CIS Arista MLS EOS 4.X Router STIG Benchmark

v1.0.0 - 08-27-2025

Terms of Use

Please see the below link for our current terms of use:

https://www.cisecurity.org/cis-securesuite/cis-securesuite-membership-terms-of-use/

For information on referencing and/or citing CIS Benchmarks in 3rd party documentation (including using portions of Benchmark Recommendations) please contact CIS Legal (legalnotices@cisecurity.org) and request guidance on copyright usage.

NOTE: It is ***NEVER*** acceptable to host a CIS Benchmark in ***ANY*** format (PDF, etc.) on a 3rd party (non-CIS owned) site.

Table of Contents

Overview

Target Technology Details

Arista MLS EOS 4.X Router Secure Technical Implementation Guide (STIG)

Version: 2 Release: 2 Benchmark

Date: 02 Apr 2025

Intended Audience

This benchmark is intended for system and application administrators, security specialists, auditors, help desk, and platform deployment personnel who plan to develop, deploy, assess, or secure solutions that incorporate Arista MLS EOS 4.X Router and are looking to comply with the STIG guidance

Recommendation Definitions

The following defines the various components included in a CIS recommendation as applicable. If any of the components are not applicable it will be noted, or the component will not be included in the recommendation.

Title

Concise description for the recommendation's intended configuration.

Assessment Status

An assessment status is included for every recommendation. The assessment status indicates whether the given recommendation can be automated or requires manual steps to implement. Both statuses are equally important and are determined and supported as defined below:

Automated

Represents recommendations for which assessment of a technical control can be fully automated and validated to a pass/fail state. Recommendations will include the necessary information to implement automation.

Manual

Represents recommendations for which assessment of a technical control cannot be fully automated and requires all or some manual steps to validate that the configured state is set as expected. The expected state can vary depending on the environment.

Profile

A collection of recommendations (or rules) for securing a technology or a supporting platform. STIG Benchmark profiles are used to identify which Vulnerability Severity Category Code (CAT) each rule is associated with.

Description

The Rule Title from the STIG.

Rationale Statement

Detailed reasoning for the recommendation to provide the user a clear and concise understanding on the importance of the recommendation.

Audit Procedure

Systematic instructions for determining if the target system complies with the recommendation.

Remediation Procedure

Systematic instructions for applying recommendations to the target system to bring it into compliance according to the recommendation.

Additional Information

References from the STIG Rule if applicable.

Profile Definitions

The following configuration profiles are defined by this Benchmark:

- **SEVERITY: CAT I**

  Items in this profile exhibit one or more of the following characteristics:

  - are considered to be high severity
  - are intended for environments or use cases where following STIG based security guidance is paramount.
  - acts as defense in depth measure.
  - may negatively inhibit the utility or performance of the technology.

  This profile is intended for servers and workstations.

- **SEVERITY: CAT II**

  Items in this profile exhibit one or more of the following characteristics:

  - are considered to be medium severity
  - are intended for environments or use cases where following STIG based security guidance is paramount.
  - acts as defense in depth measure.
  - may negatively inhibit the utility or performance of the technology.

  This profile is intended for servers and workstations.

- **SEVERITY: CAT III**

  Items in this profile exhibit one or more of the following characteristics:

  - are considered to be low severity
  - are intended for environments or use cases where following STIG based security guidance is paramount.
  - acts as defense in depth measure.
  - may negatively inhibit the utility or performance of the technology.

  This profile is intended for servers and workstations.

Acknowledgements

The Recommendations in this Benchmark are a representation of the Rules in the unclassified DISA STIG for Arista MLS EOS 4.X Router.

Recommendations

1 STIG RULES

Arista Cloud EOS

Arista MLS EOS 4.X Router Secure Technical Implementation Guide (STIG)

Version: 2 Release: 2 Benchmark

Date: 02 Apr 2025

CLASSIFICATION unclassified

1.1 ARST-RT-000010 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to enforce approved authorizations for controlling the flow of information within the network based on organization-defined information flow control policies.

```
GROUP ID: V-255987
RULE ID: SV-255987r882303
```

Rationale:

Information flow control regulates where information is allowed to travel within a network and between interconnected networks. The flow of all network traffic must be monitored and controlled so it does not introduce any unacceptable risk to the network infrastructure or data. Information flow control policies and enforcement mechanisms are commonly employed by organizations to control the flow of information between designated sources and destinations (e.g., networks, individuals, and devices) within information systems.

Enforcement occurs, for example, in boundary protection devices (e.g., gateways, routers, guards, encrypted tunnels, and firewalls) that employ rule sets or establish configuration settings that restrict information system services, provide a packet filtering capability based on header information, or provide a message filtering capability based on message content (e.g., implementing key word searches or using document characteristics).

Audit:

This requirement is not applicable for the DODIN backbone.

Verify that for the Arista router configuration, access control lists (ACLs) and filters are configured to allow or deny traffic for specific source and destination addresses as well as ports and protocols. These filters must be applied inbound or outbound on the appropriate external and internal interfaces.

Example:

```
router# show ip access-lists
```

Verify IP access list configuration ACLs and filter are configured to allow or deny specific traffic.

```
!
ip access-list STIG
   10 deny ip 172.16.50.0/30 10.10.100.0/24
   20 permit ip any any
!
```

Verify the IP access list ACLs are applied to the specific Ethernet interface.

```
!
router# show ethernet Interface Eth3
!
Interface Ethernet 3
description BGP Link to Gateway Router
no router port
ip address 192.168.1.1/30
   ip access-group STIG in
!
```

If the Arista router is not configured to enforce approved authorizations for controlling the flow of information within the network based on organization-defined information flow control policies, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista routers to enforce ACLs and filters to allow or deny traffic for specific source and destination addresses as well as ports and protocols for controlling information flow.

To configure an IP access list to fulfill this function, enter the following commands, substituting organizational values for the bracketed variables and values.

Step 1:

```
Ip access-list [name]
[permit/deny] [protocol] [source address] [source port] [destination address]
[destination port]
Exit
```

Step 2:

Apply the filters inbound or outbound on the appropriate external and internal interfaces.

```
Interface [type] [number]
Ip access-group [name] [direction]
```

Note: Policy-based routing can also be implemented if needed.

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.2 ARST-RT-000020 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to reject inbound route advertisements for any Bogon prefixes.

```
GROUP ID: V-255988
RULE ID: SV-255988r882306
```

Rationale:

Accepting route advertisements for bogon prefixes can result in the local autonomous system (AS) becoming a transit for malicious traffic as it will in turn advertise these prefixes to neighbor autonomous systems.

Audit:

Review the Arista router configuration to verify it will reject routes of any Bogon prefixes.

The prefix filter must be referenced inbound on the appropriate BGP neighbor statements.

Step 1: Review the BGP Bogon Prefix Lists configured. To verify IP prefix lists are configured, execute the command "show ip prefix-list".

```
ip prefix-list BOGON_v4
  seq 1 deny 0.0.0.0/8 le 32
  seq 2 deny 10.0.0.0/8 le 32
  seq 3 deny 100.64.0.0/10 le 32
  seq 4 deny 127.0.0.0/8 le 32
  seq 5 deny 169.254.0.0/16 le 32
  seq 6 deny 172.16.0.0/12 le 32
  seq 100 permit 0.0.0.0/0 ge 8
```

Step 2: Review the prefix lists inbound to the appropriate BGP neighbor to verify the BGP config and verify the prefix is applied. Execute the command "show ip bgp nei X.2.1.1".

```
router bgp 65001
  neighbor 100.2.1.1 prefix-list BOGON_v4 in
```

If the Arista router is not configured to reject or permit inbound route advertisements for any bogon prefixes, this is a finding.

Remediation:

Step 1: Configure the BGP Bogon Prefix List.

```
LEAF-1A(config)#ip prefix-list BOGON_v4
LEAF-1A(config-ip-pfx)#seq 1 deny 0.0.0.0/8 le 32
LEAF-1A(config-ip-pfx)#seq 2 deny 10.0.0.0/8 le 32
LEAF-1A(config-ip-pfx)#seq 3 deny 100.64.0.0/10 le 32
LEAF-1A(config-ip-pfx)#seq 4 deny 127.0.0.0/8 le 32
LEAF-1A(config-ip-pfx)#seq 5 deny 169.254.0.0/16 le 32
LEAF-1A(config-ip-pfx)#seq 6 deny 172.16.0.0/12 le 32
LEAF-1A(config-ip-pfx)#seq 100 permit 0.0.0.0/0 ge 8
```

Step 2: Configure the prefix list inbound to the appropriate BGP neighbor.

```
LEAF-1A(config)#router bgp 65001
LEAF-1A(config-router-bgp)#neighbor 100.2.1.1 prefix-list BOGON_v4 in
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.3 ARST-RT-000030 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to reject inbound route advertisements for any prefixes belonging to the local autonomous system (AS).

```
GROUP ID: V-255989
RULE ID: SV-255989r882309
```

Rationale:

Accepting route advertisements belonging to the local AS can result in traffic looping or being black holed, or at a minimum using a non-optimized path.

Audit:

Review the Arista router configuration to verify that it will reject routes belonging to the local AS.

The prefix filter must be referenced inbound on the appropriate BGP neighbor statements.

Step 1: Review the prefix lists configured on the router to filter the local AS prefix (10.12.0.0/16). To verify IP prefix lists are configured, execute the command "show ip prefix-list".

```
router#sh ip prefix-list
router#ip prefix-list LOCAL_SCOPE_BOUNDARY
   seq 10 deny 10.12.0.0/16
   seq  100 permit  0.0.0.0/0 le32
```

Step 2: Review the BGP configuration to verify the prefix filter is applied inbound to the BGP neighbor. To verify the BGP config and verify the prefix is applied, execute the command "show run | section router bgp".

```
router bgp 65000
  router-id 10.11.11.11
  address-family ipv4
    no neighbor 10.11.12.2 prefix-list out
    neighbor 10.12.0.0 prefix-list LOCAL_SCOPE_BOUNDARY in
```

If the Arista router is not configured to reject inbound route advertisements belonging to the local AS, this is a finding.

Remediation:

Configure Arista eBGP routers to reject inbound route advertisements for prefixes that are not allocated to that specific customer.

Step 1: Configure the prefix-list to reject inbound route advertisements belonging to the local AS.

```
router(config)#ip prefix-list LOCAL_SCOPE_BOUNDARY
router(config-ip-pfx)#seq 10 deny 10.12.0.0/16
router(config-ip-pfx)#seq 100 permit 0.0.0.0/0 le32
```

Step 2: Configure a route-map to match the prefix-list.

```
router(config)#route-map LOCAL_AS deny
router(config-route-map-LOCAL_AS)#match IP address prefix-list
LOCAL_SCOPE_BOUNDARY
router(config-route-map-LOCAL_AS)#exit
```

Step 3: Configure the route-map to be applied inbound to the appropriate BGP neighbor.

```
router(config)#router bgp 65000
router(config-router-bgp)#neighbor 10.12.0.0 prefix-list LOCAL_SCOPE_BOUNDARY
in
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.4 ARST-RT-000040 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to reject inbound route advertisements from a customer edge (CE) router for prefixes that are not allocated to that customer.

```
GROUP ID: V-255990
RULE ID: SV-255990r882312
```

Rationale:

As a best practice, a service provider should only accept customer prefixes that have been assigned to that customer and any peering autonomous systems. A multi-homed customer with BGP speaking routers connected to the internet or other external networks could be breached and used to launch a prefix de-aggregation attack. Without ingress route filtering of customers, the effectiveness of such an attack could impact the entire IP core and its customers.

Audit:

Review the Arista router configuration to verify filters are defined to only accept routes for prefixes that belong to specific customers.

The prefix filter must be referenced inbound on the appropriate BGP neighbor statement.

Step 1: Verify the Arista router is configured with ip access-list BGP_Enclave_Gateway_Filter_IN. To verify IP access lists are configured, execute the command "show ip access-lists".

```
ip access-list BGP_Enclave_Gateway_Filter_IN
 10 permit ip 172.16.50.0/30 any
 20 permit ip 10.1.1.0/24 any
 30 permit  ip 12.15.4.9/32 any
 40 deny ip any any log
```

Step 2: Review the route-map and verify it matches the ACL. To verify route maps are configured, execute the command "show route-map".

```
route-map FILTER_INBOUND permit 10
   match ip address access-list BGP_Enclave_Gateway_Filter_IN
```

Step 3: Review the BGP configuration to verify the filter is applied inbound to the appropriate BGP neighbor. To verify the BGP config and verify the route map is applied, execute the command "show run | section router bgp".

```
router bgp 65001
 neighbor 100.2.1.1 route-map FILTER_INBOUND in
```

If the Arista router is not configured to reject inbound route advertisements from each CE router for prefixes that are not allocated to that customer, this is a finding.

Note: Routes to PE-CE links within a VPN are needed for troubleshooting end-to-end connectivity across the MPLS/IP backbone. Hence, these prefixes are an exception to this requirement.

Remediation:

Step 1: Configure Arista eBGP routers to reject inbound route advertisements from a CE router for prefixes that are not allocated to that specific customer.

```
LEAF-1A(config)#ip access-list BGP_Enclave_Gateway_Filter_IN
LEAF-1A(config-acl-BGP_Enclave_Gateway_Filter_IN)# 10 permit ip
172.16.50.0/30 any
LEAF-1A(config-acl-BGP_Enclave_Gateway_Filter_IN)# 20 permit ip 10.1.1.0/24
any
LEAF-1A(config-acl-BGP_Enclave_Gateway_Filter_IN)# 30 permit  ip 12.15.4.9/32
any
LEAF-1A(config-acl-BGP_Enclave_Gateway_Filter_IN)# 40 deny ip any any log
```

Step 2: Configure a route-map to match the ACL.

```
LEAF-1A(config-ip-pfx)#route-map FILTER_INBOUND permit 10
LEAF-1A(config-route-map-FILTER_INBOUND)#match ip address access-list
BGP_Enclave_Gateway_Filter_IN
```

Step 3: Configure the route-map to be applied inbound to the appropriate CE customer neighbor.

```
LEAF-1A(config)#router bgp 65001
LEAF-1A(config-router-bgp)#neighbor 100.2.1.1 route-map FILTER_INBOUND in
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.5 ARST-RT-000050 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to reject outbound route advertisements for any prefixes that do not belong to any customers or the local autonomous system (AS).

```
GROUP ID: V-255991
RULE ID: SV-255991r882315
```

Rationale:

Advertisement of routes by an autonomous system for networks that do not belong to any of its customers pulls traffic away from the authorized network. This causes a denial of service (DoS) on the network that allocated the block of addresses and may cause a DoS on the network that is inadvertently advertising it as the originator. It is also possible that a misconfigured or compromised router within the GIG IP core could redistribute IGP routes into BGP, thereby leaking internal routes.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify a filter is defined to only advertise routes for prefixes that belong to any customers or the local AS.

The prefix filter must be referenced outbound on the appropriate BGP neighbor statements.

Step 1: Verify the prefix list is configured on the router and is accepting only prefixes belonging to customers or the local AS prefix (10.12.0.0/16). To verify IP prefix lists are configured, execute the command "show ip prefix-list".

```
ip prefix-list ADVERTISE_ROUTES
   seq 10 permit 10.12.0.0/16
   seq 20 deny 10.17.0.0/16
   seq 30 deny 10.23.0.0/16
   seq 40 deny 10.47.0.0/16
   seq 50 deny 10.59.0.0/16
   seq 100 deny 0.0.0.0/0 le 32
```

Step 2: Verify in the BGP configuration that the filter is applied outbound for each customer to the appropriate BGP neighbor. To verify the BGP config and verify the prefix is applied, execute the command "show run | section router bgp".

```
router bgp 65001
 neighbor 100.2.1.1 prefix-list ADVERTISE_ROUTES out
```

If the Arista router is not configured to reject outbound route advertisements that do not belong to any customers or the local AS, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure all Arista eBGP routers to filter outbound route advertisements for prefixes that are not allocated to or belong to any customer or the local AS.

Configure the Arista router to deny outbound route advertisements for any prefix belonging to the local AS Boundary.

Step 1: Configure the prefix lists.

```
LEAF-1A(config)#ip prefix-list ADVERTISE_ROUTES
LEAF-1A(config-ip-pfx)#seq 10 permit 10.12.0.0/16
LEAF-1A(config-ip-pfx)#seq 20 deny 10.17.0.0/16
LEAF-1A(config-ip-pfx)#seq 30 deny 10.23.0.0/16
LEAF-1A(config-ip-pfx)#seq 40 deny 10.47.0.0/16
LEAF-1A(config-ip-pfx)#seq 50 deny 10.59.0.0/16
LEAF-1A(config-ip-pfx)#seq 100 deny 0.0.0.0/0 le 32
```

Step 2: Configure the prefix lists outbound to the appropriate BGP neighbor.

```
LEAF-1A(config)#router bgp 65001
LEAF-1A(config-router-bgp)#neighbor 100.2.1.1 prefix-list ADVERTISE_ROUTES
out
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.6 ARST-RT-000060 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista BGP router must be configured to reject route advertisements from BGP peers that do not list their autonomous system (AS) number as the first AS in the AS_PATH attribute.

```
GROUP ID: V-255992
RULE ID: SV-255992r945854
```

Rationale:

Verifying the path a route has traversed will ensure the IP core is not used as a transit network for unauthorized or possibly even internet traffic. All autonomous system boundary routers (ASBRs) must ensure updates received from eBGP peers list their AS number as the first AS in the AS_PATH attribute.

Audit:

The feature below is enabled by default.

Verify the BGP configuration to deny the updates received from eBGP peer that do not have the proper AS in the AS_PATH_attribute. To verify the BGP config and that the feature is applied, execute the command "show run all | in first".

```
router bgp 65001
  bgp enforce-first-as
```

If the router is not configured for "enforce-first-as", this is a finding.

Remediation:

Configure all Arista ASBR routers to deny updates received from eBGP peers that do not list their AS number as the first AS in the AS_PATH attribute.

```
LEAF-1A(config)#router  bgp 65001
LEAF-1A(config-router-bgp)#bgp enforce-first-as
```

Additional Information:

CCI-000032 Enforce information flow control using organization-defined security policy filters as a basis for flow control decisions for organization-defined information flows.

- NIST SP 800-53::AC-4 (8)
- NIST SP 800-53A::AC-4 (8).1 (ii)
- NIST SP 800-53 Revision 4::AC-4 (8)
- NIST SP 800-53 Revision 5::AC-4 (8) (a)

1.7 ARST-RT-000070 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista Multicast Source Discovery Protocol (MSDP) router must be configured to filter received source-active multicast advertisements for any undesirable multicast groups and sources.

```
GROUP ID: V-255993
RULE ID: SV-255993r882321
```

Rationale:

The interoperability of BGP extensions for interdomain multicast routing and MSDP enables seamless connectivity of multicast domains between autonomous systems. MP-BGP advertises the unicast prefixes of the multicast sources used by Protocol Independent Multicast (PIM) routers to perform RPF checks and build multicast distribution trees. MSDP is a mechanism used to connect multiple PIM sparse-mode domains, allowing RPs from different domains to share information about active sources. When RPs in peering multicast domains hear about active sources, they can pass on that information to their local receivers, thereby allowing multicast data to be forwarded between the domains. Configuring an import policy to block multicast advertisements for reserved, martian, single-source multicast, and any other undesirable multicast groups, as well as any source-group (S, G) states with bogon source addresses, would assist in avoiding unwanted multicast traffic from traversing the core.

Audit:

Review the Arista router configuration to determine if there is an import policy to block source-active multicast advertisements for any undesirable multicast groups, as well as any (S, G) states with undesirable source addresses.

Step 1: Verify that an inbound source-active filter is bound to each MSDP peer.

To verify the MSDP peer is configured and the source-active filter is configured inbound, execute the command "show run | sec router msdp".

```
router msdp
 peer 10.1.12.2
  sa-filter in PIM_NEIGHBOR_SA_FILTER
```

Step 2: Review the access lists referenced by the source-active filter to verify that undesirable multicast groups, auto-RP, single source multicast (SSM) groups, and advertisements from undesirable sources are blocked.

To verify IP access lists are configured, execute the command "show ip access-lists".

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

If the router is not configured with an import policy to block undesirable SA multicast advertisements, this is a finding.

Remediation:

Step 1: Configure the Arista router to implement policy to filter multicast advertisements for undesirable multicast groups and sources.

```
router msdp
 peer 10.1.12.2
  sa-filter in PIM_NEIGHBOR_SA_FILTER
```

Step 2: Configure the source active access-list.

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.8 ARST-RT-000080 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista Multicast Source Discovery Protocol (MSDP) router must be configured to filter source-active multicast advertisements to external MSDP peers to avoid global visibility of local-only multicast sources and groups.

```
GROUP ID: V-255994
RULE ID: SV-255994r882324
```

Rationale:

To avoid global visibility of local information, there are a number of source-group (S, G) states in a PIM-SM domain that must not be leaked to another domain, such as multicast sources with private address, administratively scoped multicast addresses, and the auto-RP groups (224.0.1.39 and 224.0.1.40).

Allowing a multicast distribution tree, local to the core, to extend beyond its boundary could enable local multicast traffic to leak into other autonomous systems and customer networks.

Audit:

Review the Arista router configuration to determine if there is export policy to block local source-active multicast advertisements.

Step 1: Verify that an outbound source-active filter is bound to each MSDP peer.

To verify the MSDP peer is configured and to verify the source-active filter is configured outbound, execute the command "show ip msdp peer X.1.12.2 and show ip msdp summary".

```
router msdp
 peer 10.1.12.2
  sa-filter out PIM_NEIGHBOR_SA_FILTER
```

Step 2: Review the access lists referenced by the source-active filters and verify that MSDP source-active messages being sent to MSDP peers do not leak advertisements that are local.

To verify IP access lists are configured, execute the command "show ip access-lists".

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

If the router is not configured with an export policy to block local source-active multicast advertisements, this is a finding.

Remediation:

Step 1: Configure Arista router to ensure an export policy is implemented on all MSDP routers to avoid global visibility of local multicast (S,G) states.

```
router msdp
 peer 10.1.12.2
  sa-filter in PIM_NEIGHBOR_SA_FILTER
```

Step 2: Configure the source active access-list.

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.9 ARST-RT-000090 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista MSDP router must be configured to limit the amount of source-active messages it accepts on per-peer basis.

```
GROUP ID: V-255995
RULE ID: SV-255995r882327
```

Rationale:

To reduce any risk of a denial-of-service (DoS) attack from a rogue or misconfigured MSDP router, the router must be configured to limit the number of source-active messages it accepts from each peer.

Audit:

To verify the MSDP peer and the sa-limit filter is configured, execute the command "show run | sec router msdp".

```
router msdp
 peer 10.1.12.2
   sa-limit 500
 peer 10.1.55.78
   sa-limit 900
```

If the Arista router is not configured with a peer limit, this is a finding.

Remediation:

Configure the Arista MSDP router to limit the amount of source-active messages it accepts from each peer.

```
!
router (config) #router msdp
router (config-router-msdp) #peer 10.1.1.5
router (config-router-msdp-peer 10.1.1.5) # sa-limit 500
router (config-router-msdp) #peer 10.1.55.78
router (config-router-msdp-peer 10.1.55.78) # sa-limit 900
router (config-router-msdp-peer 10.1.55.78) # exit
```

Additional Information:

CCI-001368 Enforce approved authorizations for controlling the flow of information within the system based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.10 ARST-RT-000100 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista BGP router must be configured to reject route advertisements from CE routers with an originating AS in the AS_PATH attribute that does not belong to that customer.

```
GROUP ID: V-255996
RULE ID: SV-255996r945855
```

Rationale:

Verifying the path a route has traversed will ensure that the local AS is not used as a transit network for unauthorized traffic. To ensure that the local AS does not carry any prefixes that do not belong to any customers, all PE routers must be configured to reject routes with an originating AS other than that belonging to the customer.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify the router is configured to deny updates received from CE routers with an originating AS in the AS_PATH attribute that does not belong to that customer.

Step 1: Review the router configuration and verify an as-path access-list statement is defined to only accept routes from a CE router whose AS did not originate the route.

```
show ip as-path access-list

ip as-path regex-mode asn
  ip as-path access-list NEIGHBOR_PATH permit ^35121$ any
  ip as-path access-list NEIGHBOR_PATH deny .* any
```

Step 2: Verify the as-path access list is referenced by the filter-list inbound for the appropriate BGP neighbors.

The filter-list CLI is not supported in Arista MLS. The workaround with route-map follows:

```
route-map TrafficOtherAS_Path permit 10
 match as-path NEIGHBOR_PATH
```

Step 3: To verify the BGP config and verifying the route map is applied inbound execute the command "show run | sec router bgp".

```
router bgp 65000
   neighbor 10.1.12.2 route-map TrafficOtherAS_Path in
```

If the Arista router is not configured to reject updates from CE routers with an originating AS in the AS_PATH attribute that does not belong to that customer, this is a finding.

Remediation:

Configure the Arista router to reject updates from CE routers with an originating AS in the AS_PATH attribute that does not belong to that customer.

Step 1: Configure the as-path access-list to filter the updates from the CE routers with an originating AS in the AS_PATH attribute that does not belong to that customer.

```
router (config) # ip as-path regex-mode asn
router (config) # ip as-path access-list NEIGHBOR_PATH permit ^35121$ any
router (config) # ip as-path access-list NEIGHBOR_PATH deny .* any
```

Step 2: Configure the route-map and match the as-path access-list.

```
route-map TrafficOtherAS_Path permit 10
 match as-path NEIGHBOR_PATH
```

Step 3: Apply the route-map to the appropriate neighbor.

```
router (config) # router bgp 65000
router (config-router-bgp) #neighbor 10.1.12.2 route-map TrafficOtherAS_Path
in
```

Additional Information:

CCI-000032 Enforce information flow control using organization-defined security policy filters as a basis for flow control decisions for organization-defined information flows.

- NIST SP 800-53::AC-4 (8)
- NIST SP 800-53A::AC-4 (8).1 (ii)
- NIST SP 800-53 Revision 4::AC-4 (8)
- NIST SP 800-53 Revision 5::AC-4 (8) (a)

1.11 ARST-RT-000110 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to enforce approved authorizations for controlling the flow of information between interconnected networks in accordance with applicable policy.

```
GROUP ID: V-255997
RULE ID: SV-255997r882333
```

Rationale:

Information flow control regulates authorized information to travel within a network and between interconnected networks. Controlling the flow of network traffic is critical so it does not introduce any unacceptable risk to the network infrastructure or data. An example of a flow control restriction is blocking outside traffic claiming to be from within the organization. For most routers, internal information flow control is a product of system design.

Audit:

Verify each Arista router enforces approved authorizations for controlling the flow of information between interconnected networks in accordance with applicable policy.

This requirement may be met through the use of IP access control lists.

Step 1: To verify on the Arista router that IP access lists are configured, execute the "show ip access-lists summary" command and check that the list is configured and is active on applicable interfaces.

```
router:#show ip access-lists summary
IPV4 ACL $$bgp-ttlSec-ip-vrf-default$$ [dynamic]
        Total rules configured: 1
        Configured on Ingress: bgp(default VRF)
        Active on     Ingress: bgp(default VRF)

IPV4 ACL ACL
        Total rules configured: 1

Standard IPV4 ACL ALLOWED_SOURCES
        Total rules configured: 2

IPV4 ACL AUTHORIZED_SOURCES
        Total rules configured: 3

Step 2: To verify the Arista router lists that control the flow of
information in accordance with organizational policy, enter the "show ip
access-list [name]" command and review the associated permit and deny
statements.

IP Access List ACL.
router#show ip access-list AUTHORIZED_SOURCES
  IP Access List AUTHORIZED_SOURCES
        10 permit ip 10.1.12.0/24 any
        20 deny ip 1.2.3.0/24 any log
        30 deny ip host 10.11.12.2 any log
```

If the Arista router does not enforce approved authorizations for controlling the flow of information between interconnected networks in accordance with applicable policy, this is a finding.

Remediation:

Configure the router to enforce approved authorizations for controlling the flow of information between interconnected networks in accordance with applicable policy.

To use an IP access list to fulfill this function, enter the following commands, substituting organizational values for the bracketed variables.

```
ip access-list [name]
[permit/deny] [protocol] [source address] [source port] [destination address]
[destination port]
exit

interface [type] [number]
ip access-group [name] [direction]
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.12 ARST-RT-000120 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista multicast router must be configured to disable Protocol Independent Multicast (PIM) on all interfaces that are not required to support multicast routing.

```
GROUP ID: V-255998
RULE ID: SV-255998r882336
```

Rationale:

If multicast traffic is forwarded beyond the intended boundary, it is possible that it can be intercepted by unauthorized or unintended personnel. Limiting where, within the network, a given multicast group's data is permitted to flow is an important first step in improving multicast security.

A scope zone is an instance of a connected region of a given scope. Zones of the same scope cannot overlap while zones of a smaller scope will fit completely within a zone of a larger scope. For example, Admin-local scope is smaller than Site-local scope, so the administratively configured boundary fits within the bounds of a site. According to RFC 4007 IPv6 Scoped Address Architecture (section 5), scope zones are also required to be "convex from a routing perspective"; that is, packets routed within a zone must not pass through any links that are outside of the zone. This requirement forces each zone to be one contiguous island rather than a series of separate islands.

As stated in the DOD IPv6 IA Guidance for MO3, "One should be able to identify all interfaces of a zone by drawing a closed loop on their network diagram, engulfing some routers and passing through some routers to include only some of their interfaces." Therefore, it is imperative that the network engineers have documented their multicast topology and thereby knows which interfaces are enabled for multicast. Once this is done, the zones can be scoped as required.

Audit:

If IPv4 or IPv6 multicast routing is enabled on the Arista router, verify all interfaces enabled for PIM are documented in the network's multicast topology diagram.

Review the Arista router configuration to determine which interfaces are enabled for PIM, identified via the "pim ipv4 sparse-mode" for ipv4 and "pim ipv6 sparse-mode" for ipv6 statement in the interface configuration, and compare to the topology.

```
sh run | sec pim

interface Ethernet3
   pim ipv4 sparse-mode
interface Ethernet8
   pim ipv4 sparse-mode
   pim ipv6 sparse-mode
interface Ethernet9
   pim ipv4 sparse-mode
   pim ipv6 sparse-mode
interface Vlan8
   pim ipv4 sparse-mode
```

If an interface is not required to support multicast routing and it is enabled, this is a finding.

Remediation:

Document all enabled interfaces for PIM in the network's multicast topology diagram. Disable support for PIM on interfaces that are not required to support it.

Step 1: Configure the router in global configuration mode to support multicast routing.

```
router(config)#router multicast
router(config-router-multicast)#ipv4
router(config-router-multicast-ipv4)#routing
router(config-router-multicast-ipv4)#exit
router(config-router-multicast)#exit
```

Step 2: Enable PIM on interfaces required to support multicast.

Interfaces have PIM disabled by default. To enable PIM from an interface active in a multicast network, enter "pim sparse-mode" in the interface configuration mode.

```
router(config)#interface Ethernet1
router(config-if-Et1)#pim ipv4 sparse-mode
router(config-if-Et1)#pim ipv6 sparse-mode
```

Step 3: Disable support for PIM on interfaces that are not required to support it.

```
router(config)#interface Ethernet2
router(config-if-Et2)#no pim ipv4 sparse-mode
router(config-if-Et2)#no pim ipv6 sparse-mode
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.13 ARST-RT-000130 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista multicast router must be configured to bind a Protocol Independent Multicast (PIM) neighbor filter to interfaces that have PIM enabled.

```
GROUP ID: V-255999
RULE ID: SV-255999r882339
```

Rationale:

PIM is a routing protocol used to build multicast distribution trees for forwarding multicast traffic across the network infrastructure. PIM traffic must be limited to only known PIM neighbors by configuring and binding a PIM neighbor filter to those interfaces that have PIM enabled. If a PIM neighbor filter is not applied to those interfaces that have PIM enabled, unauthorized routers can join the PIM domain, discover and use the rendezvous points, and also advertise their rendezvous points into the domain. This can result in a denial of service by traffic flooding or result in the unauthorized transfer of data.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router multicast topology diagram and determine if router interfaces are enabled for IPv4 or IPv6 multicast routing.

If the router is enabled for multicast routing, verify all interfaces enabled for PIM have a neighbor filter bound to the interface. The neighbor filter must only accept PIM control plane traffic from the documented PIM neighbors.

Step 1: Verify the ACL is configured that will specify the authorized PIM neighbors. To verify IP access lists are configured, execute the command "show ip access-lists".

```
ip access-list standard filter_1
  permit 10.13.24.9/24
  exit
```

Step 2: Verify the PIM neighbor-filter is configured on PIM-enabled interfaces. To verify interfaces are configured, execute the command "show run int YY".

```
interface vlan 4
  pim ipv4 sparse-mode
  pim ipv4 neighbor-filter filter_1
  exit
```

If PIM neighbor filters are not bound to all interfaces that have PIM enabled, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure neighbor filters to only accept PIM control plane traffic from documented PIM neighbors. Bind neighbor filters to all PIM-enabled interfaces.

Step 1: Configure an ACL that will specify the authorized PIM neighbors.

```
router(config)#ip access-list standard filter_1
router(config-std-acl-filter_1)#permit 10.13.24.9/24
router(config-std-acl-filter_1)#exit
```

Step 2: Configure a PIM neighbor-filter command and apply it on all PIM-enabled interfaces that are referencing the PIM neighbor ACL.

```
router(config)#interface vlan 4
router(config-if-Vl4)#pim ipv4 neighbor-filter filter_1
router(config-if-Vl4)#exit
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.14 ARST-RT-000140 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista multicast edge router must be configured to establish boundaries for administratively scoped multicast traffic.

```
GROUP ID: V-256000
RULE ID: SV-256000r882342
```

Rationale:

If multicast traffic is forwarded beyond the intended boundary, it is possible that it can be intercepted by unauthorized or unintended personnel.

Administrative scoped multicast addresses are locally assigned and are to be used exclusively by the enterprise network or enclave. Administrative scoped multicast traffic must not cross the enclave perimeter in either direction. Restricting multicast traffic makes it more difficult for a malicious user to access sensitive traffic.

Admin-Local scope is encouraged for any multicast traffic within a network intended for network management, as well as for control plane traffic that must reach beyond link-local destinations.

Audit:

Review the Arista router configuration and verify that admin-scope multicast traffic is blocked at the external edge.

Step 1: Verify the Arista router ACL is configured to deny packets with multicast administratively scoped destination addresses and verify IP access lists are configured. Execute the command "show ip access-lists".

```
ip access-list standard mbac1
 10 deny 239.120.10.0/24
 20 permit 224.0.0.0/4
 exit
```

Step 2: Verify the ACL is applied on the multicast boundary at the appropriate interfaces and verify interfaces are configured. Execute the command "show run int YY".

```
interface vlan 200
 multicast ipv4 boundary mbac1 out
 exit
```

If the Arista router is not configured to establish boundaries for administratively scoped multicast traffic, this is a finding.

Remediation:

Step 1: Configure the Arista router ACL to deny packets with multicast administratively scoped destination addresses.

```
router(config)#ip access-list standard mbac1
router(config-std-acl-mbac1)#10 deny 239.120.10.0/24
router(config-std-acl-mbac1)#20 permit 224.0.0.0/4
router(config-std-acl-mbac1)#exit
```

Step 2: Apply the multicast boundary at the appropriate interfaces.

```
router(config)#interface vlan 200
router(config-if-Vl200)#multicast ipv4 boundary mbac1 out
router(config-if-Vl200)#exit
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.15 ARST-RT-000150 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista router must be configured to have all inactive interfaces disabled.

```
GROUP ID: V-256001
RULE ID: SV-256001r882345
```

Rationale:

An inactive interface is rarely monitored or controlled and may expose a network to an undetected attack on that interface. Unauthorized personnel with access to the communication facility could gain access to a router by connecting to a configured interface that is not in use.

If an interface is no longer used, the configuration must be deleted and the interface disabled. For sub-interfaces, delete sub-interfaces that are on inactive interfaces and delete sub-interfaces that are themselves inactive. If the sub-interface is no longer necessary for authorized communications, it must be deleted.

Audit:

Review the Arista router configuration.

Verify the interfaces and sub-interfaces execute the commands "show ip interface brief" and "show interface status".

Example of a disabled interface:

```
interface Ethernet 8-10
 description The interface is administratively shutdown
 shutdown
```

If an interface is not being used but is configured or enabled, this is a finding.

Remediation:

Delete inactive sub-interfaces and disable and delete the configuration of any inactive ports on the router.

Deleting the sub-interface:

```
router(config)#no interface Ethernet8.100
```

Disabling the interface:

```
router(config)#interface Ethernet 8-10
router(config-if-Et8-10)#shutdown
```

Resetting the interface to the default-configuration:

```
router(config)#default interface Ethernet 8
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.16 ARST-RT-000160 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista perimeter router must be configured to protect an enclave connected to an alternate gateway by using an inbound filter that only permits packets with destination addresses within the sites address space.

```
GROUP ID: V-256002
RULE ID: SV-256002r882348
```

Rationale:

Enclaves with alternate gateway connections must take additional steps to ensure there is no compromise on the enclave network or NIPRNet. Without verifying the destination address of traffic coming from the site's alternate gateway, the perimeter router could be routing transit data from the internet into the NIPRNet. This could also make the perimeter router vulnerable to a denial-of-service (DoS) attack as well as provide a back door into the NIPRNet. The DOD enclave must ensure the ingress filter applied to external interfaces on a perimeter router connecting to an Approved Gateway is secure through filters permitting packets with a destination address belonging to the DOD enclave's address block.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration of each router interface connecting to an alternate gateway.

Verify each permit statement of the ingress filter only permits packets with destination addresses of the site's NIPRNet address space or a destination address belonging to the address block assigned by the alternate gateway network service provider.

Step 1: Verify an inbound ACL is configured to permit the packets with the destination addresses of the site's NIPRNet address space. Verify IP access lists are configured. Execute the command "show ip access-lists".

```
ip access-list NIPRNet_ACL
 permit tcp any host 10.51.12.34 eq www
 permit icmp host 10.51.12.25 host 10.51.12.65 echo
 permit icmp host 10.51.12.25 host 10.51.12.65 echo-reply
 permit 50 any host 10.51.12.28
 permit gre any host 10.51.12.28
 deny ip any any log
```

Step 2: Verify the ACL is applied inbound to the service provider-facing interface. Verify interfaces are configured. Execute the command "show run int YY".

```
interface ethernet 3
 ip access-group NIPRNet_ACL in
```

If the ingress filter permits packets with addresses other than those specified, such as destination addresses of the site's NIPRNet address space or a destination address belonging to the address block assigned by the alternate gateway network service provider, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the router for ingress filter of the perimeter router connected to an alternate gateway to only permit packets with destination addresses of the site's NIPRNet address space or a destination address belonging to the address block assigned by the alternate gateway network service provider.

Step 1: Configure an ACL inbound to the interface-facing service provider.

```
LEAF-1A(config)#ip access-list NIPRNet_ACL
LEAF-1A(config-acl-NIPRNet_ACL)#permit tcp any host 10.51.12.34 eq www
LEAF-1A(config-acl-NIPRNet_ACL)#permit icmp host 10.51.12.25 host 10.51.12.65
echo
LEAF-1A(config-acl-NIPRNet_ACL)#permit icmp host 10.51.12.25 host 10.51.12.65
echo-reply
LEAF-1A(config-acl-NIPRNet_ACL)#permit 50 any host 10.51.12.28
LEAF-1A(config-acl-NIPRNet_ACL)#permit gre any host 10.51.12.28
LEAF-1A(config-acl-NIPRNet_ACL)#deny ip any any log
```

Step 2: Apply the ACL to the internet service provider-facing interface.

```
LEAF-1A(config)#interface ethernet 3
LEAF-1A(config-if-Et3)#ip access-group NIPRNet_ACL in
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.17 ARST-RT-000170 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista perimeter router must be configured to not be a Border Gateway Protocol (BGP) peer to an alternate gateway service provider.

```
GROUP ID: V-256003
RULE ID: SV-256003r882351
```

Rationale:

ISPs use BGP to share route information with other autonomous systems (i.e., other ISPs and corporate networks). If the perimeter router was configured to BGP peer with an ISP, NIPRNet routes could be advertised to the ISP; thereby creating a backdoor connection from the internet to the NIPRNet.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration of the router connecting to the alternate gateway.

To verify no BGP neighbors are configured to the remote AS that belongs to the alternate gateway service provider and the static route is configured, execute the command "show ip route static".

ip route 192.168.67.0/24 12.15.4.9

If BGP neighbors are connecting the remote AS of the alternate gateway service provider, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure a static route on the perimeter router to reach the AS of a router connecting to an alternate gateway.

router(config)#ip route 192.168.67.0/24 12.15.4.9

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.18 ARST-RT-000180 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista perimeter router must be configured to not redistribute static routes to an alternate gateway service provider into BGP or an IGP peering with the NIPRNet or to other autonomous systems.

```
GROUP ID: V-256004
RULE ID: SV-256004r882354
```

Rationale:

If the static routes to the alternate gateway are being redistributed into an Exterior Gateway Protocol or Interior Gateway Protocol to a NIPRNet gateway, this could make traffic on NIPRNet flow to that particular router and not to the internet Access Point routers. This could not only wreak havoc with traffic flows on NIPRNet, but it could overwhelm the connection from the router to the NIPRNet gateway(s) and also cause traffic destined for outside of NIPRNet to bypass the defenses of the internet Access Points.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration of the router connecting to the alternate gateway and verify that redistribution of static routes to the alternate gateway is not occurring.

Verify the BGP and IGP configurations and remove the redistribute static statement if it is configured.

BGP Example:

To verify the BGP configuration, execute the commands "show bgp configuration active" and "show run section router bgp".

router bgp 1500

no redistribute static

OSPF Example:

To verify the OSPF configuration, execute the command "show run section router ospf".

router ospf 1

no redistribute static

RIP Example:

To verify the RIP configuration, execute the command "show run section router rip".

router rip

no redistribute static

If the static routes to the alternate gateway are being redistributed into BGP or any IGP peering with a NIPRNet gateway or another autonomous system, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the router so static routes are not redistributed to an alternate gateway into either an Exterior Gateway Protocol or Interior Gateway Protocol to the NIPRNet or to other autonomous systems.

Review the BGP and IGP configurations and remove the redistribute static statement if it is configured.

BGP Example:

router bgp 1500

no redistribute static

OSPF Example:

router ospf 1500

no redistribute static

RIP Example:

router rip

no redistribute static

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.19 ARST-RT-000190 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The out-of-band management (OOBM) Arista gateway router must be configured to have separate IGP instances for the managed network and management network.

```
GROUP ID: V-256005
RULE ID: SV-256005r882357
```

Rationale:

If the gateway router is not a dedicated device for the OOBM network, implementation of several safeguards for containment of management and production traffic boundaries must occur. Since the managed and management network are separate routing domains, configuration of separate Interior Gateway Protocol routing instances is critical on the router to segregate traffic from each network.

Audit:

This requirement is not applicable for the DODIN backbone.

Verify the OOBM interface is an adjacency in the Interior Gateway Protocol routing domain for the management network.

Verify interface configuration that the OOBM management network subnet is configured.

Step 1: To verify ospf process 100 interface is configured, execute the command "show run int YY". To verify vrf instance, execute "show vrf". Verify the OOBM vrf instance is configured.

```
vrf instance OOBM
ip routing vrf OOBM
interface Vlan 2
   description Connection to OOBM-LAN-Ethernet4
   vrf OOBM
   mtu 9214
   no routerport
   ip address 10.1.12.7/31
```

Step 2: To verify OSPF process is configured as OOBM management network, execute the command "show run section router ospf 100".

```
router ospf 100 vrf OOBM
 network 10.1.12.0/24 area 0.0.0.0
```

Step 3: To verify OSPF process 200 is enabled on a private network without any connectivity with the OSPF process in management network, execute the command "show run int YY". Verify the LAN vrf instance is configured.

```
vrf instance LAN
ip routing vrf LAN

interface Ethernet8
   vrf LAN
   description Connection to Private-LAN-Ethernet4
   mtu 9214
   no routerport
   ip address 172.16.35.135/31
```

Step 4: To verify OSPF process 200 is configured as Private-LAN network, execute the command "show run section router ospf 200".

```
router ospf 200 vrf LAN
 network 172.16.35.0/24 area 0.0.0.0
```

If the router does not enforce that Interior Gateway Protocol instances configured on the OOBM gateway router peer only with their own routing domain, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to enforce that Interior Gateway Protocol instances configured on the OOBM gateway router peer only with their own routing domain.

Configuring OSPF:

Step 1: Configure the interface and OOBM vrf instance.

```
vrf instance OOBM
ip routing vrf OOBM

LEAF-1A(config)#interface Vlan 2
LEAF-1A(config-if-Vl2)#description Connection to OOBM-LAN-Ethernet4
LEAF-1A(config-if-Vl2)#vrf OOBM
LEAF-1A(config-if-Vl2)#mtu 9214
LEAF-1A(config-if-Vl2)#no routerport
LEAF-1A(config-if-Vl2)#ip address 10.1.12.7/31
```

Step 2: Advertise the subnet in OSPF process 100.

```
LEAF-1A(config-router-rip)#router ospf 100 vrf OOBM
LEAF-1A(config-router-ospf)#network 10.1.12.0/24 area 0.0.0.0
```

Step 3: Configure the interface and LAN vrf instance.

```
vrf instance LAN
ip routing vrf LAN

LEAF-1A(config)#interface Ethernet8
LEAF-1A(config-if-Et8)#description Connection to Private-LAN-Ethernet4
LEAF-1A(config-if-Et8)#vrf LAN
LEAF-1A(config-if-Et8)#mtu 9214
LEAF-1A(config-if-Et8)#no routerport
LEAF-1A(config-if-Et8)#ip address 172.16.35.135/31
```

Step 4: Advertise the subnet in OSPF process 200.

```
LEAF-1A(config-router-ospf)#router ospf 200 vrf LAN
LEAF-1A(config-router-rip)#network 172.16.35.0/24 area 0.0.0.0
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.20 ARST-RT-000200 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The out-of-band management (OOBM) Arista gateway router must be configured to not redistribute routes between the management network routing domain and the managed network routing domain.

```
GROUP ID: V-256006
RULE ID: SV-256006r882360
```

Rationale:

If the gateway router is not a dedicated device for the OOBM network, several safeguards must be implemented for containment of management and production traffic boundaries; otherwise, it is possible that management traffic will not be separated from production traffic.

Since the managed network and the management network are separate routing domains, separate Interior Gateway Protocol routing instances must be configured on the router, one for the managed network and one for the OOBM network. In addition, the routes from the two domains must not be redistributed to each other.

Audit:

This requirement is not applicable for the DODIN backbone.

Verify the Interior Gateway Protocol instance used for the managed network does not redistribute routes into the Interior Gateway Protocol instance used for the management network and vice versa.

PROD IGP Routing Domain:

Step 1: To verify interfaces and vrf instance are configured, execute the command "show run int YY".

```
interface Et3/17/1
 description To_PROD
 ip address 10.1.12.1/24
```

Step 2: Verify the OSPF configuration, the PROD subnet is advertised, and IGP redistribution is removed in the OSPF process. To verify the OSPF configuration, execute the command "show run section router ospf".

```
router ospf 100 vrf PROD
   network 10.1.0.0/24 area 0.0.0.0
   no redistribute rip
```

OOBM IGP Routing Domain, running on the management network

Step 3: To verify interfaces are configured, execute the command "show run int YY".

```
interface Et3/17/2
 description To_OOBM
 ip address 172.16.10.1/24
```

Step 4: Verify the RIP configuration, the OOBM subnet is advertised, and IGP redistribution is removed in the RIP process. To verify the RIP configuration, execute the command "show run section router rip".

```
router rip
   network 172.16.10.0/24
   no redistribute ospf
   no shutdown
```

If the Interior Gateway Protocol instance used for the managed network redistributes routes into the Interior Gateway Protocol instance used for the management network or vice versa, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Interior Gateway Protocol instance used for the managed network to prohibit redistribution of routes into the Interior Gateway Protocol instance used for the management network and vice versa.

PROD IGP Routing Domain:

Step 1: Configure the interface.

```
interface Et3/17/1
 description To_PROD
 ip address 10.1.12.1/24
```

Step 2: Configure the OSPF process to remove the IGP redistribution.

```
router ospf 100
   network 10.1.0.0/24 area 0.0.0.0
   no redistribute rip
```

OOBM IGP Routing Domain, running on the management network

Step 3: Configure the interface.

```
interface Et3/17/1
 description To_OOBM
 ip address 172.16.10.1/24
```

Step 4: Configure the RIP process to remove the IGP redistribution.

```
router rip
   network 172.16.10.0/24
   no redistribute ospf
   no shutdown
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.21 ARST-RT-000210 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The multicast Rendezvous Point (RP) Arista router must be configured to filter Protocol Independent Multicast (PIM) Register and Join messages received from the Designated Router (DR) for any undesirable multicast groups and sources.

```
GROUP ID: V-256007
RULE ID: SV-256007r882363
```

Rationale:

Real-time multicast traffic can entail multiple large flows of data. An attacker can flood a network segment with multicast packets, over-using the available bandwidth and thereby creating a denial-of-service (DoS) condition. Hence, it is imperative that register messages are accepted only for authorized multicast groups and sources.

Satisfies: SRG-NET-000019-RTR-000013, SRG-NET-000019-RTR-000014

Audit:

Verify that the RP router is configured to filter PIM register and join messages.

Step 1: To verify the ACL is configured to filter the multicast groups, execute the command "show ip access-lists".

```
ip access-list standard ALLOWED_MULTICAST_GROUP
   10 permit 224.0.0.0/8
   20 deny any
```

Step 2: To verify the ACL is applied to the PIM process, execute the command "show run section router pim".

```
router pim sparse-mode
   ipv4
      rp address 100.2.1.6 access-list ALLOWED_MULTICAST_GROUP
```

If the RP router peering with PIM-SM routers is not configured with a PIM import policy to block registration messages for any undesirable multicast groups and sources, this is a finding.

Remediation:

Configure the RP router to filter PIM register and join messages received from a multicast DR for any undesirable multicast groups or sources.

Step 1: Configure an ACL to filter the multicast groups.

```
LEAF-1A(config)#ip access-list standard ALLOWED_MULTICAST_GROUP
LEAF-1A(config-std-acl-ALLOWED_MULTICAST_GROUP)#10 permit 224.0.0.0/8
LEAF-1A(config-std-acl-ALLOWED_MULTICAST_GROUP)#20 deny any
```

Step 2: Apply the ACL in the PIM process globally.

```
LEAF-1A(config)#router pim sparse-mode
LEAF-1A(config-router-pim-sparse)#ipv4
LEAF-1A(config-router-pim-sparse-ipv4)#rp address 100.2.1.6 access-list
ALLOWED_MULTICAST_GROUP
```

Additional Information:

CCI-001414 Enforce approved authorizations for controlling the flow of information between connected systems based on organization-defined information flow control policies.

- NIST SP 800-53::AC-4
- NIST SP 800-53A::AC-4.1 (iii)
- NIST SP 800-53 Revision 4::AC-4
- NIST SP 800-53 Revision 5::AC-4

1.22 ARST-RT-000230 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to produce audit records containing information to establish where the events occurred.

```
GROUP ID: V-256008
RULE ID: SV-256008r882366
```

Rationale:

Without establishing where events occurred, it is impossible to establish, correlate, and investigate the events leading up to an outage or attack.

In order to compile an accurate risk assessment and provide forensic analysis, it is essential for security personnel to know where events occurred, such as router components, modules, device identifiers, node names, and functionality.

Associating information about where the event occurred within the network provides a means of investigating an attack, recognizing resource utilization or capacity thresholds, or identifying an improperly configured router.

Satisfies: SRG-NET-000076-RTR-000001, SRG-NET-000077-RTR-000001, SRG-NET-000078-RTR-000001

Audit:

The Arista router must log all packets that have been dropped via the access control list (ACL).

Log output must contain an interface name as to where the packet was filtered.

Step 1: To verify the deny ACL is created with the log statement for dropped packets, execute the command "show ip access-list".

```
ip access-list test1
 permit ip 10.30.30.0/24 host 10.20.10.1
 deny ip 10.30.10.0/24 host 10.20.10.1 log
```

Step 2: To verify the ACL ingress is applied on the appropriate interface, execute the command "show run interface YY".

```
interface ethernet 3
 ip access-group test1 in

######

Variables in the syslog messages display the following values:
---------------------------------------------------------------
acl Name of ACL.
intf Name of interface that received the packet.
filter Action triggered by ACL (denied or permitted).
protocol IP protocol specified by packet.
vlan Number of VLAN receiving packet.
ether EtherType protocol specified by packet.
src-ip and dst-ip source and destination IP addresses.
src-prt and dst-prt source and destination ports.
src-mac and dst-mac source and destination MAC addresses.
```

If the logged output does not contain an interface name as to where the packet was filtered, this is a finding.

If the Arista router fails to log all packets that have been dropped via the ACL, this is a finding.

Remediation:

Configure the router to record the interface in the log record for packets being dropped.

Step 1: Configure the ACL.

```
router(config)#ip access-list test1
router(config-acl-test1)#15 permit ip 10.30.30.0/24 host 10.20.10.1
router(config-acl-test1)#15 deny ip 10.30.10.0/24 host 10.20.10.1 log
```

Step 2: Apply the ACL ingress on the appropriate interface.

```
router(config)#interface ethernet 3
router(config-if-Et3)#ip access-group test1 in
```

Additional Information:

CCI-000132 Ensure that audit records containing information that establishes where the event occurred.

- NIST SP 800-53::AU-3
- NIST SP 800-53A::AU-3.1
- NIST SP 800-53 Revision 4::AU-3
- NIST SP 800-53 Revision 5::AU-3 c

CCI-000133 Ensure that audit records containing information that establishes the source of the event.

- NIST SP 800-53::AU-3
- NIST SP 800-53A::AU-3.1
- NIST SP 800-53 Revision 4::AU-3
- NIST SP 800-53 Revision 5::AU-3 d

CCI-000134 Ensure that audit records containing information that establishes the outcome of the event.

- NIST SP 800-53::AU-3
- NIST SP 800-53A::AU-3.1
- NIST SP 800-53 Revision 4::AU-3
- NIST SP 800-53 Revision 5::AU-3 e

1.23 ARST-RT-000260 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista router must be configured to have all non-essential capabilities disabled.

```
GROUP ID: V-256009
RULE ID: SV-256009r882369
```

Rationale:

A compromised router introduces risk to the entire network infrastructure, as well as data resources that are accessible via the network. The perimeter defense has no oversight or control of attacks by malicious users within the network. Preventing network breaches from within is dependent on implementing a comprehensive defense-in-depth strategy, including securing each device connected to the network. This is accomplished by following and implementing all security guidance applicable for each node type. A fundamental step in securing each router is to enable only the capabilities required for operation.

Audit:

Review the Arista router configuration to determine if services or functions not required for operation or not related to router functionality (e.g., DNS, email client or server, FTP server, or web server) are enabled.

The Arista router commands can disable any individual features not required.

```
router(config)#no ip ftp [source] interfaceno logging console
no ip domain lookup source-interface lo0
no ntp
no mlag configuration
no dhcp server
no dns domain
snmp-server community community1 ro
vlan 1 trunk group DO_NOT_USE
logging trap 6
```

The Arista router configuration sample below demonstrates the default security configuration and available services that can be configured.

```
!
management console
   idle-timeout 0
!
management ssh
   idle-timeout 0
!
management telnet
   shutdown
   idle-timeout 0
!
aaa authentication login default local
no aaa authentication login console
aaa authentication enable default local
no aaa authorization console
no aaa authorization exec default
no aaa authorization commands all default
aaa authorization config-commands
no aaa accounting exec console
no aaa accounting commands all console
no aaa accounting exec default
no aaa accounting commands all default
!
no enable secret
no aaa root
no aaa authentication policy local allow-nopassword-remote-login
!
username admin privilege 1 nopassword
!
no radius-server key
radius-server timeout 5
radius-server retransmit 3
no radius-server deadtime
!
no snmp-server engineID local
no snmp-server chassis-id
no snmp-server contact
no snmp-server location
no snmp-server source-interface
snmp-server enable traps
default snmp-server enable traps entity
default snmp-server enable traps lldp
default snmp-server enable traps snmp
default snmp-server enable traps spanning-tree
default snmp-server enable traps test
!
spanning-tree mode mstp
spanning-tree max-age 20
spanning-tree forward-time 15
spanning-tree transmit hold-count 6
spanning-tree max-hops 20
no spanning-tree portfast bpduguard default
no spanning-tree loopguard default
spanning-tree bpduguard rate-limit default
spanning-tree mst 0 priority 32768
!
control-plane
```

```
   ip access-group default-control-plane-acl in
!
no tacacs-server key
tacacs-server timeout 5
!
no banner login
no banner motd
!
```

Note that only SSH is enabled by default. All cleartext protocols (telnet, SNMP) are disabled by default.

If unnecessary services and functions are enabled on the Arista router, this is a finding.

Remediation:

Remove unneeded services and functions from the router. Removal is recommended because the service or function may be inadvertently enabled otherwise. However, if removal is not possible, disable the service or function.

Step 1: Configure the Arista device to disable the features that are not required.

```
no logging console
no ip domain lookup source-interface lo0
no ntp
no mlag configuration
no dhcp server
```

Step 2: Configure the Arista device to disable the use of VLAN 1.

```
vlan 1 state suspend
interface e1
!
```

Step 3: Configure all unused ports to designated VLAN and suspend the VLAN to disable all unused ports.

```
 routerport trunk native vlan 1000
   routerport trunk allowed vlan except 1
```

Additional Information:

CCI-000381 Configure the system to provide only organization-defined mission essential capabilities.

- NIST SP 800-53::CM-7
- NIST SP 800-53A::CM-7.1 (ii)
- NIST SP 800-53 Revision 4::CM-7 a
- NIST SP 800-53 Revision 5::CM-7 a

1.24 ARST-RT-000280 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to authenticate all routing protocol messages using NIST-validated FIPS 198-1 message authentication code algorithm.

```
GROUP ID: V-256010
RULE ID: SV-256010r882372
```

Rationale:

A rogue router could send a fictitious routing update to convince a site's perimeter router to send traffic to an incorrect or even a rogue destination. This diverted traffic could be analyzed to learn confidential information about the site's network or used to disrupt the network's ability to communicate with other networks. This is known as a "traffic attraction attack" and is prevented by configuring neighbor router authentication for routing updates. However, using clear-text authentication provides little benefit since an attacker can intercept traffic and view the authentication key. This would allow the attacker to use the authentication key in an attack.

Since MD5 is vulnerable to "birthday" attacks and may be compromised, routing protocol authentication must use FIPS 198-1 validated algorithms and modules to encrypt the authentication key. This requirement applies to all IPv4 and IPv6 protocols that are used to exchange routing or packet forwarding information; this includes all Interior Gateway Protocols (such as OSPF, EIGRP, and IS-IS) and Exterior Gateway Protocols (such as BGP), MPLS-related protocols (such as LDP), and multicast-related protocols.

Satisfies: SRG-NET-000168-RTR-000077, SRG-NET-000168-RTR-000078

Audit:

Review the Arista router configuration to verify it is using a NIST-validated FIPS 198-1 message authentication code algorithm to authenticate routing protocol messages.

OSPFv2 Example:

Verify the Message-Digest authentication configuration on the interface for OSPF with the following command:

```
sh run int ethernet YY

interface Eth12
 ip ospf authentication message-digest
 ip ospf message-digest-key 23 sha256 0 arista123
```

If MD5 is present in the output, this is a finding.

BGP Example:

Step 1: Arista router must verify the system clock is set to the correct date/time zone and clock source as this will affect the key lifetimes.

```
router#show clock
Wed Oct 5 14:25:53 2022
Timezone: US/Pacific
Clock source: NTP server (192.168.10.25)
```

If the clock is incorrect, this is a finding.

Step 2: Arista router must verify the management security profile is configured and the session shared-secrets Lifetime durations are specified for the required rotation order and must not exceed 180 days.

```
router#sh man sec session shared-secret profile BGP-SHA1
Profile: BGP-SHA1

Current receive secret: ID: 5, Expires: December 05 2022, 14:34 UTC
Current transmit secret: ID: 5, Expires: December 05 2022, 14:34 UTC

Receive secret rotation order: 5, 10, 15
Transmit secret rotation order: 5, 10, 15

Secrets:
   ID 15
      Secret: $1c$rcKS3MQ9sre00iXfxDVMEg==
      Receive lifetime: March 05 2023, 14:34 UTC to infinite
      Transmit lifetime: March 05 2023, 14:34 UTC to infinite
   ID 10
      Secret: $1c$rcKS3MQ9srcBunxwqKkGEw==
      Receive lifetime: December 05 2022, 14:34 UTC to March 05 2023, 14:34
UTC
      Transmit lifetime: December 05 2022, 14:34 UTC to March 05 2023, 14:34
UTC
   ID 5
      Secret: $1c$rcKS3MQ9srd9RAMH9iKmPQ==
      Receive lifetime: October 05 2022, 14:34 UTC to December 05 2022, 14:34
UTC
      Transmit lifetime: October 05 2022, 14:34 UTC to December 05 2022,
14:34 UTC
```

If the management security profile is not configured, this is a finding.

If the key lifetime exceeds 180 days, this is a finding.

Step 3: Arista router must verify the BGP peer group is configured to use the BGP security profile for the configured BGP neighbor peer. BGP Graceful-restart must be configured in the event the peer restarts to allow TCP resets and prevent clearing of traffic keys. Arista router supports graceful-restart restart-time configuration <1-3600> default [300 seconds].

```
show running-config | section bgp 65000
router bgp 65000
   router-id 10.11.11.11
   graceful-restart restart-time 300
   graceful-restart
   neighbor Peer_Leaf peer group
   neighbor Peer_Leaf remote-as 65000
   neighbor Peer_Leaf next-hop-self
   neighbor Peer_Leaf send-community extended
   neighbor Peer_Leaf maximum-routes 12000
   neighbor Peer_Leaf2 peer group
   neighbor Peer_Leaf2 remote-as 200
   neighbor Peer_Leaf2 next-hop-self
   neighbor Peer_Leaf2 send-community extended
   neighbor Peer_Leaf2 maximum-routes 12000
   neighbor 1.1.1.1 password 7 kEFkx0nsheXsR5ICROtOfB==
   neighbor 1.1.1.1 maximum-routes 12000
   neighbor 2.2.2.2 peer group Peer_Leaf2
   neighbor 2.2.2.2 password shared-secret profile BGP-SHA1 algorithm aes-
128-cmac-96
   neighbor 2.2.2.2 maximum-routes 12000
   neighbor 10.11.12.2 peer group Peer_Leaf
   no neighbor 10.11.12.2 route-map out
   redistribute connected route-map loopback
```

If BGP is not configured to use the security profile, this is a finding.

Remediation:

Configure routing protocol authentication to use a NIST-validated FIPS 198-1 message authentication code algorithm.

OSPFv2:

```
router(config)#interface Eth12
router(config-int-Eth12)#ip ospf authentication message-digest
router(config-int-Eth12)#ip ospf message-digest-key 23 sha256 0 arista123
```

BGP:

Step 1: The Arista router must configure the system clock, which will affect the valid key for a given profile and should be used with caution.

```
router(config)#clock set hh:mm:ss  Current time
```

Step 2: The Arista router must be configured for management security profile <profile_name> and keys for BGP neighbor sessions.

```
router(config)#management security
router(config-man-sec)#session shared-secret profile BGP-SHA
router(config-man-sec-sh-sec-profile-BGP-SHA)#secret 5 password1 2022-10-05
14:34:01 2022-12-05 14:34:01
router(config-man-sec-sh-sec-profile-BGP-SHA)#secret 10 password2 2022-12-05
14:34:01 2023-03-05 14:34:01
router(config-man-sec-sh-sec-profile-BGP-SHA)#secret 15 password3 2023-10-05
14:34:01 <not to exceed 180 days>
```

Step 3: The Arista router must configure the BGP Neighbor to select the profile for use in TCP AO.

```
router(config)#router bgp 65000
router(config-router-bgp)#neighbor 2.2.2.2 password shared-secret profile
BGP-SHA1 algorithm aes-128-cmac-96
router(config-router-bgp)#exit
router(config)#write memory
```

Additional Information:

CCI-000803 Implement mechanisms for authentication to a cryptographic module that meet the requirements of applicable laws, Executive Orders, directives, policies, regulations, standards, and guidance for such authentication.

- NIST SP 800-53::IA-7
- NIST SP 800-53A::IA-7.1
- NIST SP 800-53 Revision 4::IA-7
- NIST SP 800-53 Revision 5::IA-7

CCI-002205 Uniquely identify and authenticate source by organization, system, application, service, and/or individual for information transfer.

- NIST SP 800-53 Revision 4::AC-4 (17)
- NIST SP 800-53 Revision 5::AC-4 (17)

1.25 ARST-RT-000290 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The MPLS router with RSVP-TE enabled must be configured with message pacing or refresh reduction to adjust maximum number of RSVP messages to an output queue based on the link speed and input queue size of adjacent core routers.

```
GROUP ID: V-256011
RULE ID: SV-256011r882375
```

Rationale:

RSVP-TE can be used to perform constraint-based routing when building LSP tunnels within the network core that will support QoS and traffic engineering requirements. RSVP-TE is also used to enable MPLS Fast Reroute, a network restoration mechanism that will reroute traffic onto a backup LSP in case of a node or link failure along the primary path. When there is a disruption in the MPLS core, such as a link flap or router reboot, the result is a significant amount of RSVP signaling, such as "PathErr" and "ResvErr" messages that need to be sent for every LSP using that link.

When RSVP messages are sent out, they are sent either hop by hop or with the router alert bit set in the IP header. This means that every router along the path must examine the packet to determine if additional processing is required for these RSVP messages. If there is enough signaling traffic in the network, it is possible for an interface to receive more packets for its input queue than it can hold, resulting in dropped RSVP messages and hence slower RSVP convergence. Increasing the size of the interface input queue can help prevent dropping packets; however, there is still the risk of having a burst of signaling traffic that can fill the queue. Solutions to mitigate this risk are RSVP message pacing or refresh reduction to control the rate at which RSVP messages are sent. RSVP refresh reduction includes the following features: RSVP message bundling, RSVP Message ID to reduce message processing overhead, Reliable delivery of RSVP messages using Message ID, and summary refresh to reduce the amount of information transmitted every refresh interval.

Audit:

Arista MLS router by default protects RSVP bandwidth by using Refresh Overhead Reduction (RFC 2961).

Review the router configuration to verify the router has been configured to prevent a burst of RSVP traffic engineering signaling messages from overflowing the input queue of any neighbor core router.

The command "refresh method bundled" is the default and enabled, and will not show up in the configuration. However, it can be turned off by "refresh method explicit".

```
sh run | sec mpls rsvp

mpls rsvp
   refresh method explicit
   no shutdown
```

If the Arista router is configured with "refresh method explicit" to disable Refresh Overhead Reduction, this is a finding.

Remediation:

Configure the router for Refresh Overhead Reduction if using RSVP.

Refresh Overhead Reduction (RFC 2961) can be enabled to support sending message IDs and refreshing state with refresh messages by setting the refresh method to "bundled".

```
(config-mpls-rsvp)# no shutdown
(config-mpls-rsvp)# refresh method bundled
```

This is also the default setting. The above command will reset back to default.

Additional Information:

CCI-001095 Manage capacity, bandwidth, or other redundancy to limit the effects of information flooding types of denial-of-service attacks.

- NIST SP 800-53::SC-5 (2)
- NIST SP 800-53A::SC-5 (2).1
- NIST SP 800-53 Revision 4::SC-5 (2)
- NIST SP 800-53 Revision 5::SC-5 (2)

1.26 ARST-RT-000300 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The PE router must be configured to enforce a Quality-of-Service (QoS) policy to limit the effects of packet flooding denial-of-service (DoS) attacks.

```
GROUP ID: V-256012
RULE ID: SV-256012r882378
```

Rationale:

DoS is a condition when a resource is not available for legitimate users. Packet flooding distributed denial-of-service (DDoS) attacks are referred to as volumetric attacks and have the objective of overloading a network or circuit to deny or seriously degrade performance, which denies access to the services that normally traverse the network or circuit. Volumetric attacks have become relatively easy to launch using readily available tools such as Low Orbit Ion Cannon or botnets.

Measures to mitigate the effects of a successful volumetric attack must be taken to ensure that sufficient capacity is available for mission-critical traffic. Managing capacity may include, for example, establishing selected network usage priorities or quotas and enforcing them using rate limiting, Quality of Service (QoS), or other resource reservation control methods. These measures may also mitigate the effects of sudden decreases in network capacity that are the result of accidental or intentional physical damage to telecommunications facilities (such as cable cuts or weather-related outages).

Audit:

Review the Arista router configuration and interview the System Administrator to verify that a mechanism for traffic prioritization and bandwidth reservation exists.

This arrangement must ensure that sufficient capacity is available for mission-critical traffic and enforce the traffic priorities specified by the Combatant Commands/Services/Agencies.

Review the Arista router configuration to verify traffic prioritization and bandwidth reservations.

```
router#sh run | sec bandwidth
interface Port-Channel33
   tx-queue 0
      bandwidth percent 20
   tx-queue 1
      bandwidth percent 40
   tx-queue 2
      bandwidth percent 15
   tx-queue 3
      bandwidth percent 25
policy-map type copp copp-system-policy
   class copp-system-lldp
      bandwidth kbps 500
```

If no such scheme exists or it is not configured, this is a finding.

Remediation:

Implement a mechanism for traffic prioritization and bandwidth reservation. This mechanism must enforce the traffic priorities specified by the Combatant Commands/Services/Agencies.

Step 1: Configure the Arista router for traffic queuing based on traffic prioritization and bandwidth reservation.

```
router(config)#vlan 4033
   trunk group 4033
!
interface Port-Channel33
   description L3-PO33->Distro1-QFX5200-32C-100G
   routerport trunk allowed vlan 2100-2102,4033
   routerport mode trunk
   routerport trunk group 4033
   qos trust dscp
   !
   tx-queue 0
      bandwidth percent 20
   !
   tx-queue 1
      bandwidth percent 40
      shape rate 40088888
   !
   tx-queue 2
      bandwidth percent 15
      shape rate 15022222
  !
   tx-queue 3
      bandwidth percent 25
      shape rate 25250000
```

Step 2: Configure the Arista router differentiated services code point (DSCP) with a 6-bit field in the IP header, which marks all traffic for protocol-specific traffic with the configured DSCP value.

```
router(config)#qos map dscp 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23 24 25 26
27 28 29 30 31 32 33 34 35 36 37 38 40 41 42 43 44 46 47 48 50 51 52 53 54 55
56 57 58 59 60 61 62 63 to traffic-class 0
qos map dscp 11 to traffic-class 1
qos map dscp 39 to traffic-class 2
qos map dscp 15 49 to traffic-class 3
!
```

Step 3: Configure the Arista router trusted routed links with the quality of service port trust mode on the Ethernet interface.

```
router(config)#interface Ethernet 2
description OSPF LINK TO DODIN ENCLAVE
   no routerport
   ip address 172.16.50.1/30
   ipv6 nd ra hop-limit 32
   ip access-group STIG in
   ip ospf authentication message-digest
   ip ospf message-digest-key 1 md5 7 OQ62NhxhqcbWEps4eZjZOg==
   ipv6 ospf encryption ipsec spi 1 esp null sha1 passphrase 7
ZauLr6BwU+Q1MGMLbbys9A==
   qos trust dscp
```

Additional Information:

CCI-001095 Manage capacity, bandwidth, or other redundancy to limit the effects of information flooding types of denial-of-service attacks.

- NIST SP 800-53::SC-5 (2)
- NIST SP 800-53A::SC-5 (2).1
- NIST SP 800-53 Revision 4::SC-5 (2)
- NIST SP 800-53 Revision 5::SC-5 (2)

1.27 ARST-RT-000320 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The PE router must be configured to enforce a Quality-of-Service (QoS) policy in accordance with the QoS GIG Technical Profile.

```
GROUP ID: V-256014
RULE ID: SV-256014r882384
```

Rationale:

Different applications have unique requirements and toleration levels for delay, jitter, bandwidth, packet loss, and availability. To manage the multitude of applications and services, a network requires a QoS framework to differentiate traffic and provide a method to manage network congestion. The Differentiated Services Model (DiffServ) is based on per-hop behavior by categorizing traffic into different classes and enabling each node to enforce a forwarding treatment to each packet as dictated by a policy.

Packet markings such as IP Precedence and its successor, Differentiated Services Code Points (DSCP), were defined along with specific per-hop behaviors for key traffic types to enable a scalable QoS solution. DiffServ QoS categorizes network traffic, prioritizes it according to its relative importance, and provides priority treatment based on the classification. It is imperative that end-to-end QoS is implemented within the IP core network to provide preferred treatment for mission-critical applications.

Audit:

Review the router configuration and verify that a QoS policy has been configured to provide preferred treatment for mission-critical applications in accordance with the QoS GIG Technical Profile.

Step 1: Verify the Arista router class-maps are configured to match on DSCP, protocols, or access control lists (ACLs) that identify traffic types based on ports.

```
router#sh qos map
qos map dscp 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
32 33 34 35 36 37 38 40 41 42 43 44 46 47 48 50 51 52 53 54 55 56 57 58 59 60
61 62 63 to traffic-class 0
qos map dscp 11 to traffic-class 1
qos map dscp 39 to traffic-class 2
qos map dscp 15 49 to traffic-class 3
!
```

Step 2: Verify the Arista router policy-map is configured to set DSCP values for the defined class-maps in accordance with the QoS GIG Technical Profile.

```
router#sh run |sec port-channel33
interface Port-Channel33
description L3-Routed-PO33->Distro1-QFX5200-32C-100G
routerport trunk allowed vlan 2100-2102,4033
routerport mode trunk
routerport trunk group 4033
qos trust dscp
!
tx-queue 0
bandwidth percent 20
!
tx-queue 1
bandwidth percent 40
shape rate 40088888
!
tx-queue 2
bandwidth percent 15
shape rate 15022222
!
tx-queue 3
bandwidth percent 25
shape rate 25250000
Step 3: Verify that an output service policy is bound to all interfaces.
router#sh run | section qos
interface Vlan33
service-policy type qos input policy3
hardware tcam
profile QOS_TEST
feature acl port ip
sequence 45
key size limit 160
key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-port src-
ip tcp-control ttl
action count drop mirror
packet ipv4 forwarding bridged
packet ipv4 forwarding routed
packet ipv4 forwarding routed multicast
packet ipv4 mpls ipv4 forwarding mpls decap
packet ipv4 mpls ipv6 forwarding mpls decap
packet ipv4 non-vxlan forwarding routed decap
packet ipv4 vxlan eth ipv4 forwarding routed decap
packet ipv4 vxlan forwarding bridged decap
!
feature acl port ip egress mpls-tunnelled-match
sequence 100
!
feature acl port ipv6
sequence 25
key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-ops-3b
l4-src-port src-ipv6-high src-ipv6-low tcp-control
action count drop mirror
packet ipv6 forwarding bridged
packet ipv6 forwarding routed
packet ipv6 forwarding routed multicast
packet ipv6 ipv6 forwarding routed decap
!
feature acl port ipv6 egress
```

```
sequence 110
key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port l4-src-
port src-ipv6-high src-ipv6-low tcp-control
action count drop
packet ipv6 forwarding bridged
packet ipv6 forwarding routed
!
feature acl port mac
sequence 55
key size limit 160
key field dst-mac ether-type src-mac
action count drop mirror
packet ipv4 forwarding bridged
packet ipv4 forwarding routed
packet ipv4 forwarding routed multicast
packet ipv4 mpls ipv4 forwarding mpls decap
packet ipv4 mpls ipv6 forwarding mpls decap
packet ipv4 non-vxlan forwarding routed decap
packet ipv4 vxlan forwarding bridged decap
packet ipv6 forwarding bridged
packet ipv6 forwarding routed
packet ipv6 forwarding routed decap
packet ipv6 forwarding routed multicast
packet ipv6 ipv6 forwarding routed decap
packet mpls forwarding bridged decap
packet mpls ipv4 forwarding mpls
packet mpls ipv6 forwarding mpls
packet mpls non-ip forwarding mpls
packet non-ip forwarding bridged
!
```

Step 4: Verify the Arista router is configured for a minimum four queues, 0 through 3, for (Port-Channel33) as round robin, with voice strict-priority. The allocated bandwidth for queue (0) 19.6%, queue (1) 39.6%, queue (2) 14.9%, and queue (3) 24.9%. The bandwidth percentages allow for control-plane and protocol management traffic. These configurations allow burst traffic levels and shape rates for maximum outbound traffic bandwidth per queue.

```
router#sh qos int po33
Port-Channel33:
Trust Mode: DSCP
Default COS: 0
Default DSCP: 0
Port shaping rate: enabled
Tx Bandwidth Bandwidth Shape Rate Priority ECN/WRED
Queue (percent) Guaranteed (units) (units)

7 - - ( - ) - ( - ) SP D

6 - - ( - ) - ( - ) SP D

5 - - ( - ) - ( - ) SP D

4 - - ( - ) - ( - ) SP D

3 25 - ( - ) 24.9 (Gbps) SP D

2 15 - ( - ) 14.8 (Gbps) RR D

1 40 - ( - ) 39.6 (Gbps) RR D

0 20 - ( - ) - ( - ) RR D
Legend:
RR -> Round Robin
SP -> Strict Priority
-> Not Applicable / Not Configured
% -> Percentage of reference
ECN/WRED: L -> Queue Length ECN Enabled W -> WRED Enabled D -> Disabled
```

Note: The GTP QOS document (GTP-0009) can be downloaded via the following link: https://intellipedia.intelink.gov/wiki/Portal:GIG_Technical_Guidance/GTG_GTPs/GTP_Development_List

If the Arista router is not configured to implement a QoS policy in accordance with the QoS GIG Technical Profile, this is a finding.

Remediation:

Configure a QoS policy on each router in accordance with the QoS GIG Technical Profile.

Step 1: Configure the Arista router class-maps to match on DSCP Quality of Service Differentiated Service Code Points (DSCP) values to identify four traffic-class into Class 0 (0-7, 16-38, 40-44, 46-48, 50-63) Class 1 (11) Class 2 (39) Class 3 (15, 49).

```
router(config)#qos map
qos map dscp 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
32 33 34 35 36 37 38 40 41 42 43 44 46 47 48 50 51 52 53 54 55 56 57 58 59 60
61 62 63 to traffic-class 0
qos map dscp 11 to traffic-class 1
qos map dscp 39 to traffic-class 2
qos map dscp 15 49 to traffic-class 3
!
```

Step 2: Configure the Arista router bandwidth and shape rates based on four queues defined by DSCP and the defined class-maps in accordance with the QoS GIG Technical Profile.

```
router(config)#interface Port-Channel33

router(config-if-po33)#description PO33->Distro1-QFX5200-32C-100G

   routerport trunk allowed vlan 2100-2102,4033
   routerport mode trunk
   routerport trunk group 4033
   qos trust dscp
   !
   tx-queue 0
      bandwidth percent 20
   !
   tx-queue 1
      bandwidth percent 40
      shape rate 40088888
   !
   tx-queue 2
      bandwidth percent 15
      shape rate 15022222
   !
   tx-queue 3
      bandwidth percent 25
      shape rate 25250000
!
```

Step 3: Configure the Arista router for queues 0 through 3 for Interface (Port-Channel33) as round robin, with voice strict-priority, and then allocate bandwidth for four queues. queue (0) 19.6%, queue (1) 39.6%, queue (2) 14.9%, and queue (3) 24.9%. allowing for control-plane and protocol management traffic. These configurations allow burst traffic levels and shape rates for maximum outbound traffic bandwidth per queue.

```
router#sh qos int po33
Port-Channel33:
   Trust Mode: DSCP
   Default COS: 0
   Default DSCP: 0
   Port shaping rate: enabled
  Tx    Bandwidth     Bandwidth                    Shape Rate         Priority
ECN/WRED
 Queue  (percent)     Guaranteed (units)            (units)
 ------------------------------------------------------------------------------
------------
   7        -             - ( - )              -            ( - )      SP
D
   6        -             - ( - )              -            ( - )      SP
D
   5        -             - ( - )              -            ( - )      SP
D
   4        -             - ( - )              -            ( - )      SP
D
   3        25          - ( - )           24.9      (Gbps)   SP        D
   2        15          - ( - )           14.9      (Gbps)   RR       D
   1        40          - ( - )           39.6      (Gbps)   RR       D
   0        20          - ( - )              -           ( - )       RR
D

Legend:
RR -> Round Robin
SP -> Strict Priority
 - -> Not Applicable / Not Configured
 % -> Percentage of reference
ECN/WRED: L -> Queue Length ECN Enabled     W -> WRED Enabled     D ->
Disabled
```

Additional Information:

CCI-001095 Manage capacity, bandwidth, or other redundancy to limit the effects of information flooding types of denial-of-service attacks.

- NIST SP 800-53::SC-5 (2)
- NIST SP 800-53A::SC-5 (2).1
- NIST SP 800-53 Revision 4::SC-5 (2)
- NIST SP 800-53 Revision 5::SC-5 (2)

1.28 ARST-RT-000310 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The PE router must be configured to enforce a Quality-of-Service (QoS) policy in accordance with the QoS DODIN Technical Profile.

```
GROUP ID: V-256013
RULE ID: SV-256013r882381
```

Rationale:

Different applications have unique requirements and toleration levels for delay, jitter, bandwidth, packet loss, and availability. To manage the multitude of applications and services, a network requires a QoS framework to differentiate traffic and provide a method to manage network congestion. The Differentiated Services Model (DiffServ) is based on per-hop behavior by categorizing traffic into different classes and enabling each node to enforce a forwarding treatment to each packet as dictated by a policy.

Packet markings such as IP Precedence and its successor, Differentiated Services Code Points (DSCP), were defined along with specific per-hop behaviors for key traffic types to enable a scalable QoS solution. DiffServ QoS categorizes network traffic, prioritizes it according to its relative importance, and provides priority treatment based on the classification. It is imperative that end-to-end QoS is implemented within the IP core network to provide preferred treatment for mission-critical applications.

Audit:

Step 1: Verify the Arista router class-maps are configured to match on DSCP, protocols, or access control lists (ACLs) that identify traffic types based on ports.

```
router#sh qos map
qos map dscp 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
32 33 34 35 36 37 38 40 41 42 43 44 46 47 48 50 51 52 53 54 55 56 57 58 59 60
61 62 63 to traffic-class 0
qos map dscp 11 to traffic-class 1
qos map dscp 39 to traffic-class 2
qos map dscp 15 49 to traffic-class 3
!
```

Step 2: Verify the Arista router policy-map is configured to set DSCP values for the defined class-maps in accordance with the QoS GIG Technical Profile.

```
router#sh run |sec port-channel33
interface Port-Channel33
   description L3-Routed-PO33->Distro1-QFX5200-32C-100G
   routerport trunk allowed vlan 2100-2102,4033
   routerport mode trunk
   routerport trunk group 4033
   qos trust dscp
   !
   tx-queue 0
      bandwidth percent 20
   !
   tx-queue 1
      bandwidth percent 40
      shape rate 40088888
   !
   tx-queue 2
      bandwidth percent 15
      shape rate 15022222
   !
   tx-queue 3
      bandwidth percent 25
      shape rate 25250000
!
```

Step 3: Verify an output service policy is bound to all interfaces.

```
router#sh run | section qos
interface Vlan33
   service-policy type qos input policy3
hardware tcam
   profile QOS_TEST
      feature acl port ip
         sequence 45
         key size limit 160
         key field dscp dst-ip ip-frag ip-protocol l4-dst-port l4-ops l4-src-
port src-ip tcp-control ttl
         action count drop mirror
         packet ipv4 forwarding bridged
         packet ipv4 forwarding routed
         packet ipv4 forwarding routed multicast
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan eth ipv4 forwarding routed decap
         packet ipv4 vxlan forwarding bridged decap
      !
feature acl port ip egress mpls-tunnelled-match
         sequence 100
      !
      feature acl port ipv6
         sequence 25
         key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port
l4-ops-3b l4-src-port src-ipv6-high src-ipv6-low tcp-control
         action count drop mirror
         packet ipv6 forwarding bridged
         packet ipv6 forwarding routed
         packet ipv6 forwarding routed multicast
         packet ipv6 ipv6 forwarding routed decap
      !
      feature acl port ipv6 egress
         sequence 110
         key field dst-ipv6 ipv6-next-header ipv6-traffic-class l4-dst-port
l4-src-port src-ipv6-high src-ipv6-low tcp-control
         action count drop
         packet ipv6 forwarding bridged
         packet ipv6 forwarding routed
      !
      feature acl port mac
         sequence 55
         key size limit 160
         key field dst-mac ether-type src-mac
         action count drop mirror
         packet ipv4 forwarding bridged
         packet ipv4 forwarding routed
         packet ipv4 forwarding routed multicast
         packet ipv4 mpls ipv4 forwarding mpls decap
         packet ipv4 mpls ipv6 forwarding mpls decap
         packet ipv4 non-vxlan forwarding routed decap
         packet ipv4 vxlan forwarding bridged decap
         packet ipv6 forwarding bridged
         packet ipv6 forwarding routed
         packet ipv6 forwarding routed decap
         packet ipv6 forwarding routed multicast
```

```
         packet ipv6 ipv6 forwarding routed decap
         packet mpls forwarding bridged decap
         packet mpls ipv4 forwarding mpls
         packet mpls ipv6 forwarding mpls
         packet mpls non-ip forwarding mpls
         packet non-ip forwarding bridged
      !
```

Step 4: Verify the Arista router is configured for a minimum of four queues, 0 through 3, for (Port-Channel33) as round robin, with voice strict-priority. The allocated bandwidth is queue (0) 19.6%, queue (1) 39.6%, queue (2) 14.9%, and queue (3) 24.9%. The bandwidth percentages allow for control-plane and protocol management traffic. These configurations allow burst traffic levels and shape rates for maximum outbound traffic bandwidth per queue.

```
router#sh qos int po33
Port-Channel33:
   Trust Mode: DSCP
   Default COS: 0
   Default DSCP: 0
   Port shaping rate: enabled
  Tx    Bandwidth     Bandwidth                    Shape Rate         Priority
ECN/WRED
 Queue  (percent)     Guaranteed (units)            (units)
 ------------------------------------------------------------------------------
------------
   7        -             - ( - )              -            ( - )      SP
D
   6        -             - ( - )              -            ( - )      SP
D
   5        -             - ( - )              -            ( - )      SP
D
   4        -             - ( - )              -            ( - )      SP
D
   3        25          - ( - )           24.9      (Gbps)   SP        D
   2        15          - ( - )           14.8      (Gbps)   RR       D
   1        40          - ( - )           39.6      (Gbps)   RR       D
   0        20          - ( - )              -           ( - )       RR
D

Legend:
RR -> Round Robin
SP -> Strict Priority
 - -> Not Applicable / Not Configured
 % -> Percentage of reference
ECN/WRED: L -> Queue Length ECN Enabled     W -> WRED Enabled     D ->
Disabled
```

If the Arista router is not configured to implement a QoS policy in accordance with the QoS DODIN Technical Profile, this is a finding.

Remediation:

Step 1: Configure the Arista router class-maps to match on DSCP Quality of Service values to identify four traffic-class into Class 0 (0-7, 16-38, 40-44, 46-48, 50-63) Class 1 (11) Class 2 (39) Class 3 (15, 49).

```
router(config)#qos map
qos map dscp 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
32 33 34 35 36 37 38 40 41 42 43 44 46 47 48 50 51 52 53 54 55 56 57 58 59 60
61 62 63 to traffic-class 0
qos map dscp 11 to traffic-class 1
qos map dscp 39 to traffic-class 2
qos map dscp 15 49 to traffic-class 3
!
```

Step 2: Configure the Arista router bandwidth and shape rates based on four queues defined by DSCP and the defined class-maps in accordance with the QoS GIG Technical Profile.

```
router(config)#interface Port-Channel33
router(config-if-po33)#description PO33->Distro1-QFX5200-32C-100G
   routerport trunk allowed vlan 2100-2102,4033
   routerport mode trunk
   routerport trunk group 4033
   qos trust dscp
   !
   tx-queue 0
      bandwidth percent 20
   !
   tx-queue 1
      bandwidth percent 40
      shape rate 40088888
   !
   tx-queue 2
      bandwidth percent 15
      shape rate 15022222
   !
   tx-queue 3
      bandwidth percent 25
      shape rate 25250000
!
```

Step 3: Configure the Arista router for queues 0 through 3 for Interface (Port-Channel33) as round robin, with voice strict-priority, and then allocate bandwidth for four queues: queue (0) 19.6%, queue (1) 39.6%, queue (2) 14.9%, and queue (3) 24.9%, allowing for control-plane and protocol management traffic. These configurations allow burst traffic levels and shape rates for maximum outbound traffic bandwidth per queue.

```
router#sh qos int po33
Port-Channel33:
   Trust Mode: DSCP
   Default COS: 0
   Default DSCP: 0
   Port shaping rate: enabled
  Tx    Bandwidth     Bandwidth                    Shape Rate         Priority
ECN/WRED
 Queue  (percent)     Guaranteed (units)            (units)
 ------------------------------------------------------------------------------
------------
   7        -             - ( - )              -            ( - )      SP
D
   6        -             - ( - )              -            ( - )      SP
D
   5        -             - ( - )              -            ( - )      SP
D
   4        -             - ( - )              -            ( - )      SP
D
   3        25          - ( - )           24.9      (Gbps)   SP        D
   2        15          - ( - )           14.9      (Gbps)   RR       D
   1        40          - ( - )           39.6      (Gbps)   RR       D
   0        20          - ( - )              -           ( - )       RR
D

Legend:
RR -> Round Robin
SP -> Strict Priority
 - -> Not Applicable / Not Configured
 % -> Percentage of reference
ECN/WRED: L -> Queue Length ECN Enabled     W -> WRED Enabled     D ->
Disabled
```

Additional Information:

CCI-001095 Manage capacity, bandwidth, or other redundancy to limit the effects of information flooding types of denial-of-service attacks.

- NIST SP 800-53::SC-5 (2)
- NIST SP 800-53A::SC-5 (2).1
- NIST SP 800-53 Revision 4::SC-5 (2)
- NIST SP 800-53 Revision 5::SC-5 (2)

1.29 ARST-RT-000330 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista perimeter router must be configured to deny network traffic by default and allow network traffic by exception.

```
GROUP ID: V-256015
RULE ID: SV-256015r882387
```

Rationale:

A deny-all, permit-by-exception network communications traffic policy ensures that only connections that are essential and approved are allowed.

This requirement applies to both inbound and outbound network communications traffic. All inbound and outbound traffic must be denied by default. Firewalls and perimeter routers should only allow traffic through that is explicitly permitted. The initial defense for the internal network is to block any traffic at the perimeter that is attempting to make a connection to a host residing on the internal network. In addition, allowing unknown or undesirable outbound traffic by the firewall or router will establish a state that will permit the return of this undesirable traffic inbound.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify the access control list (ACL) or filter is configured to allow specific ports and protocols and deny all other traffic.

The filter must be configured inbound on all external interfaces.

Step 1: Verify the ACL is configured to allow traffic per the requirement and deny all by default. Execute the command "sh ip access-list".

```
ip access-list INBOUND
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit tcp 10.10.10.0/24 any eq www https
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 2: Verify the ACL is applied inbound on all external interfaces. Execute the command "sh run int Eth YY".

```
interface ethernet 13
  ip access-group INBOUND in
```

If the ACL or filter is not configured to allow specific ports and protocols and deny all other traffic, this is a finding.

If the filter is not configured inbound on all external interfaces, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista perimeter router to deny network traffic by default and allow network traffic by exception.

Step 1: Configure an ACL inbound to allow traffic per the requirement and deny all by default.

```
ip access-list INBOUND
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit tcp 10.10.10.0/24 any eq www https
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 2: Apply the ACL inbound on all external interfaces.

```
router(config)#interface ethernet 13
router(config-if-Et13)#ip access-group INBOUND in
```

Additional Information:

CCI-001109 Deny network communications traffic by default and allow network communications traffic by exception at managed interfaces; and/or for organization-defined systems.

- NIST SP 800-53::SC-7 (5)
- NIST SP 800-53A::SC-7 (5).1 (i) (ii)
- NIST SP 800-53 Revision 4::SC-7 (5)
- NIST SP 800-53 Revision 5::SC-7 (5)

1.30 ARST-RT-000340 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista router must be configured to restrict traffic destined to itself.

```
GROUP ID: V-256016
RULE ID: SV-256016r882390
```

Rationale:

The route processor handles traffic destined to the router, the key component used to build forwarding paths that is also instrumental with all network management functions. Hence, any disruption or denial-of-service (DoS) attack to the route processor can result in mission-critical network outages.

Audit:

Review the access control list (ACL) or filter for the Arista router receive path and verify that it will only process specific management plane and control plane traffic from specific sources.

Note: If the platform does not support the receive path filter, verify all layer 3 interfaces have an ingress ACL to control what packets are allowed to be destined to the router for processing.

Step 1: Review the Arista router configuration for Control Plane ACL, which can be modified to permit or deny additional protocols that can be matched via an extended access-list for management traffic. Sample Default CP ACL:

```
!
ip access-list ENCLAVE_GATEWAY_FILTER
   10 permit ip any 172.16.0.0/16
!
ip access-list STIG
   10 deny ip 172.16.50.0/30 10.10.100.0/24
!
ip access-list control-plane-modified
   !! Line 180 added for AMLS-L3-000260 which requires eBGP GTSM or
equivalent
   statistics per-entry
   10 permit icmp any any
   20 permit ip any any tracked
   30 permit udp any any eq bfd ttl eq 255
   40 permit udp any any eq bfd-echo ttl eq 254
   50 permit ospf any any
   60 permit tcp any any eq ssh telnet www snmp bgp https msdp
   70 permit udp any any eq bootps bootpc snmp rip ntp
   80 permit tcp any any eq mlag ttl eq 255
   90 permit udp any any eq mlag ttl eq 255
   100 permit vrrp any any
   110 permit ahp any any
   120 permit pim any any
   130 permit igmp any any
   140 permit tcp any any range 5900 5910
   150 permit tcp any any range 50000 50100
   160 permit udp any any range 51000 51100
   170 permit tcp any any eq mlag-arp-sync ttl eq 255
   180 permit tcp 192.168.1.0/30 192.168.1.0/30 eq bgp ttl eq 255 log
```

Step 2: Verify the control plane policy is configured to restricting the LLDP traffic to CPU.

```
router#show running-config | section policy-map
policy-map type copp copp-system-policy
   class copp-system-lldp
   bandwidth kbps 500
```

Step 3: To verify the ACL is configured to allow the traffic per the requirement and deny all by default, execute the command "sh ip access-list".

```
router#show ip access-list
ip access-list INBOUND
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit tcp 10.10.10.0/24 any eq www https
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 4: To verify the ACL is applied inbound on all external interfaces, execute the command "sh run int Eth YY".

```
router#show running-config interface Ethernet 13
interface ethernet 13
  ip access-group INBOUND in
```

If the Arista router is not configured with a receive-path filter to restrict traffic destined to itself, this is a finding.

Remediation:

Configure all Arista routers with receive path filters to restrict traffic destined to the router.

Step 1: Configure the Control plane policy to restrict the LLDP traffic to CPU.

```
router(config)#policy-map type copp copp-system-policy
router(config-pmap-copp-system-policy)#class copp-system-lldp
router(config-pmap-c-copp-system-policy-copp-system-lldp)#bandwidth kbps 500
```

Step 2: Configure an ACL inbound to allow traffic per the requirement and deny all by default.

```
ip access-list INBOUND
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit tcp 10.10.10.0/24 any eq www https
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 3: Apply the ACL inbound on all external interfaces.

```
router(config)#interface ethernet 13
router(config-if-Et13)#ip access-group INBOUND in
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.31 ARST-RT-000350 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to drop all fragmented Internet Control Message Protocol (ICMP) packets destined to itself.

```
GROUP ID: V-256017
RULE ID: SV-256017r882393
```

Rationale:

Fragmented ICMP packets can be generated by hackers for DoS attacks such as Ping O' Death and Teardrop. It is imperative that all fragmented ICMP packets are dropped.

Audit:

Review the access control list (ACL) or filter for the Arista router receive path.

Verify it will drop all fragmented ICMP packets destined to itself.

Step 1: To verify the ACL is configured to filter the fragmented ICMP packets destined to itself, execute the command "sh ip access-list".

```
ip access-list ICMP_FRAGMENTS
 10 deny ip any any fragments
 20 permit ip any any
```

Step 2: To verify the ACL is applied to the external interfaces, execute the command "sh run int Eth YY".

```
interface ethernet 5
 ip access-group ICMP_FRAGMENTS in
```

If the Arista router is not configured with a receive-path filter to drop all fragmented ICMP packets, this is a finding.

Note: If the platform does not support the receive path filter, verify all layer 3 interfaces have an ingress ACL to control what packets are allowed to be destined to the router for processing.

Remediation:

Ensure all Arista routers have their receive path filter configured to drop all fragmented ICMP packets.

Step 1: Configure the ACL to filter the fragmented ICMP packets destined to itself.

```
LEAF-1A(config)#ip access-list ICMP_FRAGMENTS
LEAF-1A(config-acl-ICMP_FRAGMENTS)# 10 deny ip any any fragments
LEAF-1A(config-acl-ICMP_FRAGMENTS)# 20 permit ip any any
LEAF-1A(config-acl-ICMP_FRAGMENTS)# exit
```

Step 2: Apply the ACL to the external interfaces.

```
LEAF-1A(config)#interface ethernet 5
LEAF-1A(config-if-Et5)# ip access-group ICMP_FRAGMENTS in
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.32 ARST-RT-000380 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to filter egress traffic at the internal interface on an inbound direction.

```
GROUP ID: V-256019
RULE ID: SV-256019r882399
```

Rationale:

Access lists are used to separate data traffic into that which it will route (permitted packets) and that which it will not route (denied packets). Secure configuration of routers makes use of access lists for restricting access to services on the router itself as well as for filtering traffic passing through the router.

Inbound versus Outbound: It should be noted that some operating systems default access lists are applied to the outbound queue. The more secure solution is to apply the access list to the inbound queue for three reasons:

- The router can protect itself before damage is inflicted.
- The input port is still known and can be filtered upon.
- It is more efficient to filter packets before routing them.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify the egress ACL is bound to the internal interface in an inbound direction.

Step 1: To verify the egress ACL is bound to the internal interface in an inbound direction, execute the command "sh ip access-list".

```
ip access-list WAN_OUT
   10 permit tcp any host 180.20.10.1 eq ssh telnet
   20 permit ip any 190.16.10.0/24
   30 permit udp any 67.56.10.2 eq bootps snmp
   40 deny tcp any 208.73.210.0 0.0.1.255
   50 deny udp any 208.73.210.0 0.0.1.255
   60 deny icmp any any
   70 permit ip any any
```

Step 2: To verify the ACL is applied inbound on all external interfaces, execute the command "sh run int Eth YY".

```
interface ethernet 8
  ip access-group WAN_OUT in
```

If the Arista router is not configured to filter traffic leaving the network at the internal interface in an inbound direction, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Step 1: Configure an egress ACL bound to the internal interface in an inbound direction to filter traffic leaving the network.

```
Leaf2(config)#ip access-list WAN_OUT
Leaf2(config-acl-WAN_OUT)#10 permit tcp any host 180.20.10.1 eq ssh telnet
Leaf2(config-acl-WAN_OUT)#20 permit ip any 190.16.10.0/24
Leaf2(config-acl-WAN_OUT)#30 permit udp any host 67.56.10.2 eq bootps snmp
Leaf2(config-acl-WAN_OUT)#40 deny tcp any 208.73.210.0 0.0.1.255
Leaf2(config-acl-WAN_OUT)#50 deny udp any 208.73.210.0 0.0.1.255
Leaf2(config-acl-WAN_OUT)#60 deny icmp any any
Leaf2(config-acl-WAN_OUT)#70 permit ip any any
```

Step 2: Apply the ACL inbound on all external interfaces.

```
Leaf2(config)#interface ethernet 8
Leaf2(config-if-Et8)#  ip access-group WAN_OUT in
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.33 ARST-RT-000390 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to reject outbound route advertisements for any prefixes belonging to the IP core.

```
GROUP ID: V-256020
RULE ID: SV-256020r882402
```

Rationale:

Outbound route advertisements belonging to the core can result in traffic either looping or being black holed, or at a minimum, using a non-optimized path.

Audit:

Review the Arista router configuration to verify a filter is defined to block route advertisements for prefixes that belong to the IP core.

The prefix filter must be referenced outbound on the appropriate BGP neighbor statements.

Step 1: The following example creates an outbound route advertise filter and configures CE Arista MLS to advertise the filter to IP Core PE (100.1.0.128). An IP prefix list named FILTER_OUT is created to specify the 172.16.1.0/24 subnet for outbound route advertisements filtering.

```
ip prefix-list FILTER_OUT seq 10 permit 172.16.1.0/24
```

Step 2: Verify the outbound prefix list is applied to the appropriate BGP neighbor in the BGP process. Execute the command "sh run section router bgp".

```
router bgp 65001
 neighbor 100.1.0.128 remote-as 65200
 neighbor 100.1.0.128 prefix-list FILTER_OUT out
 exit
```

If the Arista router is not configured to reject outbound route advertisements that belong to the IP core, this is a finding.

Remediation:

Configure all eBGP Arista routers to filter outbound route advertisements belonging to the IP core.

Step 1: Configure an outbound route advertise filter and configure CE Arista MLS to advertise the filter to IP Core PE (100.1.0.128). Also configure an IP prefix list named FILTER_OUT to specify the 172.16.1.0/24 subnet for outbound route advertisements filtering.

```
LEAF-1A(config)#ip prefix-list FILTER_OUT seq 10 permit 172.16.1.0/24
```

Step 2: Apply the prefix-list outbound with the BGP neighbor in BGP process.

```
LEAF-1A(config)#router bgp 65001
LEAF-1A(config-router-bgp)#neighbor 100.1.0.128 remote-as 65200
LEAF-1A(config-router-bgp)#neighbor 100.1.0.128 prefix-list FILTER_OUT out
LEAF-1A(config-router-bgp)# exit
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.34 ARST-RT-000370 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to filter ingress traffic at the external interface on an inbound direction.

```
GROUP ID: V-256018
RULE ID: SV-256018r882396
```

Rationale:

Access lists are used to separate data traffic into that which it will route (permitted packets) and that which it will not route (denied packets). Secure configuration of routers makes use of access lists for restricting access to services on the router itself as well as for filtering traffic passing through the router.

Inbound versus Outbound: It should be noted that some operating systems default access lists are applied to the outbound queue. The more secure solution is to apply the access list to the inbound queue for three reasons:

- The router can protect itself before damage is inflicted.
- The input port is still known and can be filtered upon.
- It is more efficient to filter packets before routing them.

Satisfies: SRG-NET-000205-RTR-000003, SRG-NET-000205-RTR-000004

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify the ingress ACL is bound to the external interface in an inbound direction.

Step 1: To verify the ingress ACL is bound to the external interface in an inbound direction, execute the command "sh ip access-list".

```
ip access-list INBOUND
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit ip 10.10.10.0/24 192.168.10.0/24
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 2: To verify the ACL is applied inbound on all external interfaces, execute the command "sh run int Eth YY".

```
interface ethernet 13
  ip access-group INBOUND in
```

If the Arista router is not configured to filter traffic entering the network at the external interface in an inbound direction, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Bind the ingress ACL to the external interface (inbound).

Step 1: Configure the ACL.

```
LEAF-1A(config)#ip access-list INBOUND
LEAF-1A(config-acl-INBOUND)#10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq
ssh telnet
LEAF-1A(config-acl-INBOUND)#20 permit ip 10.10.10.0/24 192.168.10.0/24
LEAF-1A(config-acl-INBOUND)#30 permit udp 10.20.20.0/24 any eq bootps snmp
LEAF-1A(config-acl-INBOUND)#40 deny ip any any log
```

Step 2: Apply the ACL inbound on all external interfaces.

```
LEAF-1A(config)#interface ethernet 13
LEAF-1A(config-if-Et13)#ip access-group INBOUND in
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.35 ARST-RT-000400 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista router must be configured to block any traffic that is destined to IP core infrastructure.

```
GROUP ID: V-256021
RULE ID: SV-256021r882405
```

Rationale:

IP/MPLS networks providing VPN and transit services must provide, at the least, the same level of protection against denial-of-service (DoS) attacks and intrusions as Layer 2 networks. Although the IP core network elements are hidden, security should never rely entirely on obscurity.

IP addresses can be guessed. Core network elements must not be accessible from any external host. Protecting the core from any attack is vital for the integrity and privacy of customer traffic as well as the availability of transit services. A compromise of the IP core can result in an outage or, at a minimum, non-optimized forwarding of customer traffic. Protecting the core from an outside attack also prevents attackers from using the core to attack any customer. Hence, it is imperative that all routers at the edge deny traffic destined to any address belonging to the IP core infrastructure.

Audit:

Review the Arista router configuration to verify an ingress ACL is applied to all CE-facing interfaces.

Verify the ingress ACL rejects and logs packets destined to the IP core address block.

Note: Internet Control Message Protocol (ICMP) echo requests and traceroutes will be allowed to the edge from external adjacent peers.

Step 1: Verify the ingress ACL is configured to drop any traffic with destination address assigned to the IP core infrastructure. Execute the command "sh ip access-list".

```
ip access-list DROP_INBOUND
deny ip any 172.16.0.0/16 log
permit icmp any any
permit ip any any
```

Step 2: To verify the ingress ACL applied to all CE facing interfaces inbound to drop all the traffic coming toward the CE, execute the command "sh run int Eth YY".

```
interface Ethernet 2
ip access-group DROP_INBOUND in
```

If the Arista PE router is not configured to block any traffic with a destination address assigned to the IP core infrastructure, this is a finding.

Remediation:

Configure protection for the IP core to be implemented at the edges by blocking any traffic with a destination address assigned to the IP core infrastructure.

Step 1: Configure an ingress ACL to drop any traffic with destination address assigned to the IP core infrastructure.

```
router(config)#ip access-list DROP_INBOUND
router(config-acl-DROP_INBOUND)#deny ip any 172.16.0.0/16 log
router(config-acl-DROP_INBOUND)#permit icmp any any
router(config-acl-DROP_INBOUND)#permit ip any any
router(config-acl-DROP_INBOUND)#exit
```

Step 2: Apply the ACL to all CE-facing interfaces inbound to drop all the traffic coming toward the CE.

```
router(config)#interface Ethernet 2
router(config-if-Et2)#ip access-group DROP_INBOUND in
router(config-if-Et2)#end
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.36 ARST-RT-000410 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured with Unicast Reverse Path Forwarding (uRPF) loose mode enabled on all CE-facing interfaces.

```
GROUP ID: V-256022
RULE ID: SV-256022r882408
```

Rationale:

The uRPF feature is a defense against spoofing and denial-of-service (DoS) attacks by verifying if the source address of any ingress packet is reachable. To mitigate attacks that rely on forged source addresses, all provider edge routers must enable uRPF loose mode to guarantee that all packets received from a CE router contain source addresses that are in the route table.

Audit:

Review the Arista router configuration to determine if uRPF loose mode is enabled on all CE-facing interfaces.

To verify the interface configuration uRPF loose mode is enabled on all CE-facing interfaces, execute the command "sh run int Eth YY".

```
interface Ethernet 3/17/1
ip address 10.10.22.1/30
ip verify unicast source reachable-via any
```

If uRPF loose mode is not enabled on all CE-facing interfaces, this is a finding.

Remediation:

Enable uRPF loose mode on all CE-facing interfaces.

Configure uRPF loose mode on all CE-facing interfaces.

```
router(config)#interface Ethernet 3/17/1
router(config-if-Et3/17/1)#ip verify unicast source reachable-via any
router(config-if-Et3/17/1)#end
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.37 ARST-RT-000420 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The out-of-band management (OOBM) Arista gateway router must be configured to forward only authorized management traffic to the Network Operations Center (NOC).

```
GROUP ID: V-256023
RULE ID: SV-256023r991786
```

Rationale:

The OOBM network is an IP network used exclusively for the transport of OAM&P data from the network being managed to the OSS components located at the NOC. Its design provides connectivity to each managed network device, enabling network management traffic to flow between the managed network elements and the NOC. This allows the use of paths separate from those used by the managed network.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the network topology diagram to determine connectivity between the managed network and the NOC.

Review the OOBM gateway router configuration to validate the path that the management traffic traverses.

Step 1: To verify only management traffic is forwarded through the OOBM interface or IPsec tunnel, execute the command "sh ip access-list".

```
ip access-list OOBM_to_MGMT
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh telnet
   20 permit ip 10.10.10.0/24 192.168.10.0/24
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 2: To verify the ACL is applied outbound on the OOBM interface, execute the command "sh run int Eth YY".

```
interface ethernet 1
  description OOBM to MGMT link
  ip access-group OOBM_to_MGMT out
```

If traffic other than authorized management traffic is permitted through the OOBM interface or IPsec tunnel, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure filters based on port, source IP address, and destination IP address to permit only authorized management traffic into IPsec tunnels or the OOBM interface used for forwarding management data.

Step 1: To configure an ACL to allow only management traffic to be forwarded through the OOBM interface or IPsec tunnel, execute the command "sh ip access-list".

```
ip access-list OOBM_to_MGMT
   10 permit tcp 10.10.10.0/24 host 10.20.10.1 eq ssh
   20 permit ip 10.10.10.0/24 192.168.10.0/24
   30 permit udp 10.20.20.0/24 any eq bootps snmp
```

Step 2: To apply the ACL outbound on the OOBM interface, execute the command "sh run int Eth YY".

```
interface ethernet 1
  description OOBM to MGMT link
  ip access-group OOBM_to_MGMT out
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.38 ARST-RT-000430 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The out-of-band management (OOBM) Arista gateway router must be configured to block any traffic destined to itself that is not sourced from the OOBM network or the NOC.

```
GROUP ID: V-256024
RULE ID: SV-256024r991787
```

Rationale:

If the gateway router is not a dedicated device for the OOBM network, several safeguards must be implemented for containment of management and production traffic boundaries. It is imperative that hosts from the managed network are not able to access the OOBM gateway router.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the access control list (ACL) or filter for the router receive path.

Verify only traffic sourced from the OOBM network or the NOC is allowed to access the router.

Note: If the platform does not support the receive path filter, verify all non-OOBM interfaces have an ingress ACL to restrict access to that interface address or any of the router's loopback addresses to only traffic sourced from the management network. An exception would be to allow packets destined to these interfaces used for troubleshooting, such as ping and traceroute.

Step 1: To verify the ACL is configured, execute "show ip access-list OOBM_ACL".

Step 2: Determine the NOC management network subnet, which is 172.16.12.0/24.

Step 3: Verify the ACL restricts all management plane traffic.

```
ip access-list OOBM_ACL
   permit tcp 192.168.10.0/24 any eq ssh
   permit udp host 172.16.12.42 any eq snmp
   permit udp host 172.16.12.41 any eq ntp
   permit icmp 172.16.12.0/24 any
   deny ip any any log
```

Step 4: To verify the ACL is applied ingress on the OOBM interface, execute the command "sh run int Eth YY".

```
interface ethernet 1
  description LAN link
  ip access-group OOBM_ACL in
```

If the Arista router does not block any traffic destined to itself that is not sourced from the OOBM network or the NOC, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Ensure traffic from the managed network is not able to access the OOBM gateway router using either receive path or interface ingress ACLs.

Step 1: Configure the ACL to restrict all management plane traffic.

```
ip access-list OOBM_ACL
   permit tcp 192.168.10.0/24 any eq ssh
   permit udp host 172.16.12.42 any eq snmp
   permit udp host 172.16.12.41 any eq ntp
   permit icmp 172.16.12.0/24 any
   deny ip any any log
```

Step 2: Apply the ACL ingress. Execute the command "sh run int Eth YY".

```
interface ethernet 1
  description LAN link
  ip access-group OOBM_ACL in
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.39 ARST-RT-000440 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to only permit management traffic that ingresses and egresses the out-of-band management (OOBM) interface.

```
GROUP ID: V-256025
RULE ID: SV-256025r991788
```

Rationale:

The OOBM access router will connect to the management interface of the managed network elements. The management interface can be a true OOBM interface or a standard interface functioning as the management interface. In either case, the management interface of the managed network element will be directly connected to the OOBM network.

An OOBM interface does not forward transit traffic, thereby providing complete separation of production and management traffic. Since all management traffic is immediately forwarded into the management network, it is not exposed to possible tampering. The separation also ensures that congestion or failures in the managed network do not affect the management of the device. If the device does not have an OOBM port, the interface functioning as the management interface must be configured so that management traffic does not leak into the managed network and that production traffic does not leak into the management network.

Audit:

Step 1: To verify the Arista router managed interface has an inbound and outbound ACL configured, execute "show run int Eth YY".

```
router#show interface Ethernet 3
interface ethernet 3
 ip access-group FILTER_INBOUND in
 ip access-group FILTER_OUTBOUND out
```

Step 2: To verify the ingress filter only allows management, IGP, and ICMP traffic, execute "show ip access-list".

```
router#show ip access-list
ip access-list FILTER_INBOUND
   10 permit ospf any any
   20 permit icmp any any echo
   30 permit icmp any any echo-reply
   40 permit ip 10.10.10.0/24 any
   50 deny ip any any log
!
ip access-list FILTER_OUTBOUND
   10 permit ospf any any
   20 permit icmp any any echo
   30 permit icmp any any echo-reply
   40 permit ip any 10.10.10.0/24
   50 deny ip any any log
!
```

Note: If the management interface is a true OOBM interface, this requirement is not applicable.

If the Arista router does not restrict traffic that ingresses and egresses the management interface, this is a finding.

Remediation:

If the Arista management interface is a routed interface, it must be configured with both an ingress and egress ACL.

Step 1: Configure the ingress filter to only allow management, IGP, and ICMP traffic.

```
LEAF-1A(config)#ip access-list FILTER_INBOUND
LEAF-1A(config-acl-FILTER_INBOUND)#permit ospf any any
LEAF-1A(config-acl-FILTER_INBOUND)#permit icmp any any echo
LEAF-1A(config-acl-FILTER_INBOUND)#permit icmp any any echo-reply
LEAF-1A(config-acl-FILTER_INBOUND)#permit ip 10.10.10.0/24 any
LEAF-1A(config-acl-FILTER_INBOUND)#deny ip any any log
```

Step 2: Configure the outbound filter to only allow management, IGP, and ICMP traffic.

```
LEAF-1A(config)#ip access-list FILTER_OUTBOUND
LEAF-1A(config-acl-FILTER_OUTBOUND)#permit ospf any any
LEAF-1A(config-acl-FILTER_OUTBOUND)#permit icmp any any echo
LEAF-1A(config-acl-FILTER_OUTBOUND)#permit icmp any any echo-reply
LEAF-1A(config-acl-FILTER_OUTBOUND)#permit ip any 10.10.10.0/24
LEAF-1A(config-acl-FILTER_OUTBOUND)#deny ip any any log
```

Step 3: Set the managed interface to have an inbound and outbound ACL configured.

```
LEAF-1A(config)#interface ethernet 3
LEAF-1A(config-if-Et3)# ip access-group FILTER_INBOUND in
LEAF-1A(config-if-Et3)# ip access-group FILTER_OUTBOUND out
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.40 ARST-RT-000450 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Arista perimeter router must be configured to restrict it from accepting outbound IP packets that contain an illegitimate address in the source address field via egress filter or by enabling Unicast Reverse Path Forwarding (uRPF).

```
GROUP ID: V-256026
RULE ID: SV-256026r945858
```

Rationale:

A compromised host in an enclave can be used by a malicious platform to launch cyberattacks on third parties. This is a common practice in "botnets", a collection of compromised computers using malware to attack other computers or networks. DDoS attacks frequently leverage IP source address spoofing to send packets to multiple hosts that in turn will then send return traffic to the hosts with the IP addresses that were forged. This can generate significant amounts of traffic. Therefore, protection measures to counteract IP source address spoofing must be taken. When uRPF is enabled in strict mode, the packet must be received on the interface that the device would use to forward the return packet; thereby mitigating IP source address spoofing.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the router configuration to verify uRPF or an egress filter has been configured on all internal interfaces to restrict the router from accepting outbound IP packets that contain an illegitimate address in the source address field.

To verify restrict uRPF is configured on the interface to restrict the router from accepting outbound IP packets that contain an illegitimate address in the source address field, execute the command "sh run int eth YY".

```
interface Ethernet3/17/1
ip address 172.16.43.3/24
ip verify unicast source reachable-via rx
```

If uRPF or an egress filter to restrict the router from accepting outbound IP packets that contain an illegitimate address in the source address field has not been configured on all internal interfaces, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to ensure an egress filter or uRPF is configured to restrict the router from accepting any outbound IP packet that contains an external IP address in the source field.

```
LEAF-1A(config)#interface Ethernet3/17/1
LEAF-1A(config)#ip address 172.16.43.3/24
LEAF-1A(config)#ip verify unicast source reachable-via rx
```

Additional Information:

CCI-001094 Restrict the ability of individuals to launch organization-defined denial of service attacks against other systems.

- NIST SP 800-53::SC-5 (1)
- NIST SP 800-53A::SC-5 (1).1
- NIST SP 800-53 Revision 4::SC-5 (1)
- NIST SP 800-53 Revision 5::SC-5 (1)

1.41 ARST-RT-000470 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to use a unique key for each autonomous system (AS) that it peers with.

```
GROUP ID: V-256027
RULE ID: SV-256027r945862
```

Rationale:

If the same keys are used between eBGP neighbors, the chance of a hacker compromising any of the BGP sessions increases. It is possible that a malicious user exists in one autonomous system who would know the key used for the eBGP session. This user would then be able to hijack BGP sessions with other trusted neighbors.

Audit:

Interview the ISSM and router administrator to determine if unique keys are being used.

Verify the BGP router AS is configured for a unique key. Run the command "sh run | section router bgp".

```
router bgp [NN]
neighbor [ip address] password [type] [password-string]
```

If unique keys are not being used, this is a finding.

Remediation:

Configure all eBGP Arista routers with unique keys for each eBGP neighbor that it peers with.

To configure BGP authentication, in the BGP configuration mode interface, when adding neighbors, include the following statement:

```
router bgp NN
neighbor 100.1.0.0  password 0 [password-string]
```

Additional Information:

CCI-002205 Uniquely identify and authenticate source by organization, system, application, service, and/or individual for information transfer.

- NIST SP 800-53 Revision 4::AC-4 (17)
- NIST SP 800-53 Revision 5::AC-4 (17)

1.42 ARST-RT-000480 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The PE router providing MPLS Layer 2 Virtual Private Network (L2VPN) services must be configured to authenticate targeted Label Distribution Protocol (LDP) sessions used to exchange virtual circuit (VC) information using a FIPS-approved message authentication code algorithm.

```
GROUP ID: V-256028
RULE ID: SV-256028r882426
```

Rationale:

LDP provides the signaling required for setting up and tearing down pseudowires (virtual circuits used to transport Layer 2 frames) across an MPLS IP core network. Using a targeted LDP session, each PE router advertises a virtual circuit label mapping that is used as part of the label stack imposed on the frames by the ingress PE router during packet forwarding. Authentication provides protection against spoofed TCP segments that can be introduced into the LDP sessions.

Audit:

Review the Arista router configuration to determine if LDP messages are being authenticated for the targeted LDP sessions.

Step 1: Verify the Arista router configuration to verify LDP is configured globally and router-id is set.

```
mpls ldp
  router-id [x.x.x.x | interface] LoopbackY
  no shutdown
```

Step 2: Verify the Arista router configuration to ensure the password is configured for LDP neighbor.

```
mpls ldp
  password [type] [password]
```

Step 3: Enable the mpls globally or per interface.

For Global:

```
mpls ip
```

For interfaces:

```
interface Ethernet 1
 mpls ip
```

If authentication is not being used for the LDP sessions using a FIPS-approved message authentication code algorithm, this is a finding.

Remediation:

Implement authentication for all targeted LDP sessions using a FIPS-approved message authentication code algorithm.

Step 1: Configure the mpls LDP on the Arista MLS.

```
PE11(config)#mpls ldp
PE11(config-mpls-ldp)#router-id interface Loopback0
PE11(config-mpls-ldp)#no shutdown
```

Step 2: Enable the authentication for LDP neighbors.

```
PE11(config-mpls-ldp)#password 0 xxxxx
```

Additional Information:

CCI-001958 Authenticate organization-defined devices and/or types of devices before establishing a local, remote, and/or network connection.

- NIST SP 800-53 Revision 4::IA-3
- NIST SP 800-53 Revision 5::IA-3

1.43 ARST-RT-000490 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must not be configured to have any zero-touch deployment feature enabled when connected to an operational network.

```
GROUP ID: V-256029
RULE ID: SV-256029r882429
```

Rationale:

Network devices that are configured via a zero-touch deployment or auto-loading feature can have their startup configuration or image pushed to the device for installation via TFTP or Remote Copy (rcp). Loading an image or configuration file from the network is taking a security risk because the file could be intercepted by an attacker who could corrupt the file, resulting in a denial of service.

Audit:

Review the Arista MLS device configuration to determine if a configuration auto-loading or zero-touch deployment feature is enabled.

Execute the command "sh zerotouch".

```
Zerotouch Mode: Disabled
```

If a configuration auto-loading feature or zero-touch deployment feature is enabled, this is a finding.

Remediation:

Disable all configuration auto-loading or zero-touch deployment features.

```
LEAF-1A(config)#zerotouch disable
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.44 ARST-RT-000510 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to have gratuitous ARP disabled on all external interfaces.

```
GROUP ID: V-256030
RULE ID: SV-256030r882432
```

Rationale:

A gratuitous ARP is an ARP broadcast in which the source and destination MAC addresses are the same. It is used to inform the network about a host IP address. A spoofed gratuitous ARP message can cause network mapping information to be stored incorrectly, causing network malfunction.

Audit:

Review the configuration to determine if gratuitous ARP is disabled on all external interfaces.

By default, Arista router interfaces reject gratuitous ARP request packets. To verify the gratuitous ARP is disabled, execute the command "sh run int ethernet YY".

These commands enable/disable gratuitous ARP packet acceptance on.

Enable

```
interface Ethernet 2
  arp gratuitous accept

Disable

interface Ethernet 2
  no arp gratuitous accept
```

If gratuitous ARP is enabled on any external interface, this is a finding.

Remediation:

Disable gratuitous ARP on all external interfaces.

Disable

```
interface Ethernet 2
  no arp gratuitous accept
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.45 ARST-RT-000520 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista router must be configured to have IP directed broadcast disabled on all interfaces.

```
GROUP ID: V-256031
RULE ID: SV-256031r884234
```

Rationale:

An IP-directed broadcast is a datagram sent to the broadcast address of a subnet that is not directly attached to the sending machine. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet, where it is converted into a link-layer broadcast. Because of the nature of the IP addressing architecture, only the last router in the chain, which is connected directly to the target subnet, can conclusively identify a directed broadcast.

IP-directed broadcasts are used in the extremely common and popular smurf, or denial-of-service (DoS) attacks. In a smurf attack, the attacker sends Internet Control Message Protocol (ICMP) echo requests from a falsified source address to a directed broadcast address, causing all the hosts on the target subnet to send replies to the falsified source. By sending a continuous stream of such requests, the attacker can create a much larger stream of replies, which can completely inundate the host whose address is being falsified. This service should be disabled on all interfaces when not needed to prevent smurf and DoS attacks.

Directed broadcast can be enabled on internal facing interfaces to support services such as Wake-On-LAN. Case scenario may also include support for legacy applications where the content server and the clients do not support multicast. The content servers send streaming data using UDP broadcast. Used in conjunction with the IP multicast helper-map feature, broadcast data can be sent across a multicast topology. The broadcast streams are converted to multicast and vice versa at the first-hop routers and last-hop routers before entering and leaving the multicast transit area respectively. The last-hop router must convert the multicast to broadcast. Hence, this interface must be configured to forward a broadcast packet (i.e., a directed broadcast address is converted to the nodes' broadcast address).

Audit:

Review the Arista router configuration to determine if IP directed broadcast is enabled.

By default, IP directed broadcast is disabled on Arista multi-layer router. To verify the IP directed broadcast is enabled, execute the command:

```
sh run int ethernet <YY>

interface Ethernet 2
 ip address 10.1.12.1/24
 no ip directed-broadcast
```

If IP directed broadcast is enabled on layer 3 interfaces, this is a finding.

Remediation:

Disable IP directed broadcasts on all layer 3 interfaces.

```
LEAF-1A(config)#interface Ethernet 2
LEAF-1A(config-if-Et2)# ip address 10.1.12.1/24
LEAF-1A(config-if-Et2)# no ip directed-broadcast
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.46 ARST-RT-000530 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to have Internet Control Message Protocol (ICMP) unreachable notifications disabled on all external interfaces.

```
GROUP ID: V-256032
RULE ID: SV-256032r882438
```

Rationale:

The ICMP supports IP traffic by relaying information about paths, routes, and network conditions. Routers automatically send ICMP messages under a wide variety of conditions. Host unreachable ICMP messages are commonly used by attackers for network mapping and diagnosis.

Audit:

Review the device configuration to determine if controls have been defined to ensure the router does not send ICMP unreachable notifications out to any external interfaces.

Step 1: To verify the ACL is configured to determine the router does not send ICMP unreachable notifications out to any external interfaces, execute the command "sh ip access-list".

```
ip access-group DENY_ICMP_UNREACHABLE
 deny icmp any any unreachable
 permit ip any any
```

Step 2: To verify the ACL is applied outbound on interfaces, execute the command "sh run int Eth YY".

```
interface Ethernet 2
 ip access-group DENY_ICMP_UNREACHABLE out
```

If ICMP unreachable notifications are enabled on any external interfaces, this is a finding.

Remediation:

Disable ICMP unreachable notifications on all external interfaces.

Step 1: The Arista router can be configured to filter out the ICMP Unreachable for (Type 3) code 0 (Network unreachable) IPv4 and IPv6 packets with the following command:

```
router(config)#ip icmp rate-limit-unreachable 0
router(config)#ipv6 icmp rate-limit-unreachable 0
```

Step 2: The Arista router can be configured to filter out the ICMP Unreachable for (Type 3) code 1 (Network unreachable) IPv4 and IPv6 packets with the following command:

```
router(config)#ip access-list BLK-ICMP-Unreachables
10 deny icmp any any host-unreachable
20 permit ip any any
!
```

Step 3: This would need to be applied on the egress interface (for example as in et1 below):

```
router(config)#interface ethernet1
no routerport
ip address 32.1.1.12/24
ip access-group BLK-ICMP-Unreachables out
!
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.47 ARST-RT-000540 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to have Internet Control Message Protocol (ICMP) mask replies disabled on all external interfaces.

```
GROUP ID: V-256033
RULE ID: SV-256033r882441
```

Rationale:

The ICMP supports IP traffic by relaying information about paths, routes, and network conditions. Routers automatically send ICMP messages under a wide variety of conditions. Mask Reply ICMP messages are commonly used by attackers for network mapping and diagnosis.

Audit:

Review the Arista device configuration to determine if controls have been defined to ensure the router does not send ICMP Mask Reply messages out to any external interfaces.

EOS by default does not respond to ICMP Type 17 or 18.

Step 1: To verify the ACL is configured to determine the router does not send ICMP Mask Reply messages out to any external interfaces, execute the command "sh ip access-list".

```
ip access-group DENY_ICMP_MASK_REPLY
 deny icmp any any mask-reply
 permit ip any any
```

Step 2: To verify the ACL is applied outbound on interfaces, execute the command "sh run int Eth YY".

```
interface Ethernet 2
 ip access-group DENY_ICMP_MASK_REPLY out
```

If ICMP Mask Reply messages are enabled on any external interfaces, this is a finding.

Remediation:

Step 1: Disable ICMP mask replies on all external interfaces.

```
ip access-group DENY_ICMP_MASK_REPLY
 deny icmp any any mask-reply
 permit ip any any
```

Step 2: Apply the ACL outbound on interfaces.

```
interface Ethernet 2
 ip access-group DENY_ICMP_MASK_REPLY out
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.48 ARST-RT-000550 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must be configured to have Internet Control Message Protocol (ICMP) redirects disabled on all external interfaces.

```
GROUP ID: V-256034
RULE ID: SV-256034r882444
```

Rationale:

The ICMP supports IP traffic by relaying information about paths, routes, and network conditions. Routers automatically send ICMP messages under a wide variety of conditions. Redirect ICMP messages are commonly used by attackers for network mapping and diagnosis.

Audit:

Review the device configuration to determine if controls have been defined to ensure the router does not send ICMP Redirect messages out to any external interfaces.

Step 1: To verify the ACL is configured to determine the router does not send ICMP Redirect messages out to any external interfaces, execute the command "sh ip access-list".

```
ip access-group DENY_REDIRECT
 deny icmp any any redirect
 permit ip any any
```

Step 2: To verify the ACL is applied outbound on interface, execute the command "sh run int Eth YY".

```
interface Ethernet 2
 ip access-group DENY_REDIRECT out
```

If ICMP Redirect messages are enabled on any external interfaces, this is a finding.

Remediation:

Step 1: Disable ICMP redirects on all external interfaces.

```
ip access-group DENY_REDIRECT
 deny icmp any any redirect
 permit ip any any
```

Step 2: Apply the ACL outbound on interfaces.

```
interface Ethernet 2
description EXTERNAL INTERFACE
 ip access-group DENY_REDIRECT in
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.49 ARST-RT-000570 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista BGP router must be configured to limit the prefix size on any inbound route advertisement to /24 or the least significant prefixes issued to the customer.

```
GROUP ID: V-256036
RULE ID: SV-256036r882450
```

Rationale:

The effects of prefix de-aggregation can degrade router performance due to the size of routing tables and also result in black-holing legitimate traffic. Initiated by an attacker or a misconfigured router, prefix de-aggregation occurs when the announcement of a large prefix is fragmented into a collection of smaller prefix announcements.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify there is a filter to reject inbound route advertisements that are greater than /24 or the least significant prefixes issued to the customer, whichever is larger.

Step 1: To verify there is a filter to reject inbound route advertisements that are greater than /24 or the least significant prefixes issued to the customer, whichever is larger, execute the command "sh ip prefix-list".

```
ip prefix-list ADVERTISE_ROUTES deny 0.0.0.0/0 ge 25
ip prefix-list ADVERTISE_ROUTES permit 0.0.0.0/0 le 32
```

Step 2: Verify the prefix-list is applied in BGP process. Execute the command "sh run section router bgp".

```
router bgp 65000
 neighbor 10.1.12.2 prefix-list ADVERTISE_ROUTES in
```

If the Arista router is not configured to limit the prefix size on any inbound route advertisement to /24 or the least significant prefixes issued to the customer, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Ensure all eBGP Arista routers are configured to limit the prefix size on any route advertisement to /24 or the least significant prefixes issued to the customer.

Step 1: Configure the prefix-list.

```
ip prefix-list ADVERTISE_ROUTES deny 0.0.0.0/0 ge 25
ip prefix-list ADVERTISE_ROUTES permit 0.0.0.0/0 le 32
```

Step 2: Apply the prefix-list in the BGP process inbound.

```
LEAF-1A(config)#router bgp 65000
LEAF-1A(config)# neighbor 10.1.12.2 prefix-list ADVERTISE_ROUTES in
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.50 ARST-RT-000560 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista BGP router must be configured to use the maximum prefixes feature to protect against route table flooding and prefix de-aggregation attacks.

```
GROUP ID: V-256035
RULE ID: SV-256035r882447
```

Rationale:

The effects of prefix de-aggregation can degrade router performance due to the size of routing tables and also result in black-holing legitimate traffic. Initiated by an attacker or a misconfigured router, prefix de-aggregation occurs when the announcement of a large prefix is fragmented into a collection of smaller prefix announcements.

In 1997, misconfigured routers in the Florida Internet Exchange network (AS7007) de-aggregated every prefix in their routing table and started advertising the first /24 block of each of these prefixes as their own. Faced with this additional burden, the internal routers became overloaded and crashed repeatedly. This caused prefixes advertised by these routers to disappear from routing tables and reappear when the routers came back online. As the routers came back after crashing, they were flooded with the routing table information by their neighbors. The flood of information would again overwhelm the routers and cause them to crash. This process of route flapping served to destabilize not only the surrounding network but also the entire internet. Routers trying to reach those addresses would choose the smaller, more specific /24 blocks first. This caused backbone networks throughout North America and Europe to crash.

Maximum prefix limits on peer connections combined with aggressive prefix-size filtering of customers' reachability advertisements will effectively mitigate the de-aggregation risk. BGP maximum prefix must be used on all eBGP routers to limit the number of prefixes that it should receive from a particular neighbor, whether customer or peering AS. Consider each neighbor and how many routes they should be advertising and set a threshold slightly higher than the number expected.

Audit:

Review the Arista router configuration to verify the number of received prefixes from each eBGP neighbor is controlled.

To verify in the BGP configuration that number of received prefixes from each eBGP neighbor is controlled, execute the command "sh rnu section router bgp".

```
router bgp NNN
 neighbor x.1.12.1 remote-as YYY
 neighbor x.1.12.1 maximum-routes 12000
 neighbor x.1.12.1 maximum-accepted-routes 10000
```

If the Arista router is not configured to control the number of prefixes received from each peer to protect against route table flooding and prefix deaggregation attacks, this is a finding.

Remediation:

Configure all eBGP Arista routers to use the maximum prefixes feature to protect against route table flooding and prefix deaggregation attacks.

Configure each BGP neighbor to control the number of prefixes.

```
router bgp NNN
 neighbor x.1.12.1 remote-as YYY
 neighbor x.1.12.1 maximum-routes 12000
 neighbor x.1.12.1 maximum-accepted-routes 10000
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.51 ARST-RT-000580 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The multicast Rendezvous Point (RP) Arista router must be configured to limit the multicast forwarding cache so that its resources are not saturated by managing an overwhelming number of Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) source-active entries.

```
GROUP ID: V-256037
RULE ID: SV-256037r882453
```

Rationale:

MSDP peering between networks enables sharing of multicast source information. Enclaves with an existing multicast topology using PIM-SM can configure their RP routers to peer with MSDP routers. As a first step of defense against a denial-of-service (DoS) attack, all RP routers must limit the multicast forwarding cache to ensure that router resources are not saturated managing an overwhelming number of PIM and MSDP source-active entries.

Audit:

Review the Arista router configuration to determine if forwarding cache thresholds are defined.

Step 1: To verify the ACL is configured to match the prefixes, execute the command "sh ip access-list".

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

Step 2: To verify the thresholds are defined for multicast forwarding cache for source-active entries, execute the command "sh run sec router msdp".

```
router msdp
 peer 10.1.12.2
  sa-filter in PIM_NEIGHBOR_SA_FILTER
  sa-limit 500
```

If the Arista RP router is not configured to limit the multicast forwarding cache to ensure its resources are not saturated, this is a finding.

Remediation:

Configure the Arista MSDP-enabled RP routers to limit the multicast forwarding cache for source-active entries.

Step 1: Configure the ACL.

```
ip access-list PIM_NEIGHBOR_SA_FILTER
   10 deny ip any 224.1.1.0/24
   20 deny ip any 224.1.2.0/24
   30 deny ip any 224.1.3.0/24
   40 deny ip any 224.1.4.0/24
   100 permit ip any any
```

Step 2: Apply the ACL in MSDP peer and define the multicast forwarding cache for source-active entries.

```
router msdp
 peer 10.1.12.2
  sa-filter in PIM_NEIGHBOR_SA_FILTER
  sa-limit 500
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.52 ARST-RT-000590 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista multicast Designated Router (DR) must be configured to increase the shortest-path tree (SPT) threshold or set it to infinity to minimalize source-group (S, G) state within the multicast topology where Any Source Multicast (ASM) is deployed.

```
GROUP ID: V-256038
RULE ID: SV-256038r945856
```

Rationale:

ASM can have many sources for the same groups (many-to-many). For many receivers, the path via the RP may not be ideal compared with the shortest path from the source to the receiver. By default, the last-hop router will initiate a router from the shared tree to a source-specific SPT to obtain lower latencies. This is accomplished by the last-hop router sending an (S, G) Protocol Independent Multicast (PIM) Join toward S (the source).

When the last-hop router begins to receive traffic for the group from the source via the SPT, it will send a PIM Prune message to the RP for the (S, G). The RP will then send a Prune message toward the source. The SPT routerover becomes a scaling issue for large multicast topologies that have many receivers and many sources for many groups because (S, G) entries require more memory than (*, G). Hence, it is imperative to minimize the amount of (S, G) state to be maintained by increasing the threshold that determines when the SPT routerover occurs.

Audit:

Review the Arista multicast last-hop router configuration to verify the SPT routerover threshold is increased (default is "0") or set to infinity (never router over).

```
router pim sparse-mode
   ipv4
      spt threshold infinity
```

If any Arista multicast router is not configured to increase the SPT threshold or set to infinity to minimalize (S, G) state, this is a finding.

Remediation:

Configure the Arista multicast router to increase the SPT threshold or set it to infinity to minimalize (S, G) state within the multicast topology where ASM is deployed.

```
LEAF-1A(config)#router pim sparse-mode
LEAF-1A(config-router-pim-sparse)#ipv4
LEAF-1A(config-router-pim-sparse-ipv4)#spt threshold infinity
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.53 ARST-RT-000600 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista BGP router must be configured to enable the Generalized TTL Security Mechanism (GTSM).

```
GROUP ID: V-256039
RULE ID: SV-256039r882459
```

Rationale:

GTSM is designed to protect a router's IP-based control plane from DoS attacks. Many attacks focused on CPU load and line-card overload can be prevented by implementing GTSM on all Exterior Border Gateway Protocol speaking routers.

GTSM is based on the fact that the vast majority of control plane peering is established between adjacent routers; that is, the Exterior Border Gateway Protocol peers are either between connecting interfaces or between loopback interfaces. Since TTL spoofing is considered nearly impossible, a mechanism based on an expected TTL value provides a simple and reasonably robust defense from infrastructure attacks based on forged control plane traffic.

Audit:

Review the Arista router configuration.

Arista MLS IP packets to GTSM enabled BGP peers are sent with the configured TTL value of 254.

```
router bgp NNN
 neighbor 10.1.12.2 ttl maximum-hops 2
```

If the Arista router is not configured to use GTSM for all Exterior Border Gateway Protocol peering sessions, this is a finding.

Remediation:

Configure all Exterior Border Gateway Protocol peering sessions to use GTSM.

```
router bgp 65000
 neighbor 10.1.12.2 ttl maximum-hops 2
```

Additional Information:

CCI-002385 Protect against or limit the effects of organization-defined types of denial-of-service events.

- NIST SP 800-53 Revision 4::SC-5
- NIST SP 800-53 Revision 5::SC-5 a

1.54 ARST-RT-000610 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to only allow incoming communications from authorized sources to be routed to authorized destinations.

```
GROUP ID: V-256040
RULE ID: SV-256040r882462
```

Rationale:

Unrestricted traffic may contain malicious traffic that poses a threat to an enclave or to other connected networks. Additionally, unrestricted traffic may transit a network, which uses bandwidth and other resources.

Traffic can be restricted directly by an access control list (ACL), which is a firewall function, or by Policy Routing. Policy Routing is a technique used to make routing decisions based on a number of different criteria other than just the destination network, including source or destination network, source or destination address, source or destination port, protocol, packet size, and packet classification. This overrides the router's normal routing procedures used to control the specific paths of network traffic. It is normally used for traffic engineering but can also be used to meet security requirements; for example, traffic that is not allowed can be routed to the Null0 or discard interface. Policy Routing can also be used to control which prefixes appear in the routing table.

This requirement is intended to allow network administrators the flexibility to use whatever technique is most effective.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to determine if the router allows only incoming communications from authorized sources to be routed to authorized destinations.

Step 1: Verify the ACL is defined as in the following example.

Execute the command "sh ip access-list".

```
ip access-list AUTHORIZED_SOURCES
 permit ip 10.1.12.0/24 any
 deny ip 1.2.3.0/24 any log
```

Step 2: Verify the ACL is applied inbound on the external interface.

```
interface ethernet 3
 ip access-group AUTHORIZED_SOURCES in
```

If the Arista router does not restrict incoming communications to allow only authorized sources and destinations, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to allow only incoming communications from authorized sources to be routed to authorized destinations, such as the example shown below.

```
LEAF-1A(config-acl-AUTHORIZED_SOURCES)# permit ip 10.1.12.0/24 any
LEAF-1A(config-acl-AUTHORIZED_SOURCES)# deny ip 1.2.3.0/24 any log
LEAF-1A(config-acl-AUTHORIZED_SOURCES)#exit

LEAF-1A(config)#interface ethernet 3
LEAF-1A(config-if-Et3)# ip access-group AUTHORIZED_SOURCES in
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.55 ARST-RT-000620 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to block inbound packets with source Bogon IP address prefixes.

```
GROUP ID: V-256041
RULE ID: SV-256041r882465
```

Rationale:

Bogons include IP packets on the public internet that contain addresses that are not in any range allocated or delegated by the Internet Assigned Numbers Authority (IANA) or a delegated regional internet registry (RIR) and allowed for public internet use. Bogons also include multicast, IETF reserved, and special purpose address space as defined in RFC 6890.

Security of the internet's routing system relies on the ability to authenticate an assertion of unique control of an address block. Measures to authenticate such assertions rely on the validation the address block forms as part of an existing allocated address block, and must be a trustable and unique reference in the IANA address registries. The intended use of a Bogon address would only be for the purpose of address spoofing in denial-of-service attacks. Hence, it is imperative that IP packets with a source Bogon address are blocked at the network's perimeter.

Audit:

This requirement is not applicable for the DODIN backbone.

Verify the ingress filter is blocking packets with Bogon source addresses.

Review the Arista router configuration to verify it is configured to block IP packets with a Bogon source address with "show run | section prefix-list".

```
IPv4 Bogon Prefixes
0.0.0.0/8
10.0.0.0/8
100.64.0.0/10
127.0.0.0/8
169.254.0.0/16
172.16.0.0/12
192.0.0.0/24
192.0.2.0/24
192.88.99.0/24
192.168.0.0/16
198.18.0.0/15 |
198.51.100.0/24
203.0.113.0/24
224.0.0.0/4
240.0.0.0/4

IPv6 Bogon Prefixes
::/128
::1/128
0::/96
::ffff:0:0/96
3ffe::/16
64:ff9b::/96
100::/64
2001:10::/28
2001:db8::/32
2001:2::/48
2001::/32
2001::/23
2002::/16
fc00::/7
fe80::/10
fec0::/10
ff00::/8
```

If the Arista router is not configured to block inbound IP packets containing a Bogon source address, this is a finding.

Note: At a minimum, IP packets containing a source address from the special purpose address space as defined in RFC 6890 must be blocked. The 6Bone prefix (3ffe::/16) is also considered a Bogon address. Perimeter routers connected to commercial ISPs for internet or other non-DOD network sources must be reviewed for a full Bogon list.

Step 1: Verify the ACL is configured to block the IPv4 Bogon prefixes.

```
ip access-list BOGON_PREFIXES
deny ip 0.0.0.0/8 any
deny ip 10.0.0.0/8 any
deny ip 100.64.0.0/10 any
deny ip 127.0.0.0/8 any
deny ip 169.254.0.0/16 any
deny ip 172.16.0.0/12 any
deny ip 192.0.0.0/24 any
deny ip 192.0.2.0/24 any
deny ip 192.88.99.0/24 any
deny ip 192.168.0.0/16 any
deny ip 198.18.0.0/15 any
deny ip 198.51.100.0/24 any
deny ip 203.0.113.0/24 any
deny ip 224.0.0.0/4 any
deny ip 240.0.0.0/4 any
```

Step 2: Verify the ACL is configured to block the IPv6 Bogon prefixes.

```
ipv6 access-list BOGON_PREFIXES
deny ipv6 ::/128 any
deny ipv6 ::1/128 any
deny ipv6 0::/96 any
deny ipv6 ::ffff:0:0/96 any
deny ipv6 3ffe::/16 any
deny ipv6 64:ff9b::/96 any
deny ipv6 100::/64 any
deny ipv6 2001:10::/28 any
deny ipv6 2001:db8::/32 any
deny ipv6 2001:2::/48 any
deny ipv6 2001::/32 any
deny ipv6 2001::/23 any
deny ipv6 2002::/16 any
deny ipv6 fc00::/7 any
deny ipv6 fe80::/10 any
deny ipv6 fec0::/10 any
deny ipv6 ff00::/8 any
```

Step 3: Verify the IPv4 and IPv6 access lists are applied to the external interface.

```
interface ethernet 3
ip access-group BOGON_PREFIXES in
ipv6 access-group BOGON_PREFIXES in
```

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to block inbound packets with Bogon source addresses.

Step 1: Configure the ACL to block the IPv4 Bogon prefixes.

```
LEAF-1A(config)#ip access-list BOGON_PREFIXES
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 0.0.0.0/8 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 10.0.0.0/8 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 100.64.0.0/10 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 127.0.0.0/8 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 169.254.0.0/16 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 172.16.0.0/12 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 192.0.0.0/24 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 192.0.2.0/24 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 192.88.99.0/24 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 192.168.0.0/16 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 198.18.0.0/15 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 198.51.100.0/24 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 203.0.113.0/24 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 224.0.0.0/4 any
LEAF-1A(config-acl-BOGON_PREFIXES)#deny ip 240.0.0.0/4 any
LEAF-1A(config-acl-BOGON_PREFIXES)#exit
```

Step 2: Configure the ACL to block the ipv6 Bogon prefixes.

```
LEAF-1A(config)#ipv6 access-list BOGON_PREFIXES
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 ::/128 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 ::1/128 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 0::/96 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 ::ffff:0:0/96 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 3ffe::/16 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 64:ff9b::/96 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 100::/64 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2001:10::/28 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2001:db8::/32 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2001:2::/48 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2001::/32 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2001::/23 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 2002::/16 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 fc00::/7 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 fe80::/10 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 fec0::/10 any
LEAF-1A(config-ipv6-acl-BOGON_PREFIXES)#deny ipv6 ff00::/8 any
```

Step 3: Apply the IPv4 and IPv6 Bogon access lists to the external interface.

```
LEAF-1A(config)#interface ethernet 3
LEAF-1A(config-if-Et3)#ip access-group BOGON_PREFIXES in
LEAF-1A(config-if-Et3)#ipv6 access-group BOGON_PREFIXES in
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.56 ARST-RT-000630 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista perimeter router must be configured to have Link Layer Discovery Protocols (LLDPs) disabled on all external interfaces.

```
GROUP ID: V-256042
RULE ID: SV-256042r882468
```

Rationale:

LLDPs are primarily used to obtain protocol addresses of neighboring devices and discover platform capabilities of those devices. Use of SNMP with the LLDP Management Information Base (MIB) allows network management applications to learn the device type and the SNMP agent address of neighboring devices, thereby enabling the application to send SNMP queries to those devices. LLDPs are also media- and protocol-independent as they run over the data link layer; therefore, two systems that support different network-layer protocols can still learn about each other. Allowing LLDP messages to reach external network nodes is dangerous as it provides an attacker a method to obtain information of the network infrastructure that can be useful to plan an attack.

Audit:

This requirement is not applicable for the DODIN backbone.

Review all Arista router configurations to ensure LLDPs are not included in the global configuration or LLDPs are not included for each active external interface. Examples of LLDPs are Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), and Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED).

Disable LLDP on external interface.

```
int ethernet 3
 no lldp transmit
 no lldp receive
```

On Arista multi-layer routers, the LLDP can be disabled globally.

```
no lldp run
```

If LLDPs are configured globally or on any external interface, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Disable LLDPs on all external interfaces.

```
LEAF-1A(config)#int ethernet 3
LEAF-1A(config-if-Et3)#no lldp transmit
LEAF-1A(config-if-Et3)#no lldp receive
```

Disable LLDP globally.

```
LEAF-1A(config)#no lldp run
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.57 ARST-RT-000640 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to have Proxy ARP disabled on all external interfaces.

```
GROUP ID: V-256043
RULE ID: SV-256043r882471
```

Rationale:

When Proxy ARP is enabled on a Cisco router, it allows that router to extend the network (at Layer 2) across multiple interfaces (LAN segments). Because proxy ARP allows hosts from different LAN segments to look like they are on the same segment, proxy ARP is only safe when used between trusted LAN segments. Attackers can leverage the trusting nature of proxy ARP by spoofing a trusted host and then intercepting packets. Proxy ARP should always be disabled on router interfaces that do not require it unless the router is being used as a LAN bridge.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to determine if IP Proxy ARP is disabled on all external interfaces. Execute the command "sh run int ethernet YY".

```
int ethernet 3
no ip proxy-arp
```

If IP Proxy ARP is enabled on any external interface, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Disable IP Proxy ARP on all external interfaces.

```
LEAF-1A(config)#int ethernet 3
LEAF-1A(config-if-Et3)#no ip proxy-arp
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.58 ARST-RT-000650 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to block all outbound management traffic.

```
GROUP ID: V-256044
RULE ID: SV-256044r945857
```

Rationale:

For in-band management, the management network must have its own subnet in order to enforce control and access boundaries provided by Layer 3 network nodes, such as routers and firewalls. Management traffic between the managed network elements and the management network is routed via the same links and nodes as that used for production or operational traffic. Safeguards must be implemented to ensure that the management traffic does not leak past the perimeter of the managed network.

Audit:

This requirement is not applicable for the DODIN backbone.

The Arista perimeter router of the managed network must be configured with an access control list (ACL) or filter on the egress interface to block all management traffic.

Step 1: To verify the configuration is blocking all outbound traffic destined to management network, execute the command "sh ip access-list".

```
ip access-list FILTER_MANAGEMENT_SUBNET
  deny ip any 172.20.1.0 0.0.0.255 log
  permit ip any any
```

Step 2: To verify the filter is applied on egress interface, execute the command "sh run int ethernet YY".

```
interface ethernet 3
ip access-group FILTER_MANAGEMENT_SUBNET out
```

If management traffic is not blocked at the perimeter, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista perimeter router of the managed network with an ACL or filter on the egress interface to block all outbound management traffic.

Step 1: Configure the filter to block all outbound traffic destined to the management network.

```
LEAF-1A(config-if-Et3)#ip access-list FILTER_MANAGEMENT_SUBNET
LEAF-1A(config-acl-FILTER_MANAGEMENT_SUBNET)#  deny ip any 172.20.1.0
0.0.0.255 log
LEAF-1A(config-acl-FILTER_MANAGEMENT_SUBNET)#  permit ip any any
```

Step 2: Apply the filter egress on the interface.

```
LEAF-1A(config-acl-FILTER_MANAGEMENT_SUBNET)#interface ethernet 3
LEAF-1A(config-if-Et3)#ip access-group FILTER_MANAGEMENT_SUBNET out
```

Additional Information:

CCI-001097 Monitor and control communications at the external managed interfaces to the system and at key managed interfaces within the system.

- NIST SP 800-53::SC-7 a
- NIST SP 800-53A::SC-7.1 (iii)
- NIST SP 800-53 Revision 4::SC-7 a
- NIST SP 800-53 Revision 5::SC-7 a

1.59 ARST-RT-000660 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista multicast Designated Router (DR) must be configured to filter the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Report messages to allow hosts to join only multicast groups that have been approved by the organization.

```
GROUP ID: V-256045
RULE ID: SV-256045r882477
```

Rationale:

Real-time multicast traffic can entail multiple large flows of data. Large unicast flows tend to be fairly isolated (i.e., someone doing a file download here or there), whereas multicast can have broader impact on bandwidth consumption, resulting in extreme network congestion. Hence, it is imperative that there is multicast admission control to restrict which multicast groups hosts are allowed to join via IGMP or MLD.

Audit:

Review the configuration of the Arista DR to verify it is filtering IGMP or MLD report messages, allowing hosts to join only groups that have been approved.

Note: This requirement is only applicable to Source Specific Multicast (SSM) implementation. This requirement is not applicable to Any Source Multicast (ASM) because the filtering is being performed by the Rendezvous Point router.

Step 1: Verify the ACL filters the unauthorized IGMP groups. The ACL below is blocking the IGMP group sourced 232.0.0.0/8. Execute the command "sh ip access-list".

```
ip access-list FILTER_IGMP
  10 deny igmp 232.0.0.0/8 any
  20 permit ip any any
```

Step 2: Verify the ACL is configured on internal host-facing interfaces (IGMP process) to filter IGMP.

```
router igmp
  ip igmp access-group FILTER_IGMP
```

or

```
interface ethernet 3
 ip access-group FILTER_IGMP
```

If the Arista DR is not filtering IGMP or MLD report messages, this is a finding.

Remediation:

Configure the Arista DR to filter the IGMP and MLD report messages to allow hosts to join only multicast groups that have been approved.

Step 1: Configure the ACL to filter the unauthorized IGMP groups.

```
LEAF-1A(config-if-Et3)#ip access-list FILTER_IGMP
LEAF-1A(config-acl-FILTER_IGMP)#  10 deny igmp 232.0.0.0/8 any
LEAF-1A(config-acl-FILTER_IGMP)#  20 permit ip any any
```

Step 2: Configure the IGMP filter in IGMP process.

```
LEAF-1A(config-acl-FILTER_IGMP)#router igmp
LEAF-1A(config-router-igmp)#  ip igmp access-group FILTER_IGMP
```

or

Configure the IGMP filter on internal host-facing interfaces (IGMP process) to filter IGMP.

```
LEAF-1A(config-router-igmp)#interface ethernet 3
LEAF-1A(config-if-Et3)# ip access-group FILTER_IGMP in
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.60 ARST-RT-000680 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista Multicast Source Discovery Protocol (MSDP) router must be configured to only accept MSDP packets from known MSDP peers.

```
GROUP ID: V-256047
RULE ID: SV-256047r882483
```

Rationale:

MSDP peering with customer network routers presents additional risks to the DISN Core, whether from a rogue or misconfigured MSDP-enabled router. To guard against an attack from malicious MSDP traffic, the receive path or interface filter for all MSDP-enabled RP routers must be configured to only accept MSDP packets from known MSDP peers.

Audit:

Review the Arista router configuration to determine if there is a receive path or interface filter to only accept MSDP packets from known MSDP peers.

Step 1: Verify the MSDP peers and the corresponding interfaces.

```
router msdp
  peer 10.11.12.2
  !
  peer 10.22.12.2
```

Step 2: Verify the access-list is configured inbound on MSDP peering interfaces. MSDP uses TCP port 639. Execute the command "sh ip access-list".

```
ip access-list MSDP_FILTER
  10 permit tcp host 10.1.12.2 host 10.11.17.9 eq 639
  20 permit udp host 10.1.12.2 host 10.11.17.9 eq 500
  30 permit udp 10.11.17.9 eq 500 host 10.1.12.2
  40 deny ip any any log
```

Step 3: Verify the ACL is applied on the interface. Execute the command "sh run int ethernet YY".

```
interface ethernet 3
 ip access-group MSDP_FILTER in
```

If the Arista router is not configured to only accept MSDP packets from known MSDP peers, this is a finding.

Remediation:

Ensure the receive path or interface filter for all Arista MSDP routers only accepts MSDP packets from known MSDP peers.

Step 1: Configure the MSDP peers.

```
LEAF-1A(config)#router msdp
LEAF-1A(config-router-msdp)#  peer 10.11.12.2
LEAF-1A(config-router-msdp)#  peer 10.22.12.2
```

Step 2: Configure the access-list inbound on MSDP peering interfaces. MSDP uses TCP port 639.

```
LEAF-1A(config-router-msdp-peer-10.22.12.2)#ip access-list MSDP_FILTER
LEAF-1A(config-acl-MSDP_FILTER)#  10 permit tcp host 10.1.12.2 host
10.11.17.9 eq 639
LEAF-1A(config-acl-MSDP_FILTER)#  20 permit udp host 10.1.12.2 host
10.11.17.9 eq 500
LEAF-1A(config-acl-MSDP_FILTER)#  30 permit udp 10.11.17.9 eq 500 host
10.1.12.2
LEAF-1A(config-acl-MSDP_FILTER)#  40 deny ip any any log
```

Step 3: Apply the ACL on the interface.

```
LEAF-1A(config-acl-MSDP_FILTER)#interface ethernet 3
LEAF-1A(config-if-Et3)# ip access-group MSDP_FILTER in
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.61 ARST-RT-000670 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista multicast Designated Router (DR) must be configured to filter the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Report messages to allow hosts to join a multicast group only from sources that have been approved by the organization.

```
GROUP ID: V-256046
RULE ID: SV-256046r882480
```

Rationale:

Real-time multicast traffic can entail multiple large flows of data. Large unicast flows tend to be fairly isolated (i.e., someone doing a file download here or there), whereas multicast can have broader impact on bandwidth consumption, resulting in extreme network congestion. Hence, it is imperative that there is multicast admission control to restrict which multicast groups hosts are allowed to join via IGMP or MLD.

Audit:

Review the configuration of the Arista DR to verify it is filtering IGMP or MLD report messages, allowing hosts to only join multicast groups from sources that have been approved.

Note: This requirement is only applicable to Source Specific Multicast (SSM) implementation.

Step 1: To verify the ACL filters the unauthorized IGMP joins, execute the command "sh ip access-list".

```
ip access-list standard ALLOWED_SOURCES
  10 permit 232.0.0.0/8
  20 deny any log
```

Step 2: Verify the ACL is configured on internal host-facing interfaces (pim process) to filter IGMP joins.

```
router pim sparse-mode
  ipv4
    ssm range ALLOWED_SOURCES
```

If the Arista DR is not filtering IGMP or MLD report messages, this is a finding.

Remediation:

Configure the Arista DR to filter the IGMP and MLD report messages to allow hosts to join only multicast groups from sources that have been approved.

Step 1: Configure the ACL to filter the unauthorized IGMP groups.

```
LEAF-1A(config-if-Et3)#ip access-list standard ALLOWED_SOURCES
LEAF-1A(config-std-acl-ALLOWED_SOURCES)#  10 permit 232.0.0.0/8
LEAF-1A(config-std-acl-ALLOWED_SOURCES)#  20 deny any log
```

Step 2: Configure the IGMP filter in IGMP process.

```
LEAF-1A(config)#router pim sparse-mode
LEAF-1A(config-router-pim-sparse)#  ipv4
LEAF-1A(config-router-pim-sparse-ipv4)#    ssm range ALLOWED_SOURCES
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.62 ARST-RT-000690 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista BGP router must be configured to use its loopback address as the source address for iBGP peering sessions.

```
GROUP ID: V-256048
RULE ID: SV-256048r991789
```

Rationale:

Using a loopback address as the source address offers a multitude of uses for security, access, management, and scalability of the BGP routers. It is easier to construct appropriate ingress filters for router management plane traffic destined to the network management subnet since the source addresses will be from the range used for loopback interfaces instead of a larger range of addresses used for physical interfaces. Log information recorded by authentication and syslog servers will record the router's loopback address instead of the numerous physical interface addresses.

When the loopback address is used as the source for eBGP peering, the BGP session will be harder to hijack since the source address to be used is not known globally—making it more difficult for a hacker to spoof an eBGP neighbor. By using traceroute, a hacker can easily determine the addresses for an eBGP speaker when the IP address of an external interface is used as the source address. The routers within the iBGP domain should also use loopback addresses as the source address when establishing BGP sessions.

Audit:

Review the Arista router configuration to verify a loopback address has been configured.

Verify a loopback interface is used as the source address for all iBGP sessions.

Step 1: To verify the Loopback interface is defined, execute the command "sh run int loopback YY".

```
interface loopback 0
  ip address 10.1.1.1/32
```

Step 2: To verify a loopback interface is used as the source address for all iBGP sessions, execute the command "sh run sec router bgp".

```
router bgp 65001
   router-id 10.1.1.1
   neighbor Peer_Leaf peer group
   neighbor Peer_Leaf remote-as 65001
   neighbor Peer_Leaf update-source Loopback0
   neighbor 10.2.2.2 peer group Peer_Leaf
```

If the Arista router does not use its loopback address as the source address for all iBGP sessions, this is a finding.

Remediation:

Ensure the Arista router's loopback address is used as the source address when originating traffic.

Step 1: Configure the Loopback interface.

```
LEAF-1A(config)#interface Loopback0
LEAF-1A(config-if-Lo0)#ip address 10.1.1.1/32
```

Step 2: Configure the loopback interface as source for all iBGP sessions.

```
router bgp 65001
LEAF-1A(config-router-bgp)#
LEAF-1A(config-router-bgp)#neighbor Peer_Leaf peer group
LEAF-1A(config-router-bgp)#Peer_Leaf remote-as 65001
LEAF-1A(config-router-bgp)#Peer_Leaf update-source Loopback0
LEAF-1A(config-router-bgp)#10.2.2.2 peer group Peer_Leaf
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.63 ARST-RT-000700 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The MPLS router must be configured to use its loopback address as the source address for LDP peering sessions.

```
GROUP ID: V-256049
RULE ID: SV-256049r991790
```

Rationale:

Using a loopback address as the source address offers a multitude of uses for security, access, management, and scalability of backbone routers. It is easier to construct appropriate ingress filters for router management plane traffic destined to the network management subnet since the source addresses will be from the range used for loopback interfaces instead of from a larger range of addresses used for physical interfaces. Log information recorded by authentication and syslog servers will record the router's loopback address instead of the numerous physical interface addresses.

Audit:

Review the Arista router configuration to determine if it uses its loopback address as the source address for LDP peering sessions.

To verify a loopback address has been configured as shown in the following example, execute the command "sh run int loopback YY".

interface loopback 0

ip address 10.1.1.1/32

An MPLS router will use the LDP router ID as the source address for LDP hellos and when establishing TCP sessions with LDP peers; hence, it is necessary to verify the LDP router ID is the same as the loopback address. By default, routers will assign the LDP router ID using the highest IP address on the router, with preference given to loopback addresses. If the router-id command is specified that overrides this default behavior, verify it is the IP address of the designated loopback interface.

mpls ldp

router-id interface Loopback0

no shutdown

If the Arista router is not configured to use its loopback address for LDP peering, this is a finding.

Remediation:

Configure the Arista MPLS routers to use their loopback address as the source address for LDP peering sessions.

Step 1: Configure the loopback interface.

LEAF-1A(config)#interface Loopback0

LEAF-1A(config-if-Lo0)#ip address 10.1.1.1/32

Step 2: Configure the loopback interface as LDP router-id.

LEAF-1A(config)#mpls ldp

LEAF-1A(config-mpls-ldp)#router-id interface Loopback0

LEAF-1A(config-mpls-ldp)#no shutdown

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.64 ARST-RT-000710 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The MPLS router must be configured to synchronize IGP and LDP to minimize packet loss when an IGP adjacency is established prior to LDP peers completing label exchange.

```
GROUP ID: V-256050
RULE ID: SV-256050r882492
```

Rationale:

Packet loss can occur when an IGP adjacency is established and the router begins forwarding packets using the new adjacency before the LDP label exchange completes between the peers on that link. Packet loss can also occur if an LDP session closes and the router continues to forward traffic using the link associated with the LDP peer rather than an alternate pathway with a fully synchronized LDP session. The MPLS LDP-IGP Synchronization feature provides a means to synchronize LDP with OSPF or IS-IS to minimize MPLS packet loss. When an IGP adjacency is established on a link but LDP-IGP synchronization is not yet achieved or is lost, the IGP will advertise the max-metric on that link.

Audit:

Review the Arista router OSPF or IS-IS configuration.

ISIS configuration example:

```
router isis 1
   mpls ldp sync default
```

OSPF configuration example:

```
router ospf 1
   mpls ldp sync default
```

Verify LDP will synchronize with the link-state routing protocol.

```
interface Loopback1
   description MPLS-LDP-Router-ID
   ip address 10.1.129.94/32
   mpls ldp igp sync
   ip ospf area 0.0.0.5
```

If the Arista router is not configured to synchronize IGP and LDP, this is a finding.

Remediation:

Configure the Arista MPLS router to synchronize IGP and LDP, minimizing packet loss when an IGP adjacency is established prior to LDP peers completing label exchange.

ISIS configuration example:

```
P-Router-EOS(config)#router isis 1
P-Router-EOS(config-router-isis)#mpls ldp sync default
```

OSPF configuration example:

```
LEAF-1A(config-mpls-ldp)#router ospf 1
LEAF-1A(config-router-ospf)#mpls ldp sync default
```

Configure LDP to synchronize with the link-state routing protocol.

```
LEAF-1A(config)#interface Loopback1
LEAF-1A(config-if-Lo0)#erface Loopback1
LEAF-1A(config-if-Lo0)#description MPLS-LDP-Router-ID
LEAF-1A(config-if-Lo0)#ip address 10.1.129.94/32
LEAF-1A(config-if-Lo0)#mpls ldp igp sync
LEAF-1A(config-if-Lo0)#ip ospf area 0.0.0.5
```

Additional Information:

CCI-000366 Implement the security configuration settings.

- NIST SP 800-53::CM-6 b
- NIST SP 800-53A::CM-6.1 (iv)
- NIST SP 800-53 Revision 4::CM-6 b
- NIST SP 800-53 Revision 5::CM-6 b

1.65 ARST-RT-000720 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The MPLS router must be configured to have TTL propagation disabled.

```
GROUP ID: V-256051
RULE ID: SV-256051r882495
```

Rationale:

The head end of the label-routered path (LSP), the label edge router (LER) will decrement the IP packet's time-to-live (TTL) value by one and then copy the value to the MPLS TTL field. At each label-routered router (LSR) hop, the MPLS TTL value is decremented by one. The MPLS router that pops the label (either the penultimate LSR or the egress LER) will copy the packet's MPLS TTL value to the IP TTL field and decrement it by one.

This TTL propagation is the default behavior. Because the MPLS TTL is propagated from the IP TTL, a traceroute will list every hop in the path, be it routed or label routered, thereby exposing core nodes. With TTL propagation disabled, LER decrements the IP packet's TTL value by one and then places a value of 255 in the packet's MPLS TTL field, which is then decremented by one as the packet passes through each LSR in the MPLS core. Because the MPLS TTL never drops to zero, none of the LSP hops triggers an ICMP TTL exceeded message and consequently, these hops are not recorded in a traceroute. Hence, nodes within the MPLS core cannot be discovered by an attacker.

Audit:

Review the Arista router configuration to verify TTL propagation is disabled.

Verify the router is configured to disable the TTL propagation.

```
no mpls icmp ttl-exceeded tunneling
```

If the Arista router is not configured to disable TTL propagation, this is a finding.

Remediation:

Configure Arista LERs to disable TTL propagation.

Configure the router to disable the TTL propagation in MPLS core network.

```
router(config)#no mpls icmp ttl-exceeded tunneling
```

Additional Information:

CCI-000366 Implement the security configuration settings.

- NIST SP 800-53::CM-6 b
- NIST SP 800-53A::CM-6.1 (iv)
- NIST SP 800-53 Revision 4::CM-6 b
- NIST SP 800-53 Revision 5::CM-6 b

1.66 ARST-RT-000730 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The PE router must be configured to have each Virtual Routing and Forwarding (VRF) instance bound to the appropriate physical or logical interfaces to maintain traffic separation between all MPLS L3VPNs.

```
GROUP ID: V-256052
RULE ID: SV-256052r991791
```

Rationale:

The primary security model for an MPLS L3VPN infrastructure is traffic separation. The service provider must guarantee the customer that traffic from one VPN does not leak into another VPN or into the core, and that core traffic must not leak into any VPN. Hence, it is imperative that each CE-facing interface can only be associated to one VRF—that alone is the fundamental framework for traffic separation.

Audit:

Review the design plan for deploying L3VPN and VRF-lite.

Review all CE-facing interfaces and verify the proper VRF is defined.

To verify the interfaces toward CE facing with proper VRF defined, execute the command "sh run int ethernet YY".

```
vrf instance PROD
vrf instance DEVP

ip routing vrf PROD
ip routing vrf DEVP

interface Ethernet3
   no routerport
   vrf PROD
   ip address 10.1.99.11/24

interface Ethernet4
   no routerport
   vrf DEVP
   ip address 10.11.5.11/24
```

If any VRFs are not bound to the appropriate physical or logical interface, this is a finding.

Remediation:

Configure the Arista PE router to have each VRF bound to the appropriate physical or logical interfaces to maintain traffic separation between all MPLS L3VPNs.

Configure the VRF on the CE facing interfaces.

```
PE11(config)#vrf instance PROD
PE11(config)#vrf instance DEVP
!
PE11(config)#ip routing vrf PROD
PE11(config)#ip routing vrf DEVP
!
PE11(config)#interface Ethernet3
PE11(config-if-Et3)#no routerport
PE11(config-if-Et3)#vrf PROD
PE11(config-if-Et3)#ip address 10.1.99.11/24
!
PE11(config)#interface Ethernet4
PE11(config-if-Et3)#no routerport
PE11(config-if-Et3)#vrf DEVP
PE11(config-if-Et3)#ip address 10.11.5.11/24
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.67 ARST-RT-000740 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The PE router must be configured to have each Virtual Routing and Forwarding (VRF) instance with the appropriate Route Target (RT).

```
GROUP ID: V-256053
RULE ID: SV-256053r991792
```

Rationale:

The primary security model for an MPLS L3VPN as well as a VRF-lite infrastructure is traffic separation. Each interface can only be associated to one VRF, which is the fundamental framework for traffic separation. Forwarding decisions are made based on the routing table belonging to the VRF. Control of what routes are imported into or exported from a VRF is based on the RT. It is critical that traffic does not leak from one COI tenant or L3VPN to another; hence, it is imperative that the correct RT is configured for each VRF.

Audit:

Verify the correct RT is configured for each VRF.

Review the design plan for MPLS/L3VPN and VRF-lite to determine what RTs have been assigned for each VRF.

Review the route-target import, route-target, or route-target export statements under each configured VRF and verify the correct RTs have been defined for each VRF.

To verify the correct RTs have been defined for each VRF on a PE router, execute the command "sh run sec router bgp".

```
router bgp 65000
   vrf PROD
      rd 200:200
      route-target import vpn-ipv4 200:200
      route-target export vpn-ipv4 200:200
```

Note: Import and export route-maps are normally used when finer granularity is required.

If VRFs are configured with the wrong RT, this is a finding.

Remediation:

Configure all J-PE Arista routers to have the correct VRF defined with the appropriate RT.

Configure the route-target's for import and export.

```
PE11(config)#router bgp 65000
PE11(config-router-bgp)#vrf PROD
PE11(config-router-bgp-vrf-PROD)#rd 200:200
PE11(config-router-bgp-vrf-PROD)#route-target import vpn-ipv4 200:200
PE11(config-router-bgp-vrf-PROD)#route-target export vpn-ipv4 200:200
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.68 ARST-RT-000750 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The PE router must be configured to have each VRF with the appropriate Route Distinguisher (RD).

```
GROUP ID: V-256054
RULE ID: SV-256054r991793
```

Rationale:

An RD provides uniqueness to the customer address spaces within the MPLS L3VPN infrastructure. The concept of the VPN-IPv4 and VPN-IPv6 address families consists of the RD prepended before the IP address. Hence, if the same IP prefix is used in several different L3VPNs, it is possible for BGP to carry several completely different routes for that prefix, one for each VPN.

Since VPN-IPv4 addresses and IPv4 addresses are different address families, BGP never treats them as comparable addresses. The purpose of the RD is to create distinct routes for common IPv4 address prefixes. On any given PE router, a single RD can define a VRF in which the entire address space may be used independently, regardless of the makeup of other VPN address spaces. Hence, it is imperative that a unique RD is assigned to each L3VPN and that the proper RD is configured for each VRF.

Audit:

Review the RDs that have been assigned for each VRF according to the plan provided by the ISSM.

Review all VRFs configured on CE-facing interfaces and verify the proper RD has been configured for each.

To verify the proper Route Distinguisher has been configured, execute the command "sh run sec router bgp".

```
router bgp 65000
   vrf PROD
      rd 200:200
```

If the wrong RD has been configured for any VRF, this is a finding.

Remediation:

Configure the correct RD for each VRF.

Configure the correct Route Distinguisher.

```
PE11(config)#router bgp 65000
PE11(config-router-bgp)#vrf PROD
PE11(config-router-bgp-vrf-PROD)#rd 200:200
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.69 ARST-RT-000760 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The PE router providing MPLS Virtual Private Wire Service (VPWS) must be configured to have the appropriate virtual circuit identification (VC ID) for each attachment circuit.

```
GROUP ID: V-256055
RULE ID: SV-256055r991794
```

Rationale:

VPWS is an L2VPN technology that provides a virtual circuit between two PE routers to forward Layer 2 frames between two customer-edge routers or routers through an MPLS-enabled IP core. The ingress PE router (virtual circuit head-end) encapsulates Ethernet frames inside MPLS packets using label stacking and forwards them across the MPLS network to the egress PE router (virtual circuit tail-end). During a virtual circuit setup, the PE routers exchange VC label bindings for the specified VC ID. The VC ID specifies a pseudowire associated with an ingress and egress PE router and the customer-facing attachment circuits.

To guarantee that all frames are forwarded onto the correct pseudowire and to the correct customer and attachment circuits, it is imperative that the correct VC ID is configured for each attachment circuit.

Audit:

Review the ingress and egress PE router configuration for each virtual circuit that has been provisioned.

Verify the correct and unique VCID has been configured for the appropriate attachment circuit.

```
Run the command sh run | section patch
patch panel
   patch port
      connector 1 interface Ethernet2
      connector 2 pseudowire bgp vpws evi-1 pseudowire pw1
   patch subintf
      connector 1 interface Ethernet3.1
      connector 2 pseudowire bgp vpws evi-1 pseudowire pw2
```

Run the command

```
sh run | section router bgp
router bgp 65000
   neighbor 10.0.0.1 remote-as 1
   neighbor 10.0.0.1 send-community extended
   neighbor 10.0.0.1 maximum-routes 12000
   !
   vpws evi-1
      rd 10.2.2.2:2
      route-target import export evpn 0.0.0.0:1
      mpls control-word
      !
      pseudowire pw1
         evpn vpws id local 2001 remote 1001
      !
      pseudowire pw2
         evpn vpws id local 2002 remote 1002
   !
   address-family evpn
      neighbor default encapsulation mpls next-hop-self source-interface
Loopback0
      neighbor 10.0.0.1 activate
```

If the correct VC ID has not been configured on both routers, this is a finding.

Remediation:

Assign globally unique VC IDs for each virtual circuit and configure the attachment circuits with the appropriate VC ID.

Configure the same VC ID on both ends of the VC.

```
patch panel
   patch port
      connector 1 interface Ethernet2
      connector 2 pseudowire bgp vpws evi-1 pseudowire pw1
   patch subintf
      connector 1 interface Ethernet3.1
      connector 2 pseudowire bgp vpws evi-1 pseudowire pw2

router bgp 65000
   neighbor 10.0.0.1 remote-as 1
   neighbor 10.0.0.1 send-community extended
   neighbor 10.0.0.1 maximum-routes 12000
   !
   vpws evi-1
      rd 10.2.2.2:2
      route-target import export evpn 0.0.0.0:1
      mpls control-word
      !
      pseudowire pw1
         evpn vpws id local 2001 remote 1001
      !
      pseudowire pw2
         evpn vpws id local 2002 remote 1002
   !
   address-family evpn
      neighbor default encapsulation mpls next-hop-self source-interface
Loopback0
      neighbor 10.0.0.1 activate
```

VLAN mode example:

```
interface Ethernet3
   no routerport
!
interface Ethernet3.1
   encapsulation dot1q vlan 1

Flexible Encapsulation example:

interface Ethernet3
   no routerport
!
interface Ethernet3.1
   encapsulation vlan
      client dot1q 11 network client
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.70 ARST-RT-000770 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista Multicast Source Discovery Protocol (MSDP) router must be configured to use its loopback address as the source address when originating MSDP traffic.

```
GROUP ID: V-256056
RULE ID: SV-256056r991795
```

Rationale:

Using a loopback address as the source address offers a multitude of uses for security, access, management, and scalability of MSDP routers. It is easier to construct appropriate ingress filters for router management plane traffic destined to the network management subnet since the source addresses will be from the range used for loopback interfaces instead of a larger range of addresses used for physical interfaces. Log information recorded by authentication and syslog servers will record the router's loopback address instead of the numerous physical interface addresses.

Audit:

Review the Arista router configuration to verify a loopback address has been configured.

Verify a loopback interface is used as the source address for all MSDP packets generated by the router. Execute the command "sh run sec router msdp".

```
router#show running-config | section router msdp
router msdp
   originator-id local-interface Loopback0
```

If the Arista router does not use its loopback address as the source address when originating MSDP traffic, this is a finding.

Remediation:

Ensure the Arista router originator-id is the source address loopback0 for originating traffic.

```
router(config)#router msdp
router (config-router-msdp)#originator-id local-interface loopback0
router (config-router-msdp)#exit
```

Additional Information:

CCI-004931 Establish organization-defined alternate communications paths for system operations organizational command and control.

- NIST SP 800-53 Revision 5::SC-47

1.71 ARST-RT-000780 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Arista router must be configured to advertise a hop limit of at least 32 in Router Advertisement messages for IPv6 stateless auto-configuration deployments.

```
GROUP ID: V-256057
RULE ID: SV-256057r882513
```

Rationale:

The Neighbor Discovery protocol allows a hop limit value to be advertised by routers in a Router Advertisement message being used by hosts instead of the standardized default value. If a very small value was configured and advertised to hosts on the LAN segment, communications would fail due to the hop limit reaching zero before the packets sent by a host reached its destination.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to determine if the hop limit has been configured for Router Advertisement messages. Execute the command "sh run | section hop-limit".

```
interface Ethernet3
  ipv6 nd ra hop-limit 32
```

If the router has been configured and has not been set to at least 32, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to advertise a hop limit of at least 32 in Router Advertisement messages.

```
LEAF-1A(config-if-Et3)#interface ethernet 3
LEAF-1A(config-if-Et3)#ipv6 nd ra hop-limit 32
```

Additional Information:

CCI-000366 Implement the security configuration settings.

- NIST SP 800-53::CM-6 b
- NIST SP 800-53A::CM-6.1 (iv)
- NIST SP 800-53 Revision 4::CM-6 b
- NIST SP 800-53 Revision 5::CM-6 b

1.72 ARST-RT-000790 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista router must not be configured to use IPv6 Site Local Unicast addresses.

```
GROUP ID: V-256058
RULE ID: SV-256058r882516
```

Rationale:

As currently defined, site local addresses are ambiguous and can be present in multiple sites. The address itself does not contain any indication of the site to which it belongs. The use of site-local addresses has the potential to adversely affect network security through leaks, ambiguity, and potential misrouting as documented in section 2 of RFC3879. RFC3879 formally deprecates the IPv6 site-local unicast prefix FEC0::/10 as defined in RFC3513.

Audit:

Review the Arista router configuration to ensure FEC0::/10 IP addresses are not defined.

Step 1: Verify that FECO::/10 IPv6 addresses are not configured.

```
interface ethernet 3
 no routerport
 ipv6 address FD6D:8D64:AF0C:2::/64
```

If IPv6 Site Local Unicast addresses are defined, this is a finding.

Remediation:

Configure the Arista router using authorized IPv6 addresses.

Step 1: Configure the interface with IPv6 address.

```
LEAF-1A(config-if-Et3)#interface ethernet 3
LEAF-1A(config-if-Et3)#no routerport
LEAF-1A(config-if-Et3)#ipv6 address FD6D:8D64:AF0C:2::/64
```

Additional Information:

CCI-000366 Implement the security configuration settings.

- NIST SP 800-53::CM-6 b
- NIST SP 800-53A::CM-6.1 (iv)
- NIST SP 800-53 Revision 4::CM-6 b
- NIST SP 800-53 Revision 5::CM-6 b

1.73 ARST-RT-000800 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Arista perimeter router must be configured to suppress Router Advertisements on all external IPv6-enabled interfaces.

```
GROUP ID: V-256059
RULE ID: SV-256059r882519
```

Rationale:

Many of the known attacks in stateless autoconfiguration are defined in RFC 3756 were present in IPv4 ARP attacks. To mitigate these vulnerabilities, links that have no hosts connected such as the interface connecting to external gateways must be configured to suppress router advertisements.

Audit:

This requirement is not applicable for the DODIN backbone.

Review the Arista router configuration to verify Router Advertisements are suppressed on all external IPv6-enabled interfaces.

<Example configuration for VLAN 200>

```
interface vlan 200
 ipv6 nd ra disabled all
```

If the Arista router is not configured to suppress Router Advertisements on all external IPv6-enabled interfaces, this is a finding.

Remediation:

This requirement is not applicable for the DODIN backbone.

Configure the Arista router to suppress Router Advertisements on all external IPv6-enabled interfaces.

Configure the Arista router to suppress RAs on all IPv6 enabled interface as in the following example for VLAN 200:

```
router(config)#interface vlan 200
router(config-vl200)#ipv6 nd ra disabled all
router(config-vl200)#
```

Additional Information:

CCI-000366 Implement the security configuration settings.

- NIST SP 800-53::CM-6 b
- NIST SP 800-53A::CM-6.1 (iv)
- NIST SP 800-53 Revision 4::CM-6 b
- NIST SP 800-53 Revision 5::CM-6 b

1.74 ARST-RT-000830 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The perimeter router must be configured to block all packets with any IP options.

```
GROUP ID: V-256060
RULE ID: SV-256060r945859
```

Rationale:

Packets with IP options are not fast routered and henceforth must be punted to the router processor. Hackers who initiate denial-of-service (DoS) attacks on routers commonly send large streams of packets with IP options. Dropping the packets with IP options reduces the load of IP options packets on the router. The end result is a reduction in the effects of the DoS attack on the router and on downstream routers.

Audit:

Verify the perimeter router is configured to block all packets with any IP options with the following command:

```
router#show run | section IP_Option_ACL
IP Access List IP_Option_ACL
        10 deny ip any any ip-length gt 5
        20 deny any log
!
interface Ethernet25
   description STIG_IP_Option_ACL
   ip access-group IP_Option_ACL in
!
```

If the perimeter router is not configured to block packets with IP options, this is a finding.

Remediation:

Configure the perimeter router to block packets with IP options with the following commands:

```
router#config
router(config)# ip access-list IP_Option_ACL
   10 deny ip any any ip-length gt 5
!
router(config)#interface Ethernet25
 ip access-group IP_Option_ACL in
!
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

1.75 ARST-RT-000840 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The PE router must be configured to ignore or block all packets with any IP options.

```
GROUP ID: V-256061
RULE ID: SV-256061r945860
```

Rationale:

Packets with IP options are not fast routered and therefore must be punted to the router processor. Hackers who initiate denial-of-service (DoS) attacks on routers commonly send large streams of packets with IP options. Dropping the packets with IP options reduces the load of IP options packets on the router. The end result is a reduction in the effects of the DoS attack on the router and on downstream routers.

Audit:

Verify the PE router is configured to block all packets with any IP options with the following command:

```
router#show run | section IP_Option_ACL
IP Access List IP_Option_ACL
        10 deny ip any any ip-length gt 5
        20 deny any log
!
interface Ethernet25
   description STIG_IP_Option_ACL
   ip access-group IP_Option_ACL in
!
```

If the perimeter router is not configured to block packets with IP options, this is a finding.

Remediation:

Configure the PE router to block packets with IP options with the following commands:

```
router#config
router(config)# ip access-list IP_Option_ACL
   10 deny ip any any ip-length gt 5
!
router(config)#interface Ethernet25
 ip access-group IP_Option_ACL in
!
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

Appendix: Summary Table

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1	**STIG RULES**		
1.1	ARST-RT-000010 (Manual)	☐	☐
1.2	ARST-RT-000020 (Manual)	☐	☐
1.3	ARST-RT-000030 (Manual)	☐	☐
1.4	ARST-RT-000040 (Manual)	☐	☐
1.5	ARST-RT-000050 (Manual)	☐	☐
1.6	ARST-RT-000060 (Manual)	☐	☐
1.7	ARST-RT-000070 (Manual)	☐	☐
1.8	ARST-RT-000080 (Manual)	☐	☐
1.9	ARST-RT-000090 (Manual)	☐	☐
1.10	ARST-RT-000100 (Manual)	☐	☐
1.11	ARST-RT-000110 (Manual)	☐	☐
1.12	ARST-RT-000120 (Manual)	☐	☐
1.13	ARST-RT-000130 (Manual)	☐	☐
1.14	ARST-RT-000140 (Manual)	☐	☐
1.15	ARST-RT-000150 (Manual)	☐	☐
1.16	ARST-RT-000160 (Manual)	☐	☐
1.17	ARST-RT-000170 (Manual)	☐	☐
1.18	ARST-RT-000180 (Manual)	☐	☐
1.19	ARST-RT-000190 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.20	ARST-RT-000200 (Manual)	☐	☐
1.21	ARST-RT-000210 (Manual)	☐	☐
1.22	ARST-RT-000230 (Manual)	☐	☐
1.23	ARST-RT-000260 (Manual)	☐	☐
1.24	ARST-RT-000280 (Manual)	☐	☐
1.25	ARST-RT-000290 (Manual)	☐	☐
1.26	ARST-RT-000300 (Manual)	☐	☐
1.27	ARST-RT-000320 (Manual)	☐	☐
1.28	ARST-RT-000310 (Manual)	☐	☐
1.29	ARST-RT-000330 (Manual)	☐	☐
1.30	ARST-RT-000340 (Manual)	☐	☐
1.31	ARST-RT-000350 (Manual)	☐	☐
1.32	ARST-RT-000380 (Manual)	☐	☐
1.33	ARST-RT-000390 (Manual)	☐	☐
1.34	ARST-RT-000370 (Manual)	☐	☐
1.35	ARST-RT-000400 (Manual)	☐	☐
1.36	ARST-RT-000410 (Manual)	☐	☐
1.37	ARST-RT-000420 (Manual)	☐	☐
1.38	ARST-RT-000430 (Manual)	☐	☐
1.39	ARST-RT-000440 (Manual)	☐	☐
1.40	ARST-RT-000450 (Manual)	☐	☐
1.41	ARST-RT-000470 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.42	ARST-RT-000480 (Manual)	☐	☐
1.43	ARST-RT-000490 (Manual)	☐	☐
1.44	ARST-RT-000510 (Manual)	☐	☐
1.45	ARST-RT-000520 (Manual)	☐	☐
1.46	ARST-RT-000530 (Manual)	☐	☐
1.47	ARST-RT-000540 (Manual)	☐	☐
1.48	ARST-RT-000550 (Manual)	☐	☐
1.49	ARST-RT-000570 (Manual)	☐	☐
1.50	ARST-RT-000560 (Manual)	☐	☐
1.51	ARST-RT-000580 (Manual)	☐	☐
1.52	ARST-RT-000590 (Manual)	☐	☐
1.53	ARST-RT-000600 (Manual)	☐	☐
1.54	ARST-RT-000610 (Manual)	☐	☐
1.55	ARST-RT-000620 (Manual)	☐	☐
1.56	ARST-RT-000630 (Manual)	☐	☐
1.57	ARST-RT-000640 (Manual)	☐	☐
1.58	ARST-RT-000650 (Manual)	☐	☐
1.59	ARST-RT-000660 (Manual)	☐	☐
1.60	ARST-RT-000680 (Manual)	☐	☐
1.61	ARST-RT-000670 (Manual)	☐	☐
1.62	ARST-RT-000690 (Manual)	☐	☐
1.63	ARST-RT-000700 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.64	ARST-RT-000710 (Manual)	☐	☐
1.65	ARST-RT-000720 (Manual)	☐	☐
1.66	ARST-RT-000730 (Manual)	☐	☐
1.67	ARST-RT-000740 (Manual)	☐	☐
1.68	ARST-RT-000750 (Manual)	☐	☐
1.69	ARST-RT-000760 (Manual)	☐	☐
1.70	ARST-RT-000770 (Manual)	☐	☐
1.71	ARST-RT-000780 (Manual)	☐	☐
1.72	ARST-RT-000790 (Manual)	☐	☐
1.73	ARST-RT-000800 (Manual)	☐	☐
1.74	ARST-RT-000830 (Manual)	☐	☐
1.75	ARST-RT-000840 (Manual)	☐	☐

Appendix: Change History

Date	Version	Changes for this version
08/27/2025	1.0.0	Initial CIS Release