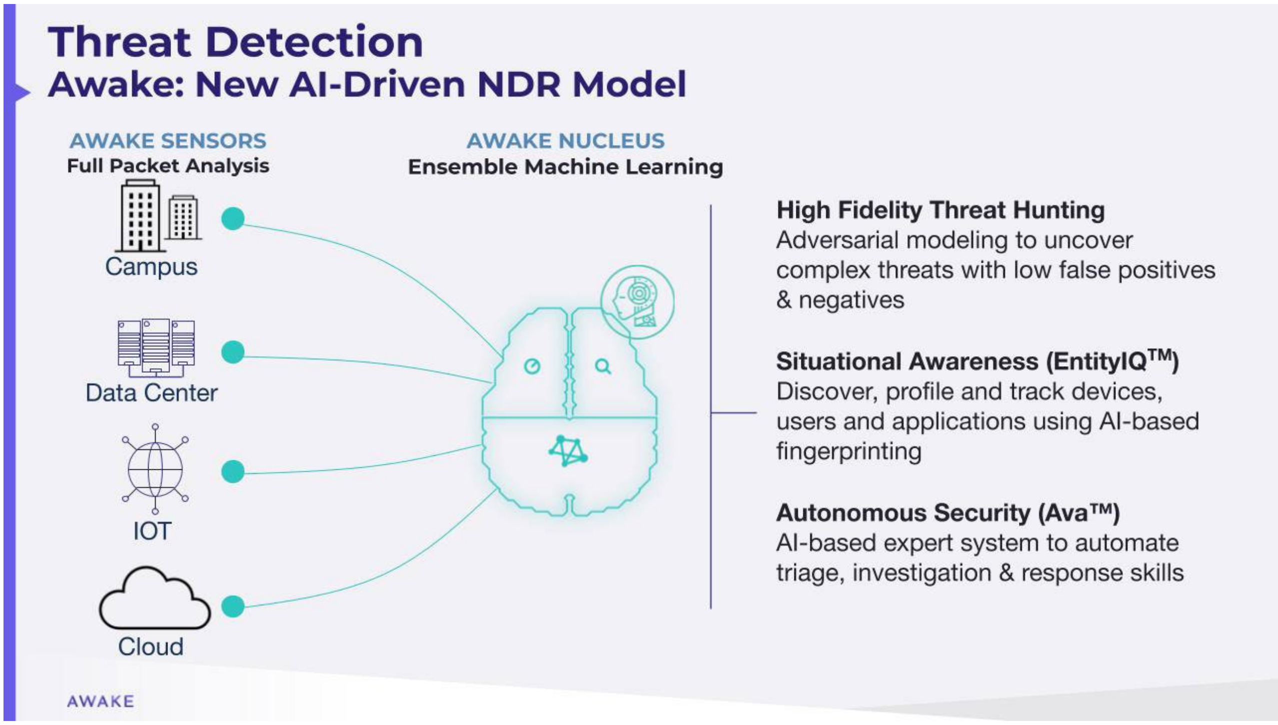

Threat Detection

Awake: New AI-Driven NDR Model

AWAKE SENSORS

Full Packet Analysis

AWAKE NUCLEUS

Ensemble Machine Learning

High Fidelity Threat Hunting

Adversarial modeling to uncover complex threats with low false positives & negatives

Situational Awareness (EntityIQ™)

Discover, profile and track devices, users and applications using AI-based fingerprinting

Autonomous Security (Ava™)

AI-based expert system to automate triage, investigation & response skills

Awake's Adversarial Modeling™ capability enables autonomous threat hunting for complex attacker tactics, techniques and procedures (TTPs), by delivering a vocabulary to express and then identify these patterns of behavior even if they occur over extended periods of time, across a variety of protocols and impact multiple network assets.

3.2 Awake's Ava™

Ava™, the world's first AI-based security expert system, performs autonomous threat hunting and incident triage. Using artificial intelligence, open source intelligence and Awake's own human expertise, Ava autonomously connects the dots across the dimensions of time, entities, and protocols, enabling the solution to present end-to-end Situations to the end user rather than atomic alerts. Analysts benefit from a decision support system that visualizes the entire scope of an attack and presents investigation and remediation options on a single screen while avoiding the effort of piecing it together themselves.

The screenshot illustrates the power of the Awake Security Platform. In this real world case study, the platform identified four Polycom devices that are acting differently from other Polycom devices. Specifically they are labelled as IP phones by Arista Awake EntityIQ and associated with in four different conference rooms. These phones are unexpectedly communicating with an AWS hosted server. Separately the AWS server is being accessed by an IT administrator ("Robert" in our anonymized case study). The Awake Security Platform stitched together these disparate events and identified a malicious activity. In this example, Robert was recording conversations via Polycom phones in four different conference rooms and uploading the recordings to an AWS server. Robert then retrieved the recording for nefarious purposes.

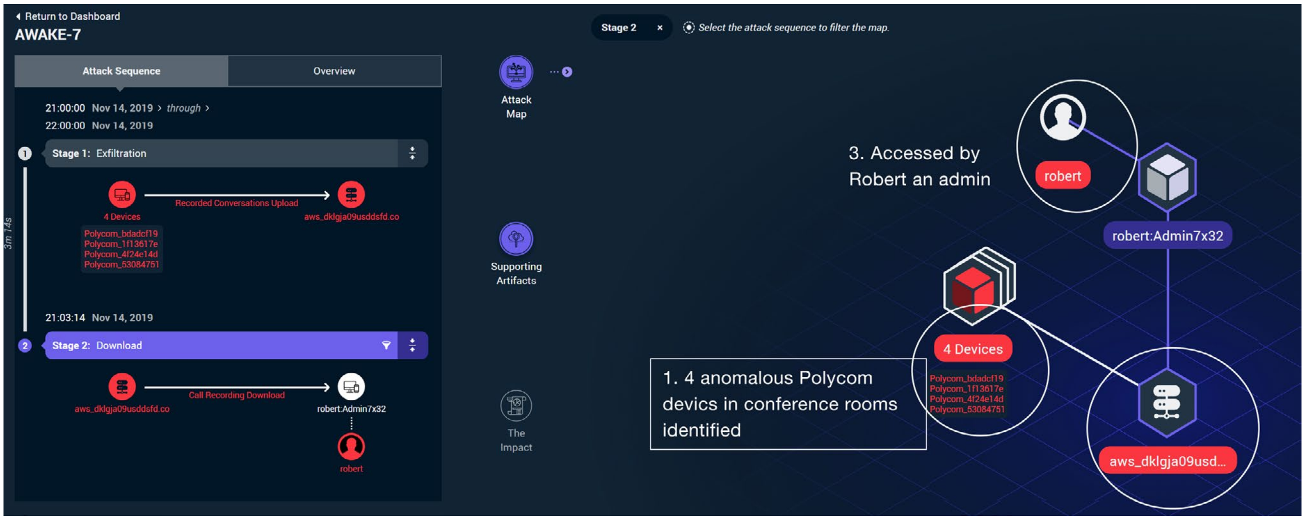

3. Accessed by Robert an admin

1. 4 anomalous Polycom devics in conference rooms identified

3.3 Awake Third-Party Integration

The Awake Security Platform also integrates with and amplifies existing solutions through integrations into industry-leading SIEM such as Splunk, business intelligence such as Microsoft Power BI, ticketing and analytics such as ServiceNow, endpoint detection such as CrowdStrike and security orchestration tools such as Palo Alto Networks' Cortex XSOAR. In addition, the platform supports a full API for custom workflows and integrations. For instance, the SIEM integration allows an analyst to pivot from an alert containing an IP or email address to a device profile with associated user(s) and roles, operating system and application details, a forensic threat timeline as well as a listing of similar device(s) for campaign analysis. Similarly, endpoint integrations allow for one click quarantining of compromised devices or retrieval of endpoint forensic data.

3.4. Deploying a Monitoring Fabric with DANZ Monitoring Fabric

Deploying a zero trust secure network requires the continuous collection and analysis of flow or packet information for situational analysis and Continuous Diagnostic and Mitigation (CDM). Traditionally, packet information is gathered through mirroring packets on a switch-by-switch basis. Network Packet Brokers (NPB) are frequently deployed but are largely proprietary and have proven to be challenging to scale for organization-wide monitoring.

DANZ Monitoring Fabric (DMF) is a next-generation NPB architected for pervasive, organization-wide visibility and security. DMF enables IT operators to pervasively monitor and mirror all traffic. DMF provides deep hop-by-hop visibility, predictive analytics and scale-out packet capture. The DMF dashboard controls the entire monitoring fabric and provides simplified network performance monitoring for real-time and historical context.

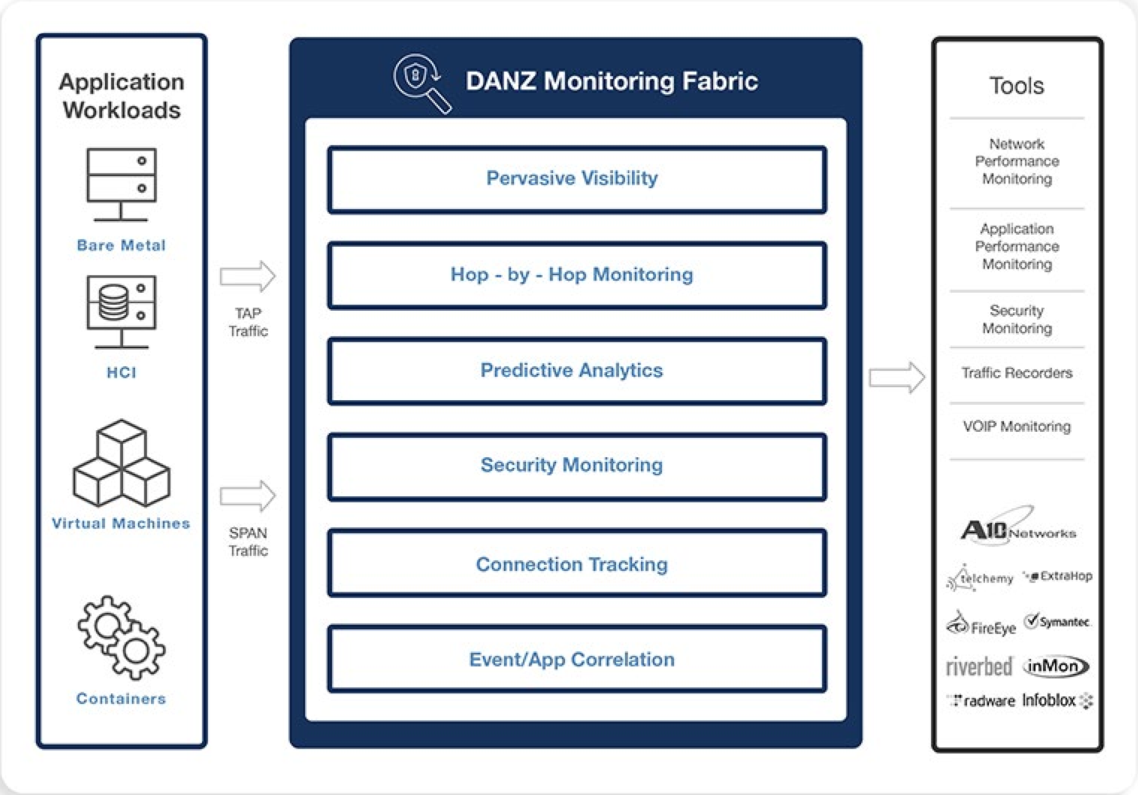

The DMF Analytics Node integrates with the DMF to provide flow analytics. The intuitive user interface quickly pinpoints suspicious traffic and anomalous behaviors. Unknown hosts, giant flows from non-business hosts, and/or traffic to web sites that have not been approved by the security and compliance teams are identified. This is especially useful to drive zero-trust decisions. Additionally, the analytics node also integrates with DMF Recorder Node to provide packet level visibility for further detailed analysis and replay.

This combination of capabilities gives a tremendous arsenal to security administrators to visualize, detect and then rectify problems in production environments.

4. Incident Response Security Expertise

While incident response is not directly a part of the NIST 800-207 framework, this is clearly a very relevant function for effective zero trust and specifically for the continuous diagnostics and mitigation. Some organizations are self-sufficient in this area, but Arista recognizes that many others want human expertise on call. Arista's Awake Security division launched its managed network detection and response (MNDR) offering with a team of individuals that have decades of experience responding to the world's most consequential breaches. Awake's MNDR solution significantly improves the maturity of your security program by delivering a comprehensive understanding of your attack surface and then monitoring as well as threat hunting across all that infrastructure whether on-premise, cloud, IoT or operational technology. The solution combines Awake's award-winning NDR technology with decades of expertise, and advanced incident response methods.

Finally, Arista can also support your organization post-breach with a variety of incident response services. It is increasingly true that the lasting damage of a breach is a consequence of how the organization responded to the breach rather than the fact that it was breached in the first place. Arista recognizes the difficulties of maintaining skilled resources, having robust investigation and response processes while also dealing with technical challenges such as unmanaged devices, IoT and cloud. These are just a few of the factors that hinder proactive preparedness for an incident. Arista's Awake Labs team can offer retainers that include pre-negotiated legal terms and rates, the right processes, both endpoint and network response technology and the expertise on demand. All of this saves time and brings in much needed expertise within minutes after a breach. This in turn helps contain impact in the event of an attack, including financial costs and reputational damage.

Conclusion

Attackers such as the perpetrators of the SolarWind attack have become more sophisticated and effective at subverting conventional malware-based threats and less reliant on traditional techniques such as phishing and exploits. In the recent case, to avoid detection, the attackers successfully injected malicious code. While the Sunburst campaign is certainly in the recent news, this is not a new methodology. In fact, today more than 50% of breaches are "malware-free". Security teams are best served by assuming the environment is in fact compromised and that the perimeter has been breached. With that mindset, they then must architect their networks and systems to be resilient in that assumed compromise state. This is the premise of Arista's zero trust security that marries an AI Driven security model with the segmentation and observability of cloud networking to being a secure cognitive client to cloud networking.

[2] *CrowdStrike 2020 Global Threat Report*

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

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