Arista Networks Multi-Domain Segmentation Service™ (MSS)

Introduction: Network Microperimeter Segmentation

Network micro-segmentation is a crucial aspect of modern cybersecurity: it provides organizations with the ability to enhance their security posture by implementing fine-grained security policies based on microperimeters defined around the identity of the endpoints or the applications rather than on traditional network boundaries like subnets and VRFs. A segmentation strategy based on microparameters is the most effective solution to minimize the overall attack surface available to any endpoint.

The rise of micro-segmentation technologies is influenced by several key trends, such as the increasing adoption of the zero trust architecture (ZTA), which drives the requirements to reduce lateral movement into increasingly smaller trust zones (microperimeters) where workloads are granularly identified and only approved connections are permitted.

Another significant trend driving the need for micro-segmentation is the growing importance of lateral attacks in network security. With the increase in ransomware attacks and the need to defend against sophisticated attackers that attempt to spread laterally, organizations require finer-grained policies that can block traffic at the workload level. This trend underscores the importance of implementing robust micro-segmentation strategies to defend against lateral movements and enhance the overall security posture of the network.

In this context, Arista defines microperimeter segmentation as the capability to insert a group-based security policy between any two groups with endpoints in the same network (e.g., VLAN) or across multiple networks. Microperimeter segmentation should be viewed as a solution stack that, besides policy enforcement, includes the following core capabilities:

- **Endpoint discovery:** This is the first step in planning the implementation of a microperimeter segmentation strategy and is used to gather advanced context and identity information to define the microperimeters.
- **Traffic session mapping:** This fundamental capability allows organizations to gather and visualize east/west bidirectional traffic flows (sessions), which are then used to generate a set of policy recommendations.
- **Policy recommendation engine:** Built on the observed traffic map, this capability enables the definition of rules to explicitly permit the traffic sessions detected in the network and to prevent the accidental breaking of existing applications. Generating a policy recommendation is an automated functionality, however the user remains responsible for auditing the recommended rules before implementing them in the network.
- **Integration with threat detection solutions** based on layer seven protocol inspection and anomaly detection technologies (such as Arista Network Detection and Response (NDR) or third-party endpoint security technologies).

Microperimeter segmentation solutions are typically based on switch infrastructure or host-based firewall technologies.

Switch-based solutions are attractive because of wire-speed performance and independence from any endpoint. Their two main drawbacks are the architecture inconsistency between campus and data center, as well as the dependency on specific VXLAN-based tagging protocols breaking interoperability with both legacy network infrastructure and traditional firewalls.

On the other hand, host-based firewall solutions have the advantage of being independent of the network (as micro-segmentation should be) and share the same architecture between campus and data center networks. Their two main drawbacks are the portability limitations across endpoints and the cost as well as the complexity of managing endpoint software at scale. Therefore, these solutions end up creating pockets of micro-segmentation, leaving most of the infrastructure unprotected.

Figure 1: Switch-based vs. Host-based Solutions

Arista Networks has defined the only multi-domain (campus, branch, and data center) segmentation technology in the industry that is consistent across any network domain and simultaneously independent of any network and any endpoint. Its advantages and components are described in the following section.

Arista Multi-domain Segmentation Services (MSS) Architecture Overview

The following key principles of the Arista MSS architecture make it stand out compared to all other solutions in the market:

1. **Consistent architecture across multiple network domains (campus, branch, data center)** predicated on a single EOS binary, common across all switching platforms, a single Arista CloudVision™ policy orchestration platform, and an aggregated Network Data Lake (Arista EOS NetDL™) infrastructure for state management and monitoring.
2. **The only microperimeter (tag-based)** segmentation solution in the industry that is both network and endpoint agnostic: it has no dependency on any network data plane or control plane protocols and at the same time it does not require any software agents on the endpoints.
3. **A unified framework for micro- as well as macro-segmentation policies** enabling security rules to be built around microperimeter objects as well as traditional network objects (subnets, VRFs), with the ability to enforce policies in the network or redirect traffic to any traditional security gateways.

Arista MSS offers a switch-based microperimeter technology stack that preserves the best attributes of switch-based and host-based firewall micro-segmentation technologies and, at the same time, overcomes their main limitations while delivering a consistent, unified segmentation architecture end-to-end across multiple domains (from the campus/branch to the data center).

Figure 2: Arista Multi-domain Segmentation Architecture

Arista MSS technology's advantages are depicted in the figure below:

Figure 3: Advantages of Arista MSS's Unified Segmentation Architecture

MSS Deployment with Traditional Network Macro-segmentation

The objective of Arista MSS is to complement traditional network segmentation using firewalls by extending granular security controls deeper into the network.

While firewalls only secure traffic between security zones, MSS fills the gap by extending security inside each zone where firewalls, for cost or performance reasons, cannot reach. The picture below illustrates the complementary positioning of MSS with traditional security gateways.

MSS goes beyond enforcement of micro-segmentation policies in the network. It empowers security policies to dynamically insert advanced firewall services into any micro-segmented connection, providing the ability to redirect micro-segmented traffic to a centralized security gateway.

Figure 4 : MSS Technology Deployment

Core Capabilities of Arista's MSS Microperimeter Segmentation

Arista MSS's core functionalities are divided into four areas corresponding to the logical steps necessary to implement a complete microperimeter segmentation strategy. Not all four capabilities are mandatory at the same time, but their aggregate provides a complete solution. They are shown in the figure below, arranged in logical order, and are described in the following.

Figure 5: Arista MSS's Core Functionalities

1. Dynamic Group Discovery

Any micro-segmentation strategy starts by defining the specific security groups that correspond to the microperimeters. This operation requires the administrators to identify the endpoints and/or the applications and to map them to specific groups. Such groups are then used as the foundation objects for the definition of the microperimeter segmentation security policies.

Endpoint discovery is the process of identifying and cataloging devices connected to a network, including computers, servers, mobile devices, and IoT devices. Endpoint discovery enables organizations to maintain an accurate inventory of network assets, assess security risks, and ensure compliance with security policies. It facilitates effective network management, monitoring, and the implementation of security measures to protect against potential threats and vulnerabilities.

CloudVision enables the administrator to either statically map endpoints (hosts or networks) to groups or leverage the integration with external asset- and endpoint-management databases to dynamically learn endpoints and how they map to groups.

In campus networks, Network Access Control (NAC) is a common security function that regulates and manages access to the network by devices and users, ensuring compliance with security policies and protecting against unauthorized access and potential threats. NAC manages the network identity services for the devices and can be used for endpoint discovery. Arista CloudVision can integrate with NAC solutions such as Arista CloudVision AGNI (Arista Guardian for Network Identity), Forescout, and Cisco Identity Services Engine (ISE).

CloudVision AGNI can also integrate with Arista NDR, which performs device group classification and threat assessment for various classes of devices (including IoT devices/BYODs). The NDR device's classification technology is based on advanced data plane traffic analysis powered by Arista AVA (Autonomous Virtual Assist).

Figure 6: Endpoint Discovery in Data Center Networks

In data center environments, Arista CloudVision derives dynamic group information from the integration with VMware vSphere, or with CMDBs such as ServiceNow and with IPAM systems such as Infoblox (as shown in the figure below). Furthermore, thanks to its versatility, the CloudVision API can extend its integrations to any database capable of exporting its tag information in the CSV format.

Figure 7: Endpoint Discovery in Data Center Networks

The figure below shows how CloudVision can import groups from external sources and allows the administrator to review and accept the groups to be used for security policies.

While dynamic group tagging is a requirement to manage scalable deployments, it is also possible to create static groups directly in CloudVision.

Figure 8: Group Import and Review in CloudVision

The recommended tool is called "MSS Planner" (see picture below), which provides a programmable spreadsheet to map networks/endpoints to tags (top portion of the image) and use those tags to visually build rules (bottom portion of the picture). The set of groups and rules can then be exported into a YAML file that can be directly imported into CloudVision.

Alternatively, the MSS Planner can also generate the CLI configuration to program the switches directly.

Figure 9: MSS Planner Tool

2. Stateful Traffic Map and Policy Recommendation Engine

Once the endpoints are mapped to microperimeter groups, it is important to obtain a precise map of the communications among these groups to be able to design security policy rules that do not break existing applications and that explicitly permit the forwarding of the discovered traffic flows. The objective of such a traffic map is to be able to recommend a set of microperimeter-based security rules—which the administrator can audit—to explicitly permit safe or valid traffic between any two groups.

This is an optional step, but it is very valuable to remove any guesswork on the effect(s) of new microperimeter-based policies on existing application flows. This stateful traffic map can also be continuously updated to both validate the policies and provide visibility into traffic that is dropped by any deny policy rules.

The Arista ZTX-7250S MSS appliance performs stateful monitoring using a special "monitor" forwarding rule configured by the Arista CloudVision MSS Manager on the network switches (mirror sources). Unsampled packets are copied from the mirror sources to a GRE tunnel to reach the ZTX appliance destination. After receiving the packets, the ZTX appliance performs session correlation and summarization: sessions are statefully validated, and directions (requester/responder) must be unequivocally identified to be able to recommend valid policies.

Figure 10: Arista ZTX-7250S MSS Appliance

Figure 11: ZTX-7250S Logical Deployment Model

The ZTX appliance then summarizes the stateful sessions and exports them to CloudVision (Network Data Lake) which in turn aggregates the IP address prefixes into groups and automatically generates a security rule set recommendation to explicitly allow the observed sessions (see figure below).

Figure 12: Stateful Observability

This key operation, exemplified in the figure above, is also referred to as traffic auditing or stateful observability.

For example, as shown in the figure, two employees connect to two different email servers repeatedly over a predefined period of time. They will start a request from an ephemeral source Layer 4 port toward designated destination port 25, and each server will respond from port 25 to the endpoint's ephemeral port. Over time client 200.1.1.10 will have sent N requests and will have received N responses from the server (for a total of N bidirectional sessions). Similarly, client 200.1.1.20 will have sent M requests and will have received M responses from the other server (for a total of M bidirectional sessions). Since each email session is bidirectional, the total number of observed unidirectional flows is (N+M)*2.

All the packets of each of these sessions ((N+M)*2 flows) are mirrored by the switches to a GRE tunnel (using a selective hardware rule applied to a specific path, as shown in the figure above) so that they can be conveyed to the (local or remote) MSS appliance for analysis and summarization.

In the email server example above, all the bidirectional sessions are aggregated into a common tuple defined as (VRF, source address, destination address, service). The counters of the aggregate sessions are also aggregated to represent the total amount of associated traffic.

The MSS appliance exports the summarized stateful sessions to CloudVision MSS Manager (see picture below), which further summarizes the IP address prefixes into groups and automatically generates a dynamic rule set based on the observed sessions. Therefore, CloudVision MSS Manager acts as a rule recommendation engine.

Figure 13: Session Summarization

As a subsequent step, network administrators can audit and modify the recommended rules before pushing them to the network devices.

Recommended policy rules can also support redirection to third-party firewalls, for example for traffic sessions crossing security zones or sent over the Internet. Aggregated session counters can be helpful in deciding which rules to offload to the network devices and which ones to implement on the firewalls.

Figure 14: Policy Orchestration and Audit

3. Policy Orchestration

After CloudVision has built its database of group objects (statically or dynamically), the groups can be used as configuration objects to build security policies.

Administrators can manage security policies both through the CLI and through the CloudVision MSS Manager dashboards. Compared to the CLI, the MSS Manager adds the ability to dynamically manage the group objects used in the policies, and to provide a rule recommendation engine based on the traffic inspection performed by the ZTX appliance.

In the following, this document focuses exclusively on examples of the CloudVision MSS Manager GUI.

The screenshot below shows the list of configured groups. In this example, four groups were created: three dynamic groups derived from VMware vCenter and one statically programmed by the administrator. The members of the groups are displayed at the top of the sidebar on the right-hand side, while at the bottom there is a list of rules utilizing the selected group.

Figure 15: Groups in MSS Manager Dashboard

To manage MSS policies CloudVision implements a dedicated MSS Studio workflow (see the figure below) which is structured in four separate sections to manage both the policies and the group objects:

- **Security Domains:** A security domain can correspond to a collection of switches sharing common security policies in a particular site, building, or data center pod. This section of the MSS Studio enables the operator to assign security policies to a security domain and to one of its VRFs (or to the default VRF if there are no VRFs configured). A screenshot is shown below:

Figure 16: Security Domains

- **Policies:** A security policy represents a container for a group of rules that is attached to a VRF and programmed into a security domain. This section allows the operator to create MSS policies comprising a set of rules, as shown in the screenshot below:

  Once a policy object is created it is possible to attach a set of rules to it. Rules contain the following fields (as shown in the figure below):

Figure 17: Policies

- **Source Group:** Matches a packet's source IP address against one or more address group objects.
- **Destination Group:** Matches the destination IP address against one or more address group objects.
- **Service:** Matches the protocol and destination L4 port(s) against a service object.
- **Direction:** If "bidirectional" (direction) is selected, CloudVision programs a second rule in hardware to match the reverse traffic (i.e., the response flow) where the source group and destination group are swapped and the source L4 port matches the service object.

Figure 18: Policy Rule Fields

- **Address Groups:** Address group objects are utilized in the source and destination fields of the security rules described above. This section allows the operator to create static groups and view the dynamic groups imported from external databases.

Figure 19: Address Groups

- **Services:** Service objects are utilized in the service field of the security rules. Notice that a service can contain more than one L4 destination port. This section supports the definition of service objects using protocols and L4 ports.

Figure 20: Policies

4. Distributed Enforcement with MSS Tagging Technology

The final deployment step is to distribute the rules and the objects to the network switches powered by EOS to enable wire-speed distributed enforcement of the MSS policies.

At this stage, it is critical to understand the advanced capabilities that the Arista MSS group tagging technology provides to be able to maximize their utility.

Arista MSS Group-based Tagging Concepts

To illustrate the flexibility of MSS Groups, for simplicity's sake in the sample campus and data center topologies displayed below, the figures represent the (wired and wireless) network with a single switch comprising multiple subnets (VLANs) and a single VRF.

Figure 21: Sample Campus Branch Topology with MSS Groups of Users, Office and IoT Devices

Figure 22: Sample Data Center Topology with MSS Groups of Application Workloads

Referring to the above examples, MSS tagging supports the following capabilities:

1. Endpoints or networks sharing a VRF and common security policies can be part of the same microperimeter (i.e., an MSS group). A group can be as small as a single endpoint/host (an individual IP address in a subnet or a /32 address) (see the group called user_id depicted in green in the figure above).
2. The group or tag information can be dynamically gleaned (e.g., from NAC or CMDB sources) or statically programmed.
3. Endpoints (or networks) in a group need to be part of the same VRF. However, endpoints can be part of the same subnet or multiple subnets in any arbitrary way.
4. MSS policies are completely abstracted from the location of the endpoints in the network. As a matter of fact, MSS can enforce policies even for traffic exchanged between groups on the same VLAN/subnet.
5. An endpoint or a network can be associated with more than one group (i.e., with more than one label): in other words, it is possible to create rules based on one or more labels (as explained in the following).

Expanding on the above capabilities that describe how MSS groups work, the next section discusses how an EOS switch leverages this endpoint grouping technology to provide wire-speed policy enforcement.

How an EOS Switch Handles Group Tags and Tag-based Rules

The MSS policy control plane in EOS translates the prefix-to-group mappings programmed by the MSS Manager into internal labels (a.k.a. MSS tags) stored in the hardware tables (which vary depending on the switch architecture).

These internal labels are locally significant to each switch and completely independent from the network (control and data plane). As a result, the MSS tagging technology can interoperate with any standard network topology and network protocol in a mixed-vendor environment. The label-to-prefix mappings, programmed by CloudVision into every switch, are stored in hardware tables (which vary depending on the switch architecture).

The labels stored in the "hardware label database" are utilized to program Ternary Content Addressable Memory (TCAM) rules where they summarize and replace sets of network prefixes as source and/or destination. The figure below shows the two distinct hardware tables that store the labels and the rules based on such labels.

Figure 23: Internal Tag/Label Use in Hardware-based Policy Rules

As depicted in the figure below, when packets enter a switch, a simultaneous source and destination IP address lookup operation occurs to derive the labels corresponding to the IP addresses (step 1 in the figure). Then the labels along with the rest of the packet fields are used to perform a lookup operation in the TCAM where rules are programmed using labels as sources and destinations (step 2). Once the TCAM applies a security action (e.g., drop or forward) the normal routing and bridging operations occur (step 3).

Once a packet matches a specific rule, multiple actions can be enforced in hardware:

- **Forward:** The traffic is allowed to be bridged or routed to the destination.
- **Drop:** The traffic is dropped (on ingress).
- **Monitor:** The traffic is allowed to be bridged or routed to the destination, but a copy is sent over to the ZTX appliance for stateful monitoring.
- **Drop+Monitor:** The traffic is dropped (on ingress), but a copy is sent over to the ZTX appliance for stateful monitoring.
- **Redirect:** The traffic is redirected to third-party security gateways for stateful/Layer 7 inspection.

Figure 24: Dynamic Action Enforcement of the Arista MSS Group Tagging Technology

In summary, MSS tag-based hardware enforcement provides multiple benefits:

1. The entire process occurs at wire speed without any latency or bandwidth penalty, and before the normal bridging and routing operations of the switch.
2. The MSS control plane in EOS optimizes the label generation to maximize the rule compression in the TCAM and enable much greater rule scalability in hardware.
3. Tags are locally significant; hence, they are data plane and control plane agnostic, enabling the graceful insertion of the MSS technology into any network and the interoperability with any firewall.
4. Supports hybrid rules with a mix of groups and IP address prefix objects along with multiple hardware-based actions, such as drop, forward, monitor, redirect to a third-party gateway, drop+monitor.

Understanding MSS Rules with Multiple Group Tags

In advanced network segmentation designs, it is important to be able to map an endpoint to more than one label to implement granular micro-perimeter policies. For example (see the picture below), an endpoint can be part of an application tier (e.g., web tier, backend tier, database tier), which is common across different applications (e.g., app1 through app N), which are replicated in independent environments (e.g., production, development, certification).

Figure 25: Multi-label Mapping

With the availability of multiple labels, you have the freedom to define coarse policies (e.g., using a single label) as well as more granular policies for individual endpoints associated with the intersection of more than one label. The ability to map an endpoint to multiple labels unlocks an unprecedented level of granular control that can be expressed with hardware-based policies.

The picture below shows an example of a network with 45 endpoints organized in three different environments, five applications, and three application tiers.

In the following, this paper uses this sample network to define a few rules to showcase the flexibility of Arista's MSS tagging system.

Figure 26: Sample Network Divided into Three Environments

Three examples with different levels of granularity:

1. The single-tag rule shown in the figure below states that Dev cannot talk to Prod, irrespective of the application being used or its application tier.

Figure 27: Rule Blocking Communication between Dev and Prod

2. In the example below, a dual-tag rule permits only the Web tier of App2 to use the HTTPS protocol to talk to the Backend Tier of App2, irrespective of the environment.

Figure 28: Double Tag Rule with App and Tier Filtering

3. Lastly, the figure below depicts a triple-tag rule, which is the most granular case, allowing App1 to initiate an HTTPS connection with App2 only in the Web Tier and only in the production environment.

Figure 29: Triple-tag Rule with Three Levels of Filtering

Conclusion

Arista's Multi-Domain Segmentation Service (MSS) is a comprehensive microperimeter segmentation solution that provides fine-grained security policies based on microperimeters defined around the identity of endpoints or applications. MSS offers a consistent architecture across multiple network domains, is both network and endpoint-agnostic, and enables the enforcement of policies in the network or redirection of traffic to traditional security gateways.

The core capabilities of Arista's MSS include:

- **Dynamic Group Discovery:** The MSS workflow starts by defining specific security groups that correspond to the microperimeters. This involves identifying and cataloging devices connected to a network, including computers, servers, mobile devices, and IoT devices. Arista CloudVision, the management platform for MSS, enables the administrator to either statically map endpoints to groups or leverage integration with external asset- and endpoint-management databases to dynamically learn endpoints and how they map to groups.
- **Stateful Traffic Map and Policy Recommendation Engine:** Once the endpoints are mapped to microperimeter groups, MSS provides a precise map of the communications among these groups to design security policy rules that do not break existing applications and explicitly permit the forwarding of the discovered traffic flows. The Arista ZTX-7250S MSS appliance performs stateful monitoring of the traffic sessions, and then CloudVision MSS Manager generates a set of recommended microperimeter-based security rules based on the observed traffic by the ZTX appliance.
- **Policy Orchestration:** CloudVision MSS Manager allows administrators to manage security policies and the related policy objects (e.g., groups, services, rules) and program them into the network.
- **Distributed Enforcement with the MSS Tagging Technology:** The final deployment step is to distribute the rules and objects to the network switches powered by EOS to enable wire-speed distributed enforcement of the MSS policies. EOS maps groups and prefixes to internal labels with an advanced set of compression algorithms that optimize hardware resource utilization and maximize scalability. Because the hardware labels are only internally significant to a switch, the MSS technology can be seamlessly inserted into any network.

In summary, Arista's Multi-Domain Segmentation Service (MSS) offers the most flexible solution in the industry for implementing microperimeter segmentation strategies, providing organizations with the ability to enhance their security posture and get a step closer to a true Zero Trust Architecture.

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