CIS Cisco IOS XR Router RTR STIG Benchmark

v1.0.0 - 09-03-2025

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

Cisco IOS XR Router RTR Secure Technical Implementation Guide (STIG)

Version: 3 Release: 2 Benchmark

Date: 24 Oct 2024

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 Cisco IOS XR Router RTR 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 Cisco IOS XR Router RTR

Recommendations

1 STIG RULES

Cisco IOS XR router

Cisco IOS XR Router RTR Secure Technical Implementation Guide (STIG)

Version: 3 Release: 2 Benchmark

Date: 24 Oct 2024

CLASSIFICATION unclassified

1.1 CISC-RT-000010 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216735
RULE ID: SV-216735r1007831
```

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.

If Virtual Routing and Forwarding (VRF) is used to segment traffic and force traffic to traverse through next generation firewalls with ACLs, this requirement is not applicable.

Review the router configuration to verify that ACLs are configured to allow or deny traffic for specific source and destination addresses as well as ports and protocols. For example, the configuration below will allow only printer traffic into subnet 10.1.23.0/24 and SQL traffic into subnet 10.1.24.0/24. ICMP is allowed for troubleshooting and OSPF is the routing protocol used within the network.

interface GigabitEthernet0/0/0/0

description link to core

ipv4 address 10.1.12.2 255.255.255.0

ipv4 access-group FILTER_SERVER_TRAFFIC ingress

…

…

…

ipv4 access-list FILTER_SERVER_TRAFFIC

10 permit tcp any 10.1.23.0 0.0.0.255 eq lpd

20 permit tcp any 10.1.23.0 0.0.0.255 eq 631

30 permit tcp any 10.1.23.0 0.0.0.255 eq 9100

40 permit tcp any 10.1.24.0 0.0.0.255 eq 1433

50 permit tcp any 10.1.24.0 0.0.0.255 eq 1434

60 permit tcp any 10.1.24.0 0.0.0.255 eq 4022

70 permit icmp any any

80 permit ospf any any

90 deny ipv4 any any

If the 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 ACLs to allow or deny traffic for specific source and destination addresses as well as ports and protocols between various subnets as required. The commands used below were used to create the configuration as shown in the check content.

RP/0/0/CPU0:R2(config)#ipv4 access-list FILTER_SERVER_TRAFFIC

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.23.0 0.0.0.255 eq lpd

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.23.0 0.0.0.255 eq 631

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.23.0 0.0.0.255 eq 9100

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.24.0 0.0.0.255 eq 1433

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.24.0 0.0.0.255 eq 1434

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any 10.1.24.0 0.0.0.255 eq 4022

RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp any any

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ospf any any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any

RP/0/0/CPU0:R2(config-ipv4-acl)#

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#int g0/0/0/0

RP/0/0/CPU0:R2(config-if)#ipv4 access-group FILTER_SERVER_TRAFFIC

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 CISC-RT-000050 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to enable routing protocol authentication using FIPS 198-1 algorithms with keys not exceeding 180 days of lifetime.

```
GROUP ID: V-216739
RULE ID: SV-216739r1007834
```

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 using FIPS 198-1 algorithms for routing updates. If the keys used for authentication are guessed, the malicious user could create havoc within the network by advertising incorrect routes and redirecting traffic. Some routing protocols allow the use of key chains for authentication. A key chain is a set of keys that is used in succession, with each having a lifetime of no more than 180 days. Changing the keys frequently reduces the risk of them eventually being guessed. If a time period occurs during which no key is activated, neighbor authentication cannot occur, and therefore routing updates will fail.

Audit:

Review the router configuration using the configuration examples below for BGP and OSPF.

Certain older protocols supporting only MD5 will incur a permanent finding for those protocols as MD5 is not FIPS compliant.

Note: The 180-day key lifetime is Not Applicable for the DODIN Backbone. The remainder of the requirement still applies.

Verify that neighbor router authentication is enabled for all routing protocols. If neighbor authentication is not enabled this is a finding.

Verify that authentication is configured to use FIPS 198-1 message authentication algorithms. If the routing protocol authentication is not configured to use FIPS 198-1 algorithms this is a finding.

Verify that the protocol key lifetime is configured to not exceed 180 days. If any protocol key lifetime is configured to exceed 180 days this is a finding.

BGP Example:

key chain <KEY-CHAIN-NAME>

key <KEY-ID>

send-id <ID>

recv-id <ID>

cryptographic-algorithm hmac-sha256

key-string <KEY>

accept-lifetime 00:00:00 Jan 1 2022 duration 180

send-lifetime 00:00:00 Jan 1 2022 duration 180

!

tcp ao

keychain BGP_KEY_CHAIN

key <KEY-ID> SendID <ID> ReceiveID <ID>

!

!

router bgp <ASN>

neighbor X.X.X.X

remote-as <ASN>

ao BGP_KEY_CHAIN

address-family ipv4 unicast

!

Note: TCP-AO is used to replace MD5 in BGP authentication.

OSPF Example:

key chain OSPF_KEY_CHAIN

key 1

key-string xxxxxxx

send-lifetime 00:00:00 Jan 1 2018 23:59:59 Mar 31 2018

accept-lifetime 00:00:00 Jan 1 2018 01:05:00 Apr 1 2018

cryptographic-algorithm hmac-sha-256

key 2

key-string yyyyyyy

send-lifetime 00:00:00 Apr 1 2018 23:59:59 Jun 30 2018

accept-lifetime 23:55:00 Mar 31 2018 01:05:00 Jul 1 2018

cryptographic-algorithm hmac-sha-256

…

…

…

router ospf 1

area 0

authentication message-digest keychain OSPF_KEY_CHAIN

Remediation:

Configure routing protocol authentication to use a NIST-validated FIPS 198-1 message authentication code algorithm with keys not exceeding 180 days of lifetime as shown in the examples.

BGP Example:

Step 1: Configure a keychain using a FIPS 198-1 algorithm with a key duration not exceeding 180 days.

key chain <KEY-CHAIN-NAME>

key <KEY-ID>

send-id <ID>

recv-id <ID>

cryptographic-algorithm hmac-sha256

key-string <KEY>

accept-lifetime 00:00:00 Jan 1 2022 duration 180

send-lifetime 00:00:00 Jan 1 2022 duration 180

Step 2: Configure BGP autonomous system to use the keychain for authentication.

tcp ao

keychain BGP_KEY_CHAIN

key <KEY-ID> SendID <ID> ReceiveID <ID>

!

!

router bgp <ASN>

neighbor X.X.X.X

remote-as <ASN>

ao BGP_KEY_CHAIN

address-family ipv4 unicast

OSPF Example:

Step 1: Configure a keychain using a FIPS 198-1 algorithm with a key duration not exceeding 180 days.

key chain OSPF_KEY_CHAIN

key 1

key-string xxxxxxx

send-lifetime 00:00:00 Jan 1 2018 23:59:59 Mar 31 2018

accept-lifetime 00:00:00 Jan 1 2018 01:05:00 Apr 1 2018

cryptographic-algorithm hmac-sha-256

key 2

key-string yyyyyyy

send-lifetime 00:00:00 Apr 1 2018 23:59:59 Jun 30 2018

accept-lifetime 23:55:00 Mar 31 2018 01:05:00 Jul 1 2018

cryptographic-algorithm hmac-sha-256

Step 2: Configure OSPF to use the keychain for authentication.

router ospf 1

area 0

authentication message-digest keychain OSPF_KEY_CHAIN

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)

1.3 CISC-RT-000060 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216740
RULE ID: SV-216740r531087
```

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 router configuration and verify that inactive interfaces have been disabled as shown below.

interface GigabitEthernet0/0/1/0

shutdown

!

interface GigabitEthernet0/0/1/1

shutdown

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

Remediation:

Disable all inactive interfaces as shown below.

R4(config)#interface GigabitEthernet3

R4(config-if)#shutdown

R4(config)#interface GigabitEthernet4

R4(config-if)#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

1.4 CISC-RT-000070 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216741
RULE ID: SV-216741r531087
```

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 router configuration to verify that the router does not have any unnecessary or non-secure services enabled. For example, the following commands should not be in the configuration:

service ipv4 tcp-small-servers max-servers 10

service ipv4 udp-small-servers max-servers 10

http client vrf xxxxx

telnet vrf default ipv4 server max-servers 1

If any unnecessary services are enabled, this is a finding.

Remediation:

Disable the following services if enabled as shown in the example below.

RP/0/0/CPU0:R3(config)#no service ipv4 tcp-small-servers

RP/0/0/CPU0:R3(config)#no service ipv4 udp-small-servers

RP/0/0/CPU0:R3(config)#no http client vrf xxxxx

RP/0/0/CPU0:R3(config)#no telnet ipv4 server

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

1.5 CISC-RT-000130 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

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

```
GROUP ID: V-216743
RULE ID: SV-216743r531087
```

Rationale:

The route processor handles traffic destined to the router—the key component used to build forwarding paths and is instrumental with all network management functions. Hence, any disruption or DoS attack to the route processor can result in mission critical network outages.

Audit:

Review the external and internal ACLs to verify that the router is configured to only allow specific management and control plane traffic from specific sources destined to itself.

Step 1: Verify ACLs has been configured as shown in the example below that matches expected control plane and management plane traffic. With the exception of ICMP, all other traffic destined to the router should be dropped.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log

!

ipv4 access-list INTERNAL_ACL_INBOUND

10 permit icmp any any

20 permit ospf host 10.1.12.1 host 10.1.12.2

30 permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq ssh

40 permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq tacacs

50 permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq snmp

60 permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq ntp

70 deny ipv4 any host 10.1.12.2 log

80 permit ipv4 any any

Note: For the internal ACL example, all routers within the hypothetical network (10.1.0.0/16) have been configured to use the loopback address to source all management traffic (not shown); hence, the loopbacks are the only allowable destination address for management traffic. In addition, all management traffic destined to the router must originate from the management network (10.2.1.0/24). With the exception of link-local control plane traffic and ICMP, all other traffic destined to any physical interface address will be dropped.

Step 2: Verify that the ACL has been applied to the appropriate interface as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

!

interface GigabitEthernet0/0/0/2

ipv4 address 10.1.12.2 255.255.255.0

ipv4 access-group INTERNAL_ACL_INBOUND ingress

If the router is not configured to restrict traffic destined to itself, this is a finding.

Remediation:

Configure the ACL for any external interfaces as shown in the example.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any host x.11.1.1 log

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any established

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

Configure the ACL for any external interfaces as shown in the example.

RP/0/0/CPU0:R3(config)#ipv4 access-list INTERNAL_ACL_INBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp any any

RP/0/0/CPU0:R3(config-ipv4-acl)#permit ospf host 10.1.12.1 host 10.1.12.2

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2eq 22

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2eq tacacs

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq snmp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq ntp

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any host 10.1.12.2 log

RP/0/0/CPU0:R3(config-ipv4-acl)#permit ip any any

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

Note: best practice is to configure the ACL statements relative to traffic destined to the router first followed by ACL statements for transit traffic.

Step 2: Apply the ACLs to the appropriate interface as shown in the example below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#exit

RP/0/0/CPU0:R3(config)#int g0/0/0/2

RP/0/0/CPU0:R3(config-if)#ipv4 access-group INTERNAL_ACL_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

1.6 CISC-RT-000140 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216744
RULE ID: SV-216744r531087
```

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 external and internal ACLs to verify that the router is configured to drop all fragmented ICMP packets destined to itself.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

25 deny icmp any host x.11.1.2 fragments log

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log

!

ipv4 access-list INTERNAL_ACL_INBOUND

5 deny icmp any any fragments

10 permit icmp any host 10.1.12.2 fragments

20 permit ospf host 10.1.12.1 host 10.1.12.2

30 permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq ssh

40 permit tcp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq tacacs

50 permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq snmp

60 permit udp 10.2.1.0 0.0.0.255 host 10.1.12.2 eq ntp

70 deny ipv4 any host 10.1.12.2 log

…

…

…

110 permit ip any any

Note: Ensure the statement to deny ICMP fragments is before any permit statements for ICMP.

If the router is not configured to drop all fragmented ICMP packets destined to itself, this is a finding.

Remediation:

Configure the external and internal ACLs to drop all fragmented ICMP packets destined to itself as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#25 deny icmp any host x.11.1.2 fragments log

RP/0/0/CPU0:R3(config)#ipv4 access-list INTERNAL_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#5 deny icmp any host 10.1.12.2 fragments log

Note: Ensure the above statement is before any permit statements for ICMP.

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

1.7 CISC-RT-000160 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216745
RULE ID: SV-216745r856436
```

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 all nodes broadcast address).

Audit:

Review the router configuration to determine if it is compliant with this requirement. IP directed broadcast command must not be found on any interface as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 directed-broadcast

If IP directed broadcast is not disabled on all interfaces, this is a finding.

Remediation:

Disable IP directed broadcast on all interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#no ipv4 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

1.8 CISC-RT-000170 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to have Internet Control Message Protocol (ICMP) unreachable messages disabled on all external interfaces.

```
GROUP ID: V-216746
RULE ID: SV-216746r856437
```

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 configuration to verify the ipv4 unreachables disable command has been configured on all external interfaces as shown in the configuration example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 unreachables disable

Note: ICMP unreachables must also be configured. On the Null0 interface if it is used to black hole traffic.

If ICMP unreachable notifications are sent from any external or Null0 interface, this is a finding.

Alternative – DODIN Backbone

Verify that the PE router is configured to rate limit ICMP unreachable messages as shown in the example below.

ip icmp ipv4 rate-limit unreachable 60000

ip icmp ipv4 rate-limit unreachable DF 1000

Note: In the example above, packet-too-big message (ICMP Type 3 Code 4) can be sent once every second, while all other destination unreachable messages can be sent once every minute. This will avoid disrupting Path MTU Discovery for traffic traversing the backbone while mitigating the risk of an ICMP unreachable denial of service attack.

IF the PE router is not configured to rate limit ICMP unreachable messages, this is a finding.

Remediation:

Disable ip unreachables on all external interfaces as shown below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 unreachables disable

Alternative – DODIN Backbone

Configure the PE router to rate limit ICMP unreachable messages as shown in the example below.

RP/0/0/CPU0:R3(config)#icmp ipv4 rate-limit unreachable df 1000

RP/0/0/CPU0:R3(config)#icmp ipv4 rate-limit unreachable 60000

RP/0/0/CPU0:R3(config)#end

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

1.9 CISC-RT-000180 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to have Internet Control Message Protocol (ICMP) mask reply messages disabled on all external interfaces.

```
GROUP ID: V-216747
RULE ID: SV-216747r856438
```

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 router configuration and verify that ipv4 mask-reply command is not enabled on any external interfaces as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 mask-reply

If the router configuration has the ipv4 mask-reply command is enabled on any external interfaces, this is a finding.

Remediation:

Disable ipv4 mask-reply on all external interfaces as shown below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#no ipv4 mask-reply

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

1.10 CISC-RT-000190 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to have Internet Control Message Protocol (ICMP) redirect messages disabled on all external interfaces.

```
GROUP ID: V-216748
RULE ID: SV-216748r856439
```

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 router configuration to verify that ipv4 redirects command has not been configured on any external interface as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 redirects

If ICMP Redirect messages are enabled on any external interfaces, this is a finding.

Remediation:

Disable ICMP redirects on all external interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#no ipv4 redirects

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

1.11 CISC-RT-000200 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco router must be configured to log all packets that have been dropped at interfaces via ACL.

```
GROUP ID: V-216749
RULE ID: SV-216749r531087
```

Rationale:

Auditing and logging are key components of any security architecture. It is essential for security personnel to know what is being done or attempted to be done, and by whom, to compile an accurate risk assessment. Auditing the actions on network devices provides a means to recreate an attack or identify a configuration mistake on the device.

Audit:

Review all ACLs used to filter traffic and verify that packets being dropped are logged as shown in the configuration below.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

25 deny icmp any host x.11.1.2 fragments log

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log

If packets being dropped at interfaces are not logged, this is a finding.

Remediation:

Configure ACLs to log packets that are dropped as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log

Additional Information:

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

1.12 CISC-RT-000210 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216750
RULE ID: SV-216750r531087
```

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.

Audit:

Review the router configuration to verify that events are logged containing information to establish where the events occurred as shown in the example below.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

25 deny icmp any host x.11.1.2 fragments log-input

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log-input

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log-input

Note: When the log-input parameter is configured on deny statements, the log record will contain the interface where ingress packet has been dropped.

If the router is not configured to produce audit records containing information to establish to establish where the events occurred, this is a finding.

Remediation:

Configure the router to log events containing information to establish where the events occurred as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

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

1.13 CISC-RT-000220 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to produce audit records containing information to establish the source of the events.

```
GROUP ID: V-216751
RULE ID: SV-216751r531087
```

Rationale:

Without establishing the source of the event, 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, security personnel need to know the source of the event.

In addition to logging where events occur within the network, the audit records must also identify sources of events such as IP addresses, processes, and node or device names.

Audit:

Review the router configuration to verify that events are logged containing information to establish the source of the events as shown in the example below.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

25 deny icmp any host x.11.1.2 fragments log

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log-input

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log-input

Note: When the log-input parameter is configured on deny statements, the log record will contain the layer 2 address of the forwarding device for any packet being dropped.

If the router is not configured to produce audit records containing information to establish the source of the events, this is a finding.

Remediation:

Configure the router to log events containing information to establish where the events occurred as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

Additional Information:

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

1.14 CISC-RT-000235 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must be configured to have Cisco Express Forwarding enabled.

```
GROUP ID: V-230037
RULE ID: SV-230037r878128
```

Rationale:

The Cisco Express Forwarding (CEF) switching mode replaces the traditional Cisco routing cache with a data structure that mirrors the entire system routing table. Because there is no need to build cache entries when traffic starts arriving for new destinations, CEF behaves more predictably when presented with large volumes of traffic addressed to many destinations—such as a SYN flood attacks that. Because many SYN flood attacks use randomized source addresses to which the hosts under attack will reply to, there can be a substantial amount of traffic for a large number of destinations that the router will have to handle. Consequently, routers configured for CEF will perform better under SYN floods directed at hosts inside the network than routers using the traditional cache.

Audit:

Review the router to verify that CEF is enabled.

IPv4 Example: ip cef

IPv6 Example: ipv6 cef

If CEF is not enabled, this is a finding.

Remediation:

Enable CEF

IPv4 Example: ip cef

IPv6 Example: ipv6 cef

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

1.15 CISC-RT-000236 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-230040
RULE ID: SV-230040r532994
```

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:

Review the router configuration to determine if the hop limit has been configured for Router Advertisement messages as shown in the example.

ipv6 hop-limit 128

If it has been configured and has not been set to at least 32, it is a finding.

Remediation:

Configure the router to advertise a hop limit of at least 32 in Router Advertisement messages.

R1(config)#ipv6 hop-limit 128

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

1.16 CISC-RT-000237 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router must not be configured to use IPv6 Site Local Unicast addresses.

```
GROUP ID: V-230043
RULE ID: SV-230043r533001
```

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 router configuration to ensure FEC0::/10 IPv6 addresses are not defined.

If IPv6 Site Local Unicast addresses are defined, this is a finding.

Remediation:

Configure the router using only authorized IPv6 addresses.

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

1.17 CISC-RT-000240 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

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

```
GROUP ID: V-216752
RULE ID: SV-216752r531087
```

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 router configuration to verify that the inbound ACL applied to all external interfaces is configured to allow specific ports and protocols and deny all other traffic.

Step 1: Verify that an inbound ACL is applied to all external interfaces as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

Step 2: Review inbound ACL to verify that it is configured to deny all other traffic that is not explicitly allowed.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log-input

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log-input

If the ACL is not configured to allow specific ports and protocols and deny all other traffic, this is a finding. If the ACL is not configured inbound on all external interfaces, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Step 1: Configure an inbound ACL to deny all other traffic by default as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any host x.11.1.1 log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any established

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

Step 2: Apply the ingress filter to all external interfaces

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#end

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)

1.18 CISC-RT-000250 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216753
RULE ID: SV-216753r531087
```

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:

Review the router configuration to verify that ACLs are configured to allow or deny traffic for specific source and destination addresses as well as ports and protocols. In the example below, the router is peering BGP with DISN. ICMP echo and echo-reply packets are allowed for troubleshooting connectivity. WWW traffic is permitted inbound to the NIPRNet host-facing web server (x.12.1.22).

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log

60 permit tcp any host x.12.1.22 eq www

70 permit tcp any any established

80 deny ipv4 any any log-input

…

…

…

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

If the router is not configured to enforce approved authorizations for controlling the flow of information between interconnected networks, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Step 1: Configure an ACL to allow or deny traffic as shown in the example below.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any host x.11.1.1 log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any host x.12.1.22 eq www

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any established

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

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

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#end

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 CISC-RT-000260 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to only allow incoming communications from authorized sources to be routed to authorized destinations.

```
GROUP ID: V-216754
RULE ID: SV-216754r856440
```

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 router configuration to determine if the router allows only incoming communications from authorized sources to be routed to authorized destinations. The hypothetical example below allows inbound NTP from server x.1.12.9 only to host x.12.1.21.

ipv4 access-list EXTERNAL_ACL_INBOUND

…

…

…

60 permit udp host x.12.1.9 host x.12.1.21 eq ntp

70 permit tcp any any established

80 deny ipv4 any any log-input

If the 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 router to allow only incoming communications from authorized sources to be routed to authorized destinations.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

…

…

…

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp host x.12.1.9 host x.12.1.21 eq ntp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any established

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

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)

1.20 CISC-RT-000270 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to block inbound packets with source Bogon IP address prefixes.

```
GROUP ID: V-216755
RULE ID: SV-216755r856441
```

Rationale:

Packets with Bogon IP source addresses should never be allowed to traverse the IP core. Bogon IP networks are RFC1918 addresses or address blocks that have never been assigned by the IANA or have been reserved.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to verify that an ingress ACL applied to all external interfaces is blocking packets with Bogon source addresses.

Step 1: Verify an ACL has been configured containing the current Bogon prefixes as shown in the example below.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 deny ipv4 0.0.0.0 0.255.255.255 any log-input

20 deny ipv4 10.0.0.0 0.255.255.255 any log-input

30 deny ipv4 100.64.0.0 0.63.255.255 any log-input

40 deny ipv4 127.0.0.0 0.255.255.255 any log-input

50 deny ipv4 169.254.0.0 0.0.255.255 any log-input

60 deny ipv4 172.16.0.0 0.15.255.255 any log-input

70 deny ipv4 192.0.0.0 0.0.0.255 any log-input

80 deny ipv4 192.0.2.0 0.0.0.255 any log-input

90 deny ipv4 192.168.0.0 0.0.255.255 any log-input

100 deny ipv4 198.18.0.0 0.1.255.255 any log-input

110 deny ipv4 198.51.100.0 0.0.0.255 any log-input

120 deny ipv4 203.0.113.0 0.0.0.255 any log-input

130 deny ipv4 224.0.0.0 31.255.255.255 any log-input

140 permit tcp any any established

150 permit tcp host x.12.1.9 host x.12.1.10 eq bgp

160 permit tcp host x.12.1.9 eq bgp host x.12.1.10

170 permit icmp host x.12.1.9 host x.12.1.10 echo

180 permit icmp host x.12.1.9 host x.12.1.10 echo-reply

…

…

…

260 deny ipv4 any any log-input

External Interfaces connected to the NIPRNet or SIPRNet

Review the inbound ACLs on external facing interfaces attached to the NIPRNet or SIPRNet to validate access control lists are configured to block inbound packets with IP sources addresses as documented in RFC5735 and RFC6598.

External Interfaces connected to a commercial ISP or other non-DoD network

Review the inbound ACLs on external facing interfaces validate access control lists are configured to block inbound packets with IP sources addresses as documented in RFC5735 and RFC6598 as well as address space that has been allocated to the RIRs but not assigned by the RIR to an ISP or other enterprise network. The full list of bogons can be found at the following link: www.team-cymru.org/Services/Bogons/fullbogons-ipv4.txt

Step 2: Verify that the inbound ACL applied to all external interfaces will block all traffic from Bogon source addresses.

interface GigabitEthernet0/0/0/1

ipv4 address x.12.1.10 255.255.255.2

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

If the router is not configured to block inbound packets with source Bogon IP addresses, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the perimeter to block inbound packets with Bogon source addresses.

Step 1: Configure an ACL containing the current Bogon prefixes as shown below.

RP/0/0/CPU0:R2(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 0.0.0.0 0.255.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 10.0.0.0 0.255.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 100.64.0.0 0.63.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 127.0.0.0 0.255.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 169.254.0.0 0.0.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 172.16.0.0 0.15.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 192.0.0.0 0.0.0.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 192.0.2.0 0.0.0.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 192.168.0.0 0.0.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 198.18.0.0 0.1.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 198.51.100.0 0.0.0.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 203.0.113.0 0.0.0.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip 224.0.0.0 31.255.255.255 any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any any established
RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.12.1.9 host x.12.1.10 eq bgp
RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.12.1.9 eq bgp host x.12.1.10
RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp host x.12.1.9 host x.12.1.10 echo
RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp host x.12.1.9 host x.12.1.10 echo-reply
…
…
…
RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input
RP/0/0/CPU0:R2(config-ipv4-acl)#end
Step 2: Apply the ACL inbound on all external interfaces.
RP/0/0/CPU0:R3(config)#int g0/0/0/1
RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in
RP/0/0/CPU0:R3(config-if)#end

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)

1.21 CISC-RT-000280 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco 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-216756
RULE ID: SV-216756r531087
```

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.

Step 1: Verify the interface connecting to ISP has an inbound ACL as shown in the example below.

interface GigabitEthernet0/0/0/2

description Link to ISP

ipv4 address x.22.1.15 255.255.255.252

ipv4 access-group ISP_ACL_INBOUND ingress

Step 2: Verify that the ACL only allows traffic to specific destination addresses (i.e. enclave’s NIPRNet address space) as shown in the example below.

ipv4 access-list ISP_ACL_INBOUND

10 permit tcp any any established

20 permit icmp host x.12.1.16 host x.12.1.17 echo

30 permit icmp host x.12.1.16 host x.12.1.17 echo-reply

40 permit tcp any host x.12.1.22 eq www

50 permit tcp any host x.12.1.23 eq www

60 permit esp any host x.12.1.24

70 permit ahp any host x.12.1.24

80 deny ipv4 any any log-input

Note: An Approved Gateway (AG) is any external connection from a DoD NIPRNet enclave to an Internet Service Provider, or network owned by a contractor, or non-DoD federal agency that has been approved by either the DoD CIO or the DoD Component CIO. This AG requirement does not apply to commercial cloud connections when the Cloud Service Provider (CSP) network is connected via the NIPRNet Boundary Cloud Access Point (BCAP).

If the ingress ACL bound to the interface connecting to an alternate gateway 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.

Step 1: Configure the ingress ACL 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 as shown in the example below.

RP/0/0/CPU0:R2(config)#ip access-list ISP_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any any established

RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp host x.12.1.16 host x.12.1.17 echo

RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp host x.12.1.16 host x.12.1.17 echo-reply

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any host x.12.1.22 eq www

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any host x.12.1.23 eq www

RP/0/0/CPU0:R2(config-ipv4-acl)#permit 50 any host x.12.1.24

RP/0/0/CPU0:R2(config-ipv4-acl)#permit 51 any host x.12.1.24

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#end

Step 2: Apply the ACL inbound on the ISP-facing interface.

RP/0/0/CPU0:R3(config)#int g0/0/0/2

RP/0/0/CPU0:R3(config-if)#ipv4 access-group ISP_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#end

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 CISC-RT-000290 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

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

```
GROUP ID: V-216757
RULE ID: SV-216757r531087
```

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.

Step 1: Configure the ingress ACL 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 as shown in the example below.

RP/0/0/CPU0:R2(config)#ip access-list ISP_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)# permit tcp any any established

RP/0/0/CPU0:R2(config-ipv4-acl)# permit icmp host x.12.1.16 host x.12.1.17 echo

RP/0/0/CPU0:R2(config-ipv4-acl)# permit icmp host x.12.1.16 host x.12.1.17 echo-reply

RP/0/0/CPU0:R2(config-ipv4-acl)# permit tcp any host x.12.1.22 eq www

RP/0/0/CPU0:R2(config-ipv4-acl)# permit tcp any host x.12.1.23 eq www

RP/0/0/CPU0:R2(config-ipv4-acl)# permit 50 any host x.12.1.24

RP/0/0/CPU0:R2(config-ipv4-acl)# permit 51 any host x.12.1.24

RP/0/0/CPU0:R2(config-ipv4-acl)# deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#end

Step 2: Apply the ACL inbound on the ISP-facing interface.

RP/0/0/CPU0:R3(config)#int g0/0/0/2

RP/0/0/CPU0:R3(config-if)#ipv4 access-group ISP_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#end

If any BGP neighbors belonging to the alternate gateway service provider exist, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Remove any BGP neighbors belonging to the alternate gateway service provider and configure a static route to forward Internet bound traffic to the alternate gateway as shown in the example below.

R5(config)#ip route 0.0.0.0 0.0.0.0 x.22.1.14

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

1.23 CISC-RT-000300 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216758
RULE ID: SV-216758r531087
```

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.

Step 1: Review the IGP and BGP configurations. If there are redistribute static statements configured as shown in examples below proceed to step 2.

OSPF Example

router ospf 1

redistribute static

EIGRP example

router eigrp 1

address-family ipv4

redistribute static

RIP example

router rip

version 2

redistribute static

BGP example

router bgp n

address-family ipv4 unicast

redistribute static

Step 2: Review the static routes that have been configured to determine if any contain the next hop address of the alternate gateway.

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

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router so that static routes are not redistributed to an alternate gateway into either a BGP or any IGP peering with the NIPRNet or to any other autonomous systems. This can be done by excluding that route in the route policy as shown in the example below.

Step 1: Configure a prefix list for any static routes with the alternate gateway as the next-hop address

RP/0/0/CPU0:R3(config)#prefix-set ISP_PREFIX

RP/0/0/CPU0:R3(config-pfx)#x.x.0.0/24 ge 8 le 32

RP/0/0/CPU0:R3(config-pfx)#end-set

Step 2: Configure a route policy

RP/0/0/CPU0:R3(config)#route-policy FILTER_ISP_ROUTES

RP/0/0/CPU0:R3(config-rpl)#if destination in ISP_PREFIX then

RP/0/0/CPU0:R3(config-rpl-if)#drop

RP/0/0/CPU0:R3(config-rpl-if)#else

RP/0/0/CPU0:R3(config-rpl-else)#pass

RP/0/0/CPU0:R3(config-rpl-else)#endif

RP/0/0/CPU0:R3(config-rpl)#end-pol

Step 3: Apply the route policy to the IGP and BGP redistribute static commands as shown in the OSPF example.

RP/0/0/CPU0:R3(config-ospf)#redistribute static route-policy FILTER_ISP_ROUTES

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

1.24 CISC-RT-000310 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco 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-217005
RULE ID: SV-217005r945858
```

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", which are 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 ACL 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.

uRPF example

interface GigabitEthernet0/0/0/1

description downstream link to LAN

ipv4 address 10.1.34.3 255.255.255.0

ipv4 verify unicast source reachable-via rx

Egress ACL example

ipv4 access-list EGRESS_FILTER

10 permit udp 10.1.15.0 0.0.0.255 any eq domain

20 permit tcp 10.1.15.0 0.0.0.255 any eq ftp

30 permit tcp 10.1.15.0 0.0.0.255 any eq ftp-data

40 permit tcp 10.1.15.0 0.0.0.255 any eq www

50 permit icmp 10.1.15.0 0.0.0.255 any

60 permit icmp 10.1.15.0 0.0.0.255 any echo

70 deny ipv4 any any

…

…

…

interface GigabitEthernet0/0/0/1

description downstream link to LAN

ipv4 address 10.1.34.3 255.255.255.0

ipv4 access-group EGRESS_FILTER ingress

If uRPF or an egress ACL 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 in an enclave, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router to ensure that an egress ACL or uRPF is configured on internal interfaces to restrict the router from accepting any outbound IP packet that contains an illegitimate address in the source field. The example below enables uRPF.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 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)

1.25 CISC-RT-000320 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to filter traffic destined to the enclave in accordance with the guidelines contained in DoD Instruction 8551.1.

```
GROUP ID: V-216760
RULE ID: SV-216760r1007832
```

Rationale:

Vulnerability assessments must be reviewed by the System Administrator, and protocols must be approved by the Information Assurance (IA) staff before entering the enclave.

Access control lists (ACLs) are the first line of defense in a layered security approach. They permit authorized packets and deny unauthorized packets based on port or service type. They enhance the posture of the network by not allowing packets to reach a potential target within the security domain. The lists provided are highly susceptible ports and services that should be blocked or limited as much as possible without adversely affecting customer requirements. Auditing packets attempting to penetrate the network but that are stopped by an ACL will allow network administrators to broaden their protective ring and more tightly define the scope of operation.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to verify that the ingress ACL is in accordance with DoD 8551.1.

Step 1: Verify that an inbound ACL is configured on all external interfaces.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

Step 2. Review the inbound ACL to verify that it is filtering traffic in accordance with DoD 8551.1.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.11.1.1 eq bgp host x.11.1.2

20 permit tcp host x.11.1.1 host x.11.1.2 eq bgp

30 permit icmp host x.11.1.1 host x.11.1.2 echo

40 permit icmp host x.11.1.1 host x.11.1.2 echo-reply

50 deny ipv4 any host x.11.1.1 log

60 permit tcp any host x.12.1.22 eq www

70 permit tcp any any established

…

… < must be in accordance with DoD Instruction 8551.1>

…

160 deny ipv4 any any log-input

If the router does not filter traffic in accordance with the guidelines contained in DoD 8551.1, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router to use an inbound ACL on all external interfaces as shown in the example below to restrict traffic in accordance with the guidelines contained in DOD Instruction 8551.1.

RP/0/0/CPU0:R3(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any host x.11.1.1 log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any host x.12.1.22 eq www

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any established

…

… < must be in accordance with DoD Instruction 8551.1>

…

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

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

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#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

1.26 CISC-RT-000330 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216761
RULE ID: SV-216761r531087
```

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 router configuration to verify that an inbound ACL is configured on all external interfaces as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.11.1.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

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

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router to use an inbound ACL on all external interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_INBOUND in

RP/0/0/CPU0:R3(config-if)#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

1.27 CISC-RT-000340 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216762
RULE ID: SV-216762r531087
```

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

interface GigabitEthernet0/0/0/1

description downstream link to LAN

ipv4 address 10.1.34.3 255.255.255.0

ipv4 access-group EGRESS_FILTER ingress

If the 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.

Configure the router to use an inbound ACL on all internal interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#int g0/0/0/1

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EGRESS_FILTER in

RP/0/0/CPU0:R3(config-if)#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

1.28 CISC-RT-000350 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to block all packets with any IP options.

```
GROUP ID: V-217006
RULE ID: SV-217006r945859
```

Rationale:

Packets with IP options are not fast switched 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:

This requirement is not applicable for the DODIN Backbone.

In Cisco IOS XR, all IPv4 packets with any header option other than the "source-route" header options are dropped. By default, ipv4 source routing is disabled. Verify that the following command is not configured: ipv4 source-route

If the router is not configured to drop all packets with IP option source routing, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router to drop all packets with IP option source routing.

RP/0/0/CPU0:R3(config)#no ipv4 source-route

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)

1.29 CISC-RT-000360 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco perimeter router must be configured to have Link Layer Discovery Protocol (LLDP) disabled on all external interfaces.

```
GROUP ID: V-216764
RULE ID: SV-216764r856442
```

Rationale:

LLDP is a neighbor discovery protocol used to advertise device capabilities, configuration information, and device identity. LLDP is media- and protocol-independent as it runs over layer 2; therefore, two network nodes that support different layer 3 protocols can still learn about each other. Allowing LLDP messages to reach external network nodes 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.

Step 1: Verify LLDP is not enabled globally via the command

lldp

By default LLDP is not enabled globally. If LLDP is enabled, proceed to step 2.

Step 2: Verify LLDP transmit is disabled on any external interface as shown in the example below.

interface GigabitEthernet0/0/0/1

ipv4 address x.1.34.3 255.255.255.252

lldp

transmit disable

Note: LLDP is enabled by default on all interfaces once it is enabled globally; hence the commands lldp transmit and lldp receive will not be visible on the interface configuration.

If LLDP transmit is enabled on any external interface, this is a finding.

Remediation:

Disable LLDP transmit on all external interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#interface g0/0/0/1

RP/0/0/CPU0:R3(config-if)#lldp transmit disable

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)

1.30 CISC-RT-000370 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco perimeter router must be configured to have Cisco Discovery Protocol (CDP) disabled on all external interfaces.

```
GROUP ID: V-216765
RULE ID: SV-216765r856443
```

Rationale:

CDP is a Cisco proprietary neighbor discovery protocol used to advertise device capabilities, configuration information, and device identity. CDP is media- and protocol-independent as it runs over layer 2; therefore, two network nodes that support different layer 3 protocols can still learn about each other. Allowing CDP messages to reach external network nodes 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.

Step 1: Verify if CDP is enabled globally via the cdp command as shown below.

hostname R3

cdp

By default CDP is disabled globally; hence, the command cdp run will not be shown in the configuration. If CDP is enabled, proceed to step 2.

Step 2: Verify CDP is not enabled on any external interface as shown in the example below.

interface GigabitEthernet0/0/0/1

cdp

ipv4 address x.1.34.3 255.255.255.252

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

Remediation:

This requirement is not applicable for the DODIN Backbone.

Disable CDP on all external interfaces via no cdp command or disable CDP globally via no cdp command.

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)

1.31 CISC-RT-000380 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to have Proxy ARP disabled on all external interfaces.

```
GROUP ID: V-216766
RULE ID: SV-216766r856444
```

Rationale:

When Proxy ARP is enabled on a 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 router configuration to determine if IP Proxy ARP is enabled on any external interfaces as shown in the example below.

interface GigabitEthernet0/0/0/1

description link to DISN

ipv4 address x.1.12.2 255.255.255.252

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 Proxy ARP on all external interfaces as shown in the example below.

RP/0/0/CPU0:R3(config)#interface g0/0/0/1

RP/0/0/CPU0:R3(config-if)#no 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)

1.32 CISC-RT-000390 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to block all outbound management traffic.

```
GROUP ID: V-216767
RULE ID: SV-216767r945857
```

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 perimeter router of the managed network must be configured with an outbound ACL on the egress interface to block all management traffic as shown in the example below.

Step 1: Verify that all external interfaces has been configured with an outbound ACL as shown in the example below.

interface GigabitEthernet0/0/0/2

ipv4 address 10.1.35.3 255.255.255.0

ipv4 access-group EXTERNAL_ACL_OUTBOUND egress

Step 2: Verify that the outbound ACL discards management traffic as shown in the example below.

ipv4 access-list EXTERNAL_ACL_OUTBOUND

10 deny tcp any any eq tacacs log-input

20 deny tcp any any eq ssh log-input

30 deny udp any any eq snmp log-input

40 deny udp any any eq snmptrap log-input

50 deny udp any any eq syslog log-input

60 permit tcp any any eq www log-input

70 deny ipv4 any any log-input

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 perimeter router of the managed network with an outbound ACL on the egress interface to block all management traffic.

Step 1: Configure an ACL to block egress management traffic.

RP/0/0/CPU0:R3(config)#Ipv4 access-list EXTERNAL_ACL_OUTBOUND

RP/0/0/CPU0:R3(config-ipv4-acl)#deny tcp any any eq tacacs log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny tcp any any eq 22 log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny udp any any eq snmp log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny udp any any eq snmptrap log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny udp any any eq syslog log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp any any eq www log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#end

Note: Permit commands would be configured to allow applicable outbound traffic. The example above is allowing web traffic.

Step 2: Configure the external interfaces with the outbound ACL.

RP/0/0/CPU0:R3(config)#int g0/0/0/2

RP/0/0/CPU0:R3(config-if)#ipv4 access-group EXTERNAL_ACL_OUTBOUND egress

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 CISC-RT-000391 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to suppress Router Advertisements on all external IPv6-enabled interfaces.

```
GROUP ID: V-230046
RULE ID: SV-230046r533007
```

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 router configuration to verify that Router Advertisements are suppressed on all external IPv6-enabled interfaces as shown in the example below.

interface gigabitethernet1/0

ipv6 address 2001::1:0:22/64

ipv6 nd ra suppress

If the router is not configured to suppress Router Advertisements on all external IPv6-enabled interfaces, this is a finding.

Remediation:

Configure the router to suppress Router Advertisements on all external IPv6-enabled interfaces as shown in the example below.

R1(config)#int g1/0

R1(config-if)#ipv6 nd ra suppress

R1(config-if)#end

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

1.34 CISC-RT-000392 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 undetermined transport packets.

```
GROUP ID: V-230049
RULE ID: SV-230049r950991
```

Rationale:

One of the fragmentation weaknesses known in IPv6 is the undetermined transport packet. This packet contains an undetermined protocol due to fragmentation. Depending on the length of the IPv6 extension header chain, the initial fragment may not contain the layer four port information of the packet.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to determine if it is configured to drop IPv6 undetermined transport packets.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops undetermined transport packets as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny ipv6 any any log undetermined-transport

20 permit ipv6 …

…

…

…

90 deny ipv6 any any log

If the router is not configured to drop IPv6 undetermined transport packets, this is a finding.

Remediation:

Configure the router to drop IPv6 undetermined transport packets as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any undetermined-transport log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

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)

1.35 CISC-RT-000393 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured drop IPv6 packets with a Routing Header type 0, 1, or 3–255.

```
GROUP ID: V-230052
RULE ID: SV-230052r856665
```

Rationale:

The routing header can be used maliciously to send a packet through a path where less robust security is in place, rather than through the presumably preferred path of routing protocols. Use of the routing extension header has few legitimate uses other than as implemented by Mobile IPv6.

The Type 0 Routing Header (RFC 5095) is dangerous because it allows attackers to spoof source addresses and obtain traffic in response, rather than the real owner of the address. Secondly, a packet with an allowed destination address could be sent through a Firewall using the Routing Header functionality, only to bounce to a different node once inside. The Type 1 Routing Header is defined by a specification called "Nimrod Routing", a discontinued project funded by DARPA. Assuming that most implementations will not recognize the Type 1 Routing Header, it must be dropped. The Type 3–255 Routing Header values in the routing type field are currently undefined and should be dropped inbound and outbound.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to determine if it is configured to drop IPv6 packets containing a Routing Header of type 0, 1, or 3-255.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets with a Routing Header type 0, 1, or 3-255

as shown in the example below.

ipv6 access-list FILTER_IPV6

10 permit ipv6 any host 2001:DB8::1:1:1234 routing-type 2

20 deny ipv6 any any log routing

30 permit ipv6 …

…

…

…

90 deny ipv6 any any log

Note: The example above allows routing-type 2 in the event Mobility IPv6 is deployed.

If the router is not configured to drop IPv6 packets containing a Routing Header of type 0, 1, or 3-255, this is a finding.

Remediation:

Configure the router to drop IPv6 packets with Routing Header of type 0, 1, or 3-255 as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 any host 2001:DB8::0:1:1:1234 routing-type 2

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any routing log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit …

…

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

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)

1.36 CISC-RT-000394 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 packets containing a Hop-by-Hop header with invalid option type values.

```
GROUP ID: V-230147
RULE ID: SV-230147r856667
```

Rationale:

These options are intended to be for the Destination Options header only. The optional and extensible natures of the IPv6 extension headers require higher scrutiny since many implementations do not always drop packets with headers that it cannot recognize and hence could cause a Denial-of-Service on the target device. In addition, the type, length, value (TLV) formatting provides the ability for headers to be very large.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to determine if it is compliant with this requirement.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets containing a Hop-by-Hop header with option type values of 0x04 (Tunnel Encapsulation Limit), 0xC9 (Home Address Destination), or 0xC3 (NSAP Address) as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny hbh any any dest-option-type 4 log

20 deny hbh any any dest-option-type 195 log

30 deny hbh any any dest-option-type home-address log

40 permit ipv6 …

…

…

…

90 deny ipv6 any any log

If the router is not configured to drop IPv6 packets containing a Hop-by-Hop header with invalid option type values, this is a finding.

Remediation:

Drop IPv6 packets containing a Hop-by-Hop header as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny hbh any any dest-option-type 4 log

RP/0/0/CPU0:R3(config-ipv6-acl)# deny hbh any any dest-option-type 195 log

RP/0/0/CPU0:R3(config-ipv6-acl)# deny hbh any any dest-option-type home-address log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …

…

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

RP/0/0/CPU0:R3(config-if)# end

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)

1.37 CISC-RT-000395 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 packets containing a Destination Option header with invalid option type values.

```
GROUP ID: V-230151
RULE ID: SV-230151r856669
```

Rationale:

These options are intended to be for the Hop-by-Hop header only. The optional and extensible natures of the IPv6 extension headers require higher scrutiny since many implementations do not always drop packets with headers that it cannot recognize. Hence, this could cause a Denial-of-Service on the target device. In addition, the type, length, value (TLV) formatting provides the ability for headers to be very large.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to determine if it is compliant with this requirement.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets containing a Destination Option header with option type values of 0x05 (Router Alert) or 0xC2 (Jumbo Payload) as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny 60 any any dest-option-type 5 log

20 deny 60 any any dest-option-type 194 log

30 permit ipv6 …

…

…

…

90 deny ipv6 any any log

If the router is not configured to drop IPv6 packets containing a Destination Option header with option type values of 0x05 (Router Alert) or 0xC2 (Jumbo Payload), this is a finding.

Remediation:

Configure the router to drop IPv6 packets containing a Destination Option header with option type values of 0x05 (Router Alert) or 0xC2 (Jumbo Payload) as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny 60 any any dest-option-type 5 log

RP/0/0/CPU0:R3(config-ipv6-acl)# deny 60 any any dest-option-type 194 log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …

…

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

RP/0/0/CPU0:R3(config-if)# end

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)

1.38 CISC-RT-000396 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 packets containing an extension header with the Endpoint Identification option.

```
GROUP ID: V-230154
RULE ID: SV-230154r856671
```

Rationale:

The optional and extensible natures of the IPv6 extension headers require higher scrutiny since many implementations do not always drop packets with headers that it cannot recognize, and hence could cause a Denial-of-Service on the target device. In addition, the type, length, value (TLV) formatting provides the ability for headers to be very large. This option type is associated with the Nimrod Routing system and has no defining RFC document.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration to determine if it is compliant with this requirement.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets containing an extension header with the Endpoint Identification option as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny any any dest-option-type 138 log

20 permit ipv6 …

…

…

90 deny ipv6 any any log

If the router is not configured to drop IPv6 packets containing an extension header with the Endpoint Identification option, this is a finding.

Remediation:

Configure the router to drop IPv6 packets containing an option type values of 0x8A (Endpoint Identification) regardless of whether it appears in a Hop-by-Hop or Destination Option header as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 138 log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …

…

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

RP/0/0/CPU0:R3(config-if)# end

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)

1.39 CISC-RT-000397 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 packets containing the NSAP address option within Destination Option header.

```
GROUP ID: V-230157
RULE ID: SV-230157r856673
```

Rationale:

The optional and extensible natures of the IPv6 extension headers require higher scrutiny since many implementations do not always drop packets with headers that it cannot recognize, and hence could cause a Denial-of-Service on the target device. In addition, the type, length, value (TLV) formatting provides the ability for headers to be very large. This option type from RFC 1888 (OSI NSAPs and IPv6) has been deprecated by RFC 4048.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration and determine if filters are bound to the applicable interfaces to drop IPv6 packets containing the NSAP address option within Destination Option header.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets containing the NSAP address option within Destination Option header as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny 60 any any dest-option-type 195 log

20 permit ipv6 …

…

…

…

90 deny ipv6 any any log

If the router is not configured to drop IPv6 packets containing the NSAP address option within Destination Option header, this is a finding.

Remediation:

Configure the router to drop IPv6 packets containing a Destination Option header with option type value of 0xC3 (NSAP address) as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny 60 any any dest-option-type 195 log

RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …

…

…

RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log

RP/0/0/CPU0:R3(config-ipv6-acl)# exit

RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2

RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress

RP/0/0/CPU0:R3(config-if)# end

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)

1.40 CISC-RT-000398 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco perimeter router must be configured to drop IPv6 packets containing a Hop-by-Hop or Destination Option extension header with an undefined option type.

```
GROUP ID: V-230160
RULE ID: SV-230160r856675
```

Rationale:

The optional and extensible natures of the IPv6 extension headers require higher scrutiny since many implementations do not always drop packets with headers that it cannot recognize, and hence could cause a Denial-of-Service on the target device. In addition, the type, length, value (TLV) formatting provides the ability for headers to be very large.

Audit:

This requirement is not applicable for the DODIN Backbone.

Review the router configuration and determine if filters are bound to the applicable interfaces to drop all inbound IPv6 packets containing an undefined option type value regardless of whether they appear in a Hop-by-Hop or Destination Option header. Undefined values are 0x02, 0x03, 0x06, 0x9 – 0xE, 0x10 – 0x22, 0x24, 0x25, 0x27 – 0x2F, and 0x31 – 0xFF.

Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.

interface interface gigabitethernet 0/2/0/2

ipv6 address 2001::1:0:22/64

ipv6 access-group FILTER_IPV6 ingress

Step 2: Verify that the ACL drops IPv6 packets containing a Hop-by-Hop or Destination Option extension header with an undefined option type as shown in the example below.

ipv6 access-list FILTER_IPV6

10 deny any any dest-option-type 2

20 deny any any dest-option-type 3

30 deny any any dest-option-type 6

40 deny any any dest-option-type 9

50 deny any any dest-option-type 10

60 deny any any dest-option-type 11

70 deny any any dest-option-type 12

80 deny any any dest-option-type 13

90 deny any any dest-option-type 14

100 deny any any dest-option-type 16

…

280 deny any any dest-option-type 34

290 deny any any dest-option-type 36

300 deny any any dest-option-type 37

310 deny any any dest-option-type 39

…

390 deny any any dest-option-type 47

400 deny any any dest-option-type 49

…

nnn deny any any dest-option-type 255

nnn permit …

…

…

…

nnn deny ipv6 any any log

Note: Because hop-by-hop and destination options have the same exact header format, they can be combined under the dest-option-type keyword. Since Hop-by-Hop and Destination Option headers have non-overlapping types, you can use dest-option-type to match either.

If the router is not configured to drop IPv6 packets containing a Hop-by-Hop or Destination Option extension header with an undefined option type, this is a finding.

Remediation:

Configure the router to drop all inbound IPv6 packets containing an undefined option type value regardless of whether they appear in a Hop-by-Hop or Destination Option header as shown in the example below.

RP/0/0/CPU0:R3(config)# ipv6 access-list FILTER_IPV6

RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 2

```
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 3
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 6
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 9
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 10
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 11
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 12
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 13
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 14
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 16
…
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 34
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 36
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 37
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 39
…
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 47
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 49
…
RP/0/0/CPU0:R3(config-ipv6-acl)# deny any any dest-option-type 255
RP/0/0/CPU0:R3(config-ipv6-acl)# permit ipv6 …
…
…
…
RP/0/0/CPU0:R3(config-ipv6-acl)# deny ipv6 any any log
RP/0/0/CPU0:R3(config-ipv6-acl)# exit
RP/0/0/CPU0:R3(config)# interface gigabitethernet 0/2/0/2
RP/0/0/CPU0:R3(config-if)# ipv6 access-group FILTER_IPV6 ingress
RP/0/0/CPU0:R3(config-if)# end
```

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)

1.41 CISC-RT-000400 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco out-of-band management (OOBM) gateway router must be configured to transport management traffic to the Network Operations Center (NOC) via dedicated circuit, MPLS/VPN service, or IPsec tunnel.

```
GROUP ID: V-216768
RULE ID: SV-216768r991934
```

Rationale:

Using dedicated paths, the OOBM backbone connects the OOBM gateway routers located at the edge of the managed network and at the NOC. Dedicated links can be deployed using provisioned circuits or MPLS Layer 2 and Layer 3 VPN services or implementing a secured path with gateway-to-gateway IPsec tunnels. The tunnel mode ensures that the management traffic will be logically separated from any other traffic traversing the same path.

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 and interface that the management traffic traverses. If an IPsec tunnel is used to transport the management traffic between the NOC and the managed network, review the configuration following the steps below.

Step 1: Note the profile referenced for the IPsec tunnel to the NOC.

interface tunnel-ipsec 30

profile IPSEC_NOC_PROFILE

tunnel source GigabitEthernet0/0/0/2

tunnel destination x.1.22.2

Step 2: Note the crypto ACL that was specified in the IPsec profile.

crypto isakmp keyring ISAKMP_KEYRING

pre-shared-key address x.1.22.2 255.255.255.255 key encrypted 150A13141C32

!

crypto isakmp policy 10

hash sha256

encryption aes 256

authentication pre-share

!

crypto ipsec transform-set IPSEC_TRANS esp-aes 256 esp-sha256-hmac

mode tunnel

!

crypto ipsec profile IPSEC_NOC_PROFILE

set pfs group16

match address MGMT_TRAFFIC_ACL

Step 3: Review the ACL defined in the crypto map and verify that the destination is the management network.

!

ipv4 access-list MGMT_TRAFFIC_ACL

10 permit ipv4 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255

!

Note: The management network is this example is 10.22.2.0/24

If management traffic is not transported between the managed network and the NOC via dedicated circuit, MPLS/VPN service, or IPsec tunnel, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Ensure that a dedicated circuit, MPLS/VPN service, or IPsec tunnel is deployed to transport management traffic between the managed network and the NOC. If an IPsec tunnel is to be used, the steps below can be used as a guideline.

Step 1: Configure the ACL for the management network as the destination. This ACL will be defined in the crypto as the interesting traffic to be forwarded into the IPsec tunnel.

RP/0/0/CPU0:R3(config)#Ipv4 access-list MGMT_TRAFFIC_ACL

RP/0/0/CPU0:R3(config-ipv4-acl)#permit ip 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255

Step 2: Create an ISAKMP policy for Phase 1 negotiations.

RP/0/0/CPU0:R3(config)#crypto isakmp policy 10

RP/0/0/CPU0:R3(config-isakmp-policy)#authentication pre-share

RP/0/0/CPU0:R3(config-isakmp-policy)#hash sha256

RP/0/0/CPU0:R3(config-isakmp-policy)#encryption aes 256

RP/0/0/CPU0:R3(config-isakmp-policy)#exit

Step 3: Configure the keyring to be used for ISAKMP to authenticate the remote side during IKE negotiation.

RP/0/0/CPU0:R3(config)#crypto isakmp keyring ISAKMP_KEYRING

RP/0/0/CPU0:R3(config-crypto-keyring)#pre-shared-key address 255.255.255.255 key xxxxx

RP/0/0/CPU0:R3(config-crypto-keyring)#exit

Step 4: Configure the IPsec transform set.

RP/0/0/CPU0:R3(config)#crypto ipsec transform-set IPSEC_TRANS esp-aes 256 esp-sha256-hmac

RP/0/0/CPU0:R3(config-transform-set IPSEC_TRANS)#mode tunnel

Step 5: Configure the IPsec profile.

RP/0/0/CPU0:R3(config)#crypto ipsec profile IPSEC_NOC_PROFILE

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#set pfs group 16

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#match MGMT_TRAFFIC_ACL transform-set IPSEC_TRANS

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#exit

Step 6: Configure the IPsec virtual interface.

RP/0/0/CPU0:R3(config)#interface tunnel-ipsec 22

RP/0/0/CPU0:R3(config-if)#profile IPSEC_NOC_PROFILE

RP/0/0/CPU0:R3(config-if)#tunnel source GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-if)#tunnel destination x.1.22.2

RP/0/0/CPU0:R3(config-if)#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

CCI-004891 Implement physically or logically separate subnetworks to isolate organization-defined critical system components and functions.

- NIST SP 800-53 Revision 5::SC-7 (29)

1.42 CISC-RT-000410 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216769
RULE ID: SV-216769r531087
```

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. Verify that only management traffic is forwarded through the OOBM interface or IPsec tunnel.

If an OOBM link is used, verify that the only authorized management traffic is transported to the NOC by reviewing the outbound ACL applied to the OOBM interface as shown in the example below.

Step 1: Note the outbound ACL applied to the OOBM interface.

interface GigabitEthernet0/0/0/2

description OOB link to NOC

ipv4 address 10.11.1.8 255.255.255.0

ipv4 access-group MGMT_TRAFFIC_ACL egress

Step 2: Review the outbound ACL and verify only management traffic is forwarded to the NOC.

ipv4 access-list MGMT_TRAFFIC_ACL

10 permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs

20 permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq ssh

30 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp

40 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmptrap

50 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog

60 permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255 echo

70 permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255 echo-reply

80 deny ipv4 any any log-input

If an IPSec tunnel is used, verify that the only authorized management traffic is transported to the NOC.

Step 1: Note the profile referenced for the IPSec tunnel to the NOC.

interface tunnel-ipsec 30

profile IPSEC_NOC_PROFILE

tunnel source GigabitEthernet0/0/0/2

tunnel destination x.1.22.2

Step 2: Note the crypto ACL that was specified in the IPSec profile.

crypto isakmp keyring ISAKMP_KEYRING

pre-shared-key address x.1.22.2 255.255.255.255 key encrypted 150A13141C32

!

crypto isakmp policy 10

hash sha256

encryption aes 256

authentication pre-share

!

crypto ipsec transform-set IPSEC_TRANS esp-aes 256 esp-sha256-hmac

mode tunnel

!

crypto ipsec profile IPSEC_NOC_PROFILE

set pfs group16

match address MGMT_TRAFFIC_ACL

Step 3: Review the crypto ACL defined in the IPSec profile and verify only management traffic is forwarded to the NOC.

ipv4 access-list MGMT_TRAFFIC_ACL

10 permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs

20 permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq ssh

30 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp

40 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmptrap

50 permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog

60 permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255

Note: ICMP is permitted for troubleshooting purposes. The IPSec SA can only identify interesting traffic via address, protocol, and port; hence, the ICMP traffic cannot be qualified via type attribute.

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 ACLs to permit only authorized management traffic into IPsec tunnels or the OOBM interface used for forwarding management data as shown in the examples below.

OOBM Link

RP/0/0/CPU0:R3(config)#ipv4 access-list MGMT_TRAFFIC_ACL

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq ssh

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmptrap

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ipv4 any any log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#exit

IPsec Tunnel

RP/0/0/CPU0:R3(config)#ipv4 access-list MGMT_TRAFFIC_ACL

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs

RP/0/0/CPU0:R3(config-ipv4-acl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq ssh

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmptrap

RP/0/0/CPU0:R3(config-ipv4-acl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog

RP/0/0/CPU0:R3(config-ipv4-acl)#permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255

RP/0/0/CPU0:R3(config-ipv4-acl)#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

1.43 CISC-RT-000420 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216770
RULE ID: SV-216770r531087
```

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 that the OOBM interface is an adjacency in the IGP domain for the management network via separate VRF as shown in the example below.

router ospf 2

vrf MGMT

area 0

interface GigabitEthernet0/0/0/0.2

!

!

!

!

router ospf 3

vrf PROD

area 0

interface GigabitEthernet0/0/0/0.3

!

!

!

!

If the router is not configured to have separate IGP instances for the managed network and management network, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Configure the router to have a separate IGP instance for the management network as shown in the example below.

RP/0/0/CPU0:R2(config)#router ospf 2 vrf MGMT

RP/0/0/CPU0:R2(config-ospf-vrf)#area 0

RP/0/0/CPU0:R2(config-ospf-vrf-ar)#interface GigabitEthernet0/0/0/0.2

RP/0/0/CPU0:R2(config-ospf-vrf-ar-if)#exit

RP/0/0/CPU0:R2(config-ospf-vrf-ar)#exit

RP/0/0/CPU0:R2(config-ospf-vrf)#exit

RP/0/0/CPU0:R2(config-ospf)#exit

RP/0/0/CPU0:R2(config)#router ospf 3 vrf PROD

RP/0/0/CPU0:R2(config-ospf-vrf)#area 0

RP/0/0/CPU0:R2(config-ospf-vrf-ar)#interface GigabitEthernet0/0/0/0.3

RP/0/0/CPU0:R2(config-ospf-vrf-ar-if)#end

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

1.44 CISC-RT-000430 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216771
RULE ID: SV-216771r531087
```

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 (IGP) instance used for the managed network does not redistribute routes into the IGP instance used for the management network, and vice versa. The example below imports OSPF routes from the production route table (VRF PROD) into the management route table (VRF MGMT) using BGP.

vrf MGMT

address-family ipv4 unicast

import route-target

4:4

8:8

!

export route-target

4:4

!

```
!
!
vrf PROD
address-family ipv4 unicast
import route-target
8:8
!
export route-target
8:8
!
!
!
…
…
…
router ospf 2
vrf MGMT
redistribute bgp 64512
area 0
interface GigabitEthernet0/0/0/0.2
!
!
!
!
router ospf 3
vrf PROD
area 0
interface GigabitEthernet0/0/0/0.3
!
!
```

```
!
!
router bgp 64512
address-family ipv4 unicast
!
address-family vpnv4 unicast
!
vrf MGMT
rd 4:4
address-family ipv4 unicast
redistribute ospf 2
!
!
vrf PROD
rd 8:8
address-family ipv4 unicast
redistribute ospf 3
!
!
!
```

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

Remediation:

This requirement is not applicable for the DODIN Backbone.

Remove the configuration that imports routes from the managed network into the management network or vice versa as shown in the example below.

```
RP/0/0/CPU0:R2(config)#vrf MGMT
RP/0/0/CPU0:R2(config-vrf)#address-family ipv4 unicast
RP/0/0/CPU0:R2(config-vrf-af)#no import route-target 8:8
RP/0/0/CPU0:R2(config-vrf-af)#end
```

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

1.45 CISC-RT-000440 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216772
RULE ID: SV-216772r531087
```

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. It is only applicable if the OOBM gateway router is not a dedicated device to the OOBM backbone.

Verify that traffic destined to itself is only sourced by the OOBM or the NOC. In the example below, the OOBM backbone network is 10.11.1.0/24, the NOC address spaces is 10.12.1.0/24, and the OOBM LAN address space at remote site connecting to the managed network is 10.13.1.0/24.

Step 1: Note the inbound ACL applied to the OOBM interfaces.

interface GigabitEthernet0/0/0/2

description OOB link to NOC

ip address 10.11.1.8 255.255.255.0

ipv4 access-group TRAFFIC_FROM_NOC ingress

!

interface Giga GigabitEthernet0/0/0/3

description link to OOBM LAN access switch

ip address 10.13.1.1 255.255.255.0

ipv4 access-group TRAFFIC_TO_NOC ingress

Step 2: Review the inbound ACL bound to any OOB interface connecting to the OOBM backbone and verify traffic destined to itself is only from the OOBM or NOC address space.

ipv4 access-list TRAFFIC_FROM_NOC

10 permit ipv4 10.11.1.0 0.255.255.255 host 10.11.1.8

20 permit ipv4 10.12.1.0 0.255.255.255 host 10.11.1.8

30 permit ipv4 10.11.1.0 0.255.255.255 host 10.13.1.1

40 permit ipv4 10.12.1.0 0.255.255.255 host 10.13.1.1

50 deny ipv4 any host 10.11.1.8 log-input

60 deny ipv4 any host 10.13.1.1 log-input

70 permit ipv4 10.11.1.0 0.0.0.255 10.13.1.0 0.0.0.255

80 permit ipv4 10.12.1.0 0.0.0.255 10.13.1.0 0.0.0.255

90 deny ipv4 any any log-input

Step 3: Review the inbound ACL bound to any OOBM LAN interfaces and verify traffic destined to itself is from the OOBM LAN address space.

ipv4 access-list TRAFFIC_TO_NOC

10 permit ipv4 10.13.1.0 0.255.255.255 host 10.13.1.1

20 permit ipv4 10.13.1.0 0.255.255.255 host 10.11.1.8

30 deny ipv4 any host 10.13.1.1 log-input

40 deny ipv4 any host 10.11.1.8 log-input

50 permit ipv4 10.13.1.0 0.255.255.255 10.11.1.0 0.0.0.255

60 permit ipv4 10.13.1.0 0.255.255.255 10.12.1.0 0.0.0.255

70 deny ipv4 any any log-input

If the 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. It is only applicable if the OOBM gateway router is not a dedicated device to the OOBM backbone.

Step 1: Configure the ACL to only allow traffic to the route processor from the OOBM backbone and the NOC.

RP/0/0/CPU0:R2(config)#ipv4 access-list TRAFFIC_FROM_NOC

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.11.1.0 0.255.255.255 host 10.11.$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.12.1.0 0.255.255.255 host 10.11.$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.11.1.0 0.255.255.255 host 10.13.$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.12.1.0 0.255.255.255 host 10.13.$

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any host 10.11.1.8 log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any host 10.13.1.1 log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.11.1.0 0.0.0.255 10.13.1.0 0.0.0$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.12.1.0 0.0.0.255 10.13.1.0 0.0.0$

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

Step 2: Configure the ACL to only allow traffic to the route processor from the OOBM LAN.

RP/0/0/CPU0:R2(config)#ipv4 access-list TRAFFIC_TO_NOC

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.13.1.0 0.255.255.255 host 10.13.$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.13.1.0 0.255.255.255 host 10.11.$

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any host 10.13.1.1 log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any host 10.11.1.8 log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.13.1.0 0.255.255.255 10.11.1.0 0$

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip 10.13.1.0 0.255.255.255 10.12.1.0 0$

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

Step 3: Apply the ACLs configured above to the appropriate OOBM interfaces as shown in the example below.

RP/0/0/CPU0:R2(config)#int g0/0/0/2

RP/0/0/CPU0:R2(config-if)#ipv4 access-group TRAFFIC_FROM_NOC in

RP/0/0/CPU0:R2(config)#int g0/0/0/3

RP/0/0/CPU0:R2(config-if)#access-group TRAFFIC_TO_NOC in

RP/0/0/CPU0:R2(config-if)#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

1.46 CISC-RT-000450 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216773
RULE ID: SV-216773r531087
```

Rationale:

The OOBM access switch 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:

This requirement is only applicable where management access to the router is via an OOBM interface which is not a true OOBM interface.

Step 1: Verify that the managed interface has an inbound and outbound ACL configured.

interface MgmtEth0/0/CPU0/0

ipv4 address 10.1.13.4 255.255.255.0

ipv4 access-group INGRESS_MANAGEMENT_ACL ingress

ipv4 access-group EGRESS_MANAGEMENT_ACL egress

Step 2: Verify that the ingress ACL only allows management and ICMP traffic.

ipv4 access-list INGRESS_MANAGEMENT_ACL

10 permit tcp any host 10.11.1.22 eq tacacs

20 permit tcp any host 10.11.1.22 eq ssh

30 permit udp any host 10.11.1.22 eq snmp

40 permit udp any host 10.11.1.22 eq snmptrap

50 permit udp any host 10.11.1.22 eq ntp

60 permit icmp any host 10.11.1.22

70 deny ipv4 any any log-input

Step 3: Verify that the egress ACL blocks any transit traffic.

ipv4 access-list EGRESS_MANAGEMENT_ACL

10 deny ipv4 any any log-input

Note: On Cisco routers, local generated packets are not inspected by outgoing interface access-lists. Hence, the above configuration would simply drop any packets not generated by the router; hence, blocking any transit traffic.

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

Remediation:

If the management interface is not a dedicated OOBM interface, it must be configured with both an ingress and egress ACL.

Step 1: Configure an ingress ACL a shown in the example below.

RP/0/0/CPU0:R2(config)#Ipv4 access-list INGRESS_MANAGEMENT_ACL

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any host 10.11.1.22 eq tacacs

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any host 10.11.1.22 eq 22

RP/0/0/CPU0:R2(config-ipv4-acl)#permit udp any host 10.11.1.22 eq snmp

RP/0/0/CPU0:R2(config-ipv4-acl)#permit udp any host 10.11.1.22 eq snmptrap

RP/0/0/CPU0:R2(config-ipv4-acl)#permit udp any host 10.11.1.22 eq ntp

RP/0/0/CPU0:R2(config-ipv4-acl)#permit icmp any host 10.11.1.22

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

Step 2: Configure an egress ACL a shown in the example below.

RP/0/0/CPU0:R2(config)#Ipv4 access-list EGRESS_MANAGEMENT_ACL

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log-input

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

Step 3: Apply the ACLs to the OOBM interfaces.

RP/0/0/CPU0:R2(config)#int MgmtEth0/0/CPU0/0

RP/0/0/CPU0:R2(config-if)#ipv4 access-group INGRESS_MANAGEMENT_ACL ingress

RP/0/0/CPU0:R2(config-if)#ipv4 access-group EGRESS_MANAGEMENT_ACL egress

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

1.47 CISC-RT-000460 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco router providing connectivity to the Network Operations Center (NOC) must be configured to forward all in-band management traffic via an IPsec tunnel.

```
GROUP ID: V-216774
RULE ID: SV-216774r531087
```

Rationale:

When the production network is managed in-band, the management network could be housed at a NOC that is located remotely at single or multiple interconnected sites. NOC interconnectivity, as well as connectivity between the NOC and the managed network, must be enabled using IPsec tunnels to provide the separation and integrity of the managed traffic.

Audit:

This requirement is not applicable for the DODIN Backbone.

Verify that all traffic from the managed network to the management network or NOC and vice-versa is secured via IPsec tunnel.

Step 1: Note the profile referenced for the IPSec tunnel to the NOC.

interface tunnel-ipsec 30

profile IPSEC_NOC_PROFILE

tunnel source GigabitEthernet0/0/0/2

tunnel destination x.1.22.2

Step 2: Note the crypto ACL that was specified in the IPSec profile.

crypto isakmp keyring ISAKMP_KEYRING

pre-shared-key address x.1.22.2 255.255.255.255 key encrypted 150A13141C32

!

crypto isakmp policy 10

hash sha256

encryption aes 256

authentication pre-share

!

crypto ipsec transform-set IPSEC_TRANS esp-aes 256 esp-sha256-hmac

mode tunnel

!

crypto ipsec profile IPSEC_NOC_PROFILE

set pfs group16

match address MGMT_TRAFFIC_ACL

Step 3: Review the ACL defined in the crypto map and verify that the destination is the management network.

!

ipv4 access-list MGMT_TRAFFIC_ACL

10 permit ipv4 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255

!

Note: The management network is this example is 10.22.2.0/24.

If the management traffic is not secured via IPsec tunnel, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.

Ensure that all traffic from the managed network to the management network is secured via IPsec tunnel as shown in the configuration examples below.

Step 1: Configure the ACL for the management network as the destination. This ACL will be defined in the crypto as the interesting traffic to be forwarded into the IPsec tunnel.

RP/0/0/CPU0:R3(config)#Ipv4 access-list MGMT_TRAFFIC_ACL

RP/0/0/CPU0:R3(config-ipv4-acl)#permit ip 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255

Step 2: Create an ISAKMP policy for Phase 1 negotiations.

RP/0/0/CPU0:R3(config)#crypto isakmp policy 10

RP/0/0/CPU0:R3(config-isakmp-policy)#authentication pre-share

RP/0/0/CPU0:R3(config-isakmp-policy)#hash sha256

RP/0/0/CPU0:R3(config-isakmp-policy)#encryption aes 256

RP/0/0/CPU0:R3(config-isakmp-policy)#exit

Step 3: Configure the keyring to be used for ISAKMP to authenticate the remote side during IKE negotiation.

RP/0/0/CPU0:R3(config)#crypto isakmp keyring ISAKMP_KEYRING

RP/0/0/CPU0:R3(config-crypto-keyring)#pre-shared-key address 255.255.255.255 key xxxxx

RP/0/0/CPU0:R3(config-crypto-keyring)#exit

Step 4: Configure the IPSec transform set.

RP/0/0/CPU0:R3(config)#crypto ipsec transform-set IPSEC_TRANS esp-aes 256 esp-sha256-hmac

RP/0/0/CPU0:R3(config-transform-set IPSEC_TRANS)#mode tunnel

Step 5: Configure the IPSec profile.

RP/0/0/CPU0:R3(config)#crypto ipsec profile IPSEC_NOC_PROFILE

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#set pfs group 16

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#match MGMT_TRAFFIC_ACL transform-set IPSEC_TRANS

RP/0/0/CPU0:R3(config- IPSEC_NOC_PROFILE)#exit

Step 6: Configure the IPSec virtual interface.

RP/0/0/CPU0:R3(config)#interface tunnel-ipsec 22

RP/0/0/CPU0:R3(config-if)#profile IPSEC_NOC_PROFILE

RP/0/0/CPU0:R3(config-if)#tunnel source GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-if)#tunnel destination x.1.22.2

RP/0/0/CPU0:R3(config-if)#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

1.48 CISC-RT-000470 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco BGP router must be configured to enable the Generalized TTL Security Mechanism (GTSM).

```
GROUP ID: V-217007
RULE ID: SV-217007r856462
```

Rationale:

As described in RFC 3682, 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 BGP configuration to verify that TTL security has been configured for each external neighbor as shown in the example below.

router bgp n

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as n

ttl-security

address-family ipv4 unicast

!

!

!

If the router is not configured to use GTSM for all Exterior Border Gateway Protocol peering sessions, this is a finding.

Remediation:

Configure TTL security on all external BGP neighbors as shown in the example below.

RP/0/0/CPU0:R2(config)#router bgp n

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.23.3

RP/0/0/CPU0:R2(config-bgp-nbr)#ttl-security

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 CISC-RT-000480 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-217008
RULE ID: SV-217008r945862
```

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:

Review the BGP configuration to determine if it is peering with multiple autonomous systems. Interview the ISSM and router administrator to determine if unique keys are being used.

router bgp n

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as y

keychain YYY_KEY_CHAIN

ttl-security

address-family ipv4 unicast

!

!

neighbor x.1.24.4

remote-as z

keychain ZZZ_KEY_CHAIN

ttl-security

address-family ipv4 unicast

!

!

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

Remediation:

Configure the router to use unique keys for each AS that it peers with as shown in the example below.

RP/0/0/CPU0:R2(config)#router bgp n

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.23.3

RP/0/0/CPU0:R2(config-bgp-nbr)#keychain YYY_KEY_CHAIN

RP/0/0/CPU0:R2(config-bgp-nbr)#neighbor x.1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#keychain ZZZ_KEY_CHAIN

RP/0/0/CPU0:R2(config-bgp-nbr)#end

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

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)

1.50 CISC-RT-000490 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216777
RULE ID: SV-216777r531087
```

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 router configuration to verify that it will reject BGP routes for any Bogon prefixes.

Step 1: verify that an inbound route policy has been configured for each external neighbor as shown in the example below.

router bgp n

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as y

keychain YYY_KEY_CHAIN

ttl-security

address-family ipv4 unicast

route-policy BGP_FILTER in

!

!

neighbor x.1.24.4

remote-as z

keychain ZZZ_KEY_CHAIN

ttl-security

address-family ipv4 unicast

route-policy BGP_FILTER in

!

!

Step 2: Review the route policy to determine if it is filtering at a minimum BOGON prefixes as shown in the example below.

route-policy BGP_FILTER

if destination in BOGON_PREFIXES then

drop

else

pass

endif

end-policy

Step 3: Review the prefix set referenced in the route policy above has been configured containing the current Bogon prefixes as shown in the example below.

prefix-set BOGON_PREFIXES

0.0.0.0/8 le 32,

10.0.0.0/8 le 32,

100.64.0.0/10 le 32,

127.0.0.0/8 le 32,

169.254.0.0/16 le 32,

172.16.0.0/12 le 32,

192.0.2.0/24 le 32,

192.88.99.0/24 le 32,

192.168.0.0/16 le 32,

198.18.0.0/15 le 32,

198.51.100.0/24 le 32,

203.0.113.0/24 le 32,

240.0.0.0/4 le 32,

224.0.0.0/4 le 32

end-set

If the router is not configured to reject inbound route advertisements for any Bogon prefixes, this is a finding.

Remediation:

Configure the router to reject inbound route advertisements for any Bogon prefixes.

Step 1: Configure a prefix set containing the current Bogon prefixes as shown below.

RP/0/0/CPU0:R2(config)#prefix-set BOGON_PREFIXES

RP/0/0/CPU0:R2(config-pfx)#0.0.0.0/8 le 32,

RP/0/0/CPU0:R2(config-pfx)#10.0.0.0/8 le 32,

RP/0/0/CPU0:R2(config-pfx)#100.64.0.0/10 le 32,

RP/0/0/CPU0:R2(config-pfx)#127.0.0.0/8 le 32,

RP/0/0/CPU0:R2(config-pfx)#169.254.0.0/16 le 32,

RP/0/0/CPU0:R2(config-pfx)#172.16.0.0/12 le 32,

RP/0/0/CPU0:R2(config-pfx)#192.0.2.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#192.88.99.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#192.168.0.0/16 le 32,

RP/0/0/CPU0:R2(config-pfx)#198.18.0.0/15 le 32,

RP/0/0/CPU0:R2(config-pfx)#198.51.100.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#203.0.113.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#240.0.0.0/4 le 32,

RP/0/0/CPU0:R2(config-pfx)#224.0.0.0/4 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 2: Configure the route policy to drop routes with BOGON prefixes as shown in the example below.

RP/0/0/CPU0:R2(config)#route-policy BGP_FILTER

RP/0/0/CPU0:R2(config-rpl)#if destination in BOGON_PREFIXES then

RP/0/0/CPU0:R2(config-rpl-if)#drop

RP/0/0/CPU0:R2(config-rpl-if)#else pass endif

RRP/0/0/CPU0:R2(config-rpl)#end-policy

RP/0/0/CPU0:R2(config)#exit

Step 3: Apply the route policy to each external BGP neighbor as shown in the example.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.23.3

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_FILTER in

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_FILTER 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

1.51 CISC-RT-000500 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216778
RULE ID: SV-216778r531087
```

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

Step 1: verify that an inbound route policy has been configured for each external neighbor as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as yy

keychain YYY_KEY_CHAIN

ttl-security

address-family ipv4 unicast

route-policy BGP_FILTER_INBOUND in

!

!

neighbor x.1.24.4

remote-as zz

keychain ZZZ_KEY_CHAIN

ttl-security

address-family ipv4 unicast

route-policy BGP_FILTER_INBOUND in

!

!

Step 2: Review the route policy to determine if it is filtering at a minimum local prefixes as shown in the example below.

route-policy BGP_FILTER_INBOUND

if destination in LOCAL_PREFIX then

drop

else

pass

endif

end-policy

Note: If bogons are also filtered per previous requirement, the route policy would look similar to the following example:

route-policy BGP_FILTER_INBOUND

if destination in BOGON_PREFIXES then

drop

elseif destination in LOCAL_PREFIX then

drop

else

pass

endif

end-policy

Step 3: Review the prefix set referenced in the route policy above to determine if it includes the local global prefix as shown in the example below.

prefix-set LOCAL_PREFIX

x.13.1.0/24 le 32

end-set

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

Remediation:

Step 1: Configure a prefix set containing the current Bogon prefixes as shown below.

RP/0/0/CPU0:R2(config)#prefix-set Step 1: Configure a prefix set containing the current Bogon prefixes as shown below.

RP/0/0/CPU0:R2(config)#prefix-set LOCAL_PREFIX

RP/0/0/CPU0:R2(config-pfx)#x.13.1.0/24 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 2: Configure the route policy to drop routes with BOGON prefixes as shown in the example below.

RP/0/0/CPU0:R2(config)#route-policy BGP_FILTER_INBOUND

RP/0/0/CPU0:R2(config-rpl)#if destination in LOCAL_PREFIX then

RP/0/0/CPU0:R2(config-rpl-if)#drop

RP/0/0/CPU0:R2(config-rpl-if)#else pass endif

RRP/0/0/CPU0:R2(config-rpl)#end-policy

RP/0/0/CPU0:R2(config)#exit

Step 3: Apply the route policy to each external BGP neighbor as shown in the example.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.23.3

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_FILTER_INBOUND in

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_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

1.52 CISC-RT-000510 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216779
RULE ID: SV-216779r531087
```

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 router configuration to verify that there are ACLs defined to only accept routes for prefixes that belong to specific customers.

Step 1: Verify that an inbound route policy has been configured for each customer neighbor as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.12.4.14

remote-as 64514

address-family ipv4 unicast

route-policy CUST1_PREFIX_FILTER in

!

!

neighbor x.12.4.16

remote-as 64516

address-family ipv4 unicast

route-policy CUST2_PREFIX_FILTER in

!

!

Step 2: Review the route policies to determine if it is accepting only prefixes belonging to each customer as shown in the example below.

route-policy CUST1_PREFIX_FILTER

if destination in CUST1_PREFIX then

pass

else

drop

endif

end-policy

!

route-policy CUST2_PREFIX_FILTER

if destination in CUST2_PREFIX then

pass

else

drop

endif

end-policy

Step 3: Review the prefix sets referenced in the route policies above to determine if they include only prefixes belonging to each customer.

prefix-set CUST1_PREFIX

x.1.1.0/24 le 32

end-set

!

prefix-set CUST2_PREFIX

x.2.1.0/24 le 32

end-set

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.

If the 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.

Remediation:

Configure the router to reject inbound route advertisements from each CE router for prefixes that are not allocated to that customer.

Step 1: Configure a prefix set for each customer containing prefixes belonging to each as shown in the example.

RP/0/0/CPU0:R2(config)#prefix-set CUST1_PREFIX

RP/0/0/CPU0:R2(config-pfx)#x.1.1.0/24 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

RP/0/0/CPU0:R2(config)#prefix-set CUST2_PREFIX

RP/0/0/CPU0:R2(config-pfx)#x.2.1.0/24 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 2: Configure a route policy filter for each customer as shown in the example.

RP/0/0/CPU0:R2(config)#route-policy CUST1_PREFIX_FILTER

RP/0/0/CPU0:R2(config-rpl)#if destination in CUST1_PREFIX then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

RP/0/0/CPU0:R2(config)#route-policy CUST2_PREFIX_FILTER

RP/0/0/CPU0:R2(config-rpl)#if destination in CUST2_PREFIX then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

Step 3: Apply the route policy to each customer neighbor as shown in the example.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.14

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy CUST1_PREFIX_FILTER in

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.16

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy CUST2_PREFIX_FILTER 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

1.53 CISC-RT-000520 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216780
RULE ID: SV-216780r531087
```

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.

Step 1: verify that an outbound route policy has been configured for each customer neighbor as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.12.4.14

remote-as 64514

address-family ipv4 unicast

route-policy CE_ADVERTISEMENTS out

!

!

neighbor x.12.4.16

remote-as 64516

address-family ipv4 unicast

route-policy CE_ADVERTISEMENTS out

!

!

Step 2: Review the route policy to determine if it is accepting only prefixes belonging to customers or the local autonomous system as shown in the example below.

route-policy CE_ADVERTISEMENTS

if destination in CE_PREFIX_ADVERTISEMENTS then

pass

else

drop

endif

end-policy

Step 3: Review the prefix sets referenced in the route policy above to determine if they include only prefixes belonging to customers or the local autonomous system as shown in the example below.

prefix-set CE_PREFIX_ADVERTISEMENTS

x.13.1.0/24 le 32,

x.13.2.0/24 le 32,

x.13.3.0/24 le 32,

x.13.4.0/24 le 32

end-set

If the 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.

Step 1: Configure a prefix set for customer and local autonomous system prefixes as shown in the example.

RP/0/0/CPU0:R2(config)#prefix-set CE_PREFIX_ADVERTISEMENTS

RP/0/0/CPU0:R2(config-pfx)#x.13.1.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#x.13.2.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#x.13.3.0/24 le 32,

RP/0/0/CPU0:R2(config-pfx)#x.13.4.0/24 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 2: Configure a route policy filter for allow customer and local autonomous system prefixes as shown in the example.

RP/0/0/CPU0:R2(config)#route-policy CE_ADVERTISEMENTS

RP/0/0/CPU0:R2(config-rpl)#if destination in CE_PREFIX_ADVERTISEMENTS then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

Step 3: Apply the route policy to each customer neighbor as shown in the example.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.14

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy route-policy CE_ADVERTISEMENTS out

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.16

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy CE_ADVERTISEMENTS out

RP/0/0/CPU0:R2(config-bgp-nbr-af)#end

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

1.54 CISC-RT-000530 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216781
RULE ID: SV-216781r531087
```

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:

Step 1: verify that an outbound route policy has been configured for each external neighbor as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as yy

address-family ipv4 unicast

route-policy BGP_FILTER_OUTBOUND out

!

!

neighbor x.1.24.4

remote-as zz

address-family ipv4 unicast

route-policy BGP_FILTER_OUTBOUND out

!

!

Step 2: Review the route policy to determine if it is filtering at a minimum IP core prefixes as shown in the example below.

route-policy BGP_FILTER_OUTBOUND

if destination in CORE_PREFIX then

drop

else

pass

endif

end-policy

Step 3: Review the prefix set referenced in the route policy above to determine if it includes the IP core prefix as shown in the example below.

prefix-set CORE_PREFIX

10.1.1.0/24 le 32

end-set

If the router is not configured to reject outbound route advertisements for prefixes belonging to the IP core, this is a finding.

Remediation:

Step 1: Configure a prefix set containing the IP core prefix as shown below.

RP/0/0/CPU0:R2(config)#prefix-set

Step 2: Configure a prefix set containing the current Bogon prefixes as shown below.

RP/0/0/CPU0:R2(config)#prefix-set CORE_PREFIX

RP/0/0/CPU0:R2(config-pfx)#10.1.1.0/24 le 32

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 3: Configure the route policy to drop route advertisements for IP core prefixes as shown in the example below.

RP/0/0/CPU0:R2(config)#route-policy BGP_FILTER_OUTBOUND

RP/0/0/CPU0:R2(config-rpl)#if destination in CORE_PREFIX then

RP/0/0/CPU0:R2(config-rpl-if)#drop

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#pass

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

Step 4: Apply the route policy to each external BGP neighbor as shown in the example.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.23.3

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_FILTER_OUTBOUND out

RP/0/0/CPU0:R2(config-bgp)#neighbor x.1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy BGP_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

1.55 CISC-RT-000540 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216782
RULE ID: SV-216782r945854
```

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:

Review the router configuration to verify the router is configured to deny updates received from eBGP peers that do not list their AS number as the first AS in the AS_PATH attribute.

By default Cisco IOS enforces the first AS in the AS_PATH attribute for all route advertisements. Review the router configuration to verify that the command bgp enforce-first-as disable is not configured as shown in the example below.

router bgp nn

bgp enforce-first-as disable

If the router is not configured to reject updates from peers that do not list their AS number as the first AS in the AS_PATH attribute, this is a finding.

Remediation:

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

RP/0/0/CPU0:R2(config)#router bgp 2

RP/0/0/CPU0:R2(config-bgp)#no bgp enforce-first-as disable

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)

1.56 CISC-RT-000550 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216783
RULE ID: SV-216783r945855
```

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 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: verify that an inbound route policy has been configured for each customer neighbor as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.12.4.14

remote-as 64514

address-family ipv4 unicast

route-policy FILTER _64514_ROUTES in

!

!

neighbor x.12.4.16

remote-as 64516

address-family ipv4 unicast

route-policy FILTER_64516_ROUTES in

!

Step 2: Verify that the route policy permits only routes from each CE router with an originating AS that does not belong to that customer.

route-policy FILTER_64514_ROUTES

if as-path originates-from 64514' then

pass

else

drop

endif

end-policy

!

route-policy FILTER_64516_ROUTES

if as-path originates-from 64516' then

pass

else

drop

endif

end-policy

Note: The inbound route policy to filter customer prefixes can be nested with the above route policy as shown in the example below.

route-policy CUST1_INBOUND_FILTER

apply CUST1_FILTER

apply FILTER_64514_ROUTES

end-policy

If the 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:

This requirement is not applicable for the DODIN Backbone.

Configure the 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 an inbound route policy for each customer to only accept routes with an originating AS that belongs to that customer as shown in the example below.

RP/0/0/CPU0:R2(config)#route-policy FILTER_64514_ROUTES

RP/0/0/CPU0:R2(config-rpl)#if as-path originates-from '64514' then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

RP/0/0/CPU0:R2(config)#route-policy FILTER_64516_ROUTES

RP/0/0/CPU0:R2(config-rpl)#if as-path originates-from '64516' then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

Step 2: Apply the appropriate inbound route policy with each peering CE router as shown in the example below.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.14

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy route-policy FILTER_64514_ROUTES in

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.16

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy FILTER_64516_ROUTES in

RP/0/0/CPU0:R2(config-bgp-nbr-af)#end

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)

1.57 CISC-RT-000560 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216784
RULE ID: SV-216784r856445
```

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 router configuration to verify that the number of received prefixes from each eBGP neighbor is controlled.

router bgp xx

address-family ipv4 unicast

!

neighbor x.1.23.3

remote-as yy

ttl-security

address-family ipv4 unicast

route-policy BGP_FILTER in

maximum-prefix nnnn 75 discard-extra-paths

!

!

neighbor x.1.24.4

remote-as zz

address-family ipv4 unicast

route-policy BGP_FILTER in

maximum-prefix nnnn 75 discard-extra-paths

!

!

!

If the router is not configured to control the number of prefixes received from each peer to protect against route table flooding and prefix de-aggregation attacks, this is a finding.

Remediation:

Configure the router to use the maximum prefixes feature to protect against route table flooding and prefix de-aggregation attacks as shown in the example below.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#maximum-prefix 444 discard-extra-paths

RP/0/0/CPU0:R2(config-bgp-nbr-af)#end

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

1.58 CISC-RT-000570 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216785
RULE ID: SV-216785r856446
```

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 router configuration to determine if it is compliant with this requirement.

Step 1: Verify that an inbound route policy has been configured for each CE router as shown in the example below.

router bgp xx

address-family ipv4 unicast

!

neighbor x.12.4.14

remote-as 64514

address-family ipv4 unicast

route-policy FILTER_LONG_PREFIXES in

!

!

neighbor x.12.4.16

remote-as 64516

address-family ipv4 unicast

route-policy FILTER_LONG_PREFIXES in

!

Step 2: Verify that the route policy permits only routes from each CE router with a prefix length of 24 or shorter or the least significant prefixes issued to the customer as shown in the example below.

route-policy FILTER_LONG_PREFIXES

if destination in PREFIX_LENGTH then

pass

else

drop

endif

end-policy

Note: The inbound route policy to filter customer prefixes can be nested with the above route policy as shown in the example below.

route-policy CUST1_INBOUND_FILTER

apply CUST1_FILTER

apply FILTER_64514_ROUTES

apply FILTER_LONG_PREFIXES

end-policy

Step 3: Review the prefix set referenced in the route policy above to determine if it only allows a prefix length 24 or shorter.

prefix-set PREFIX_LENGTH

0.0.0.0/0 ge 8 le 24

end-set

If the 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.

Configure the router to limit the prefix size on any route advertisement to /24 or the least significant prefixes issued to the customer.

Step 1: Configure a prefix set to not include prefixes are longer than /24.

RP/0/0/CPU0:R2(config)#prefix-set PREFIX_LENGTH

RP/0/0/CPU0:R2(config-pfx)#0.0.0.0/0 ge 8 le 24

RP/0/0/CPU0:R2(config-pfx)#end-set

Step 2: Configure a route policy to only accept prefixes that are /24 or shorter as shown in the example below.

RP/0/0/CPU0:R2(config)#route-policy FILTER_LONG_PREFIXES

RP/0/0/CPU0:R2(config-rpl)#if destination in PREFIX_LENGTH then

RP/0/0/CPU0:R2(config-rpl-if)#pass

RP/0/0/CPU0:R2(config-rpl-if)#else

RP/0/0/CPU0:R2(config-rpl-else)#drop

RP/0/0/CPU0:R2(config-rpl-else)#endif

RP/0/0/CPU0:R2(config-rpl)#end-policy

Step 3: Apply the route policy above inbound with each peering CE router as shown in the example below.

RP/0/0/CPU0:R2(config)#router bgp xx

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.14

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy route-policy FILTER_LONG_PREFIXES in

RP/0/0/CPU0:R2(config-bgp)#neighbor x.12.4.16

RP/0/0/CPU0:R2(config-bgp-nbr)#address-family ipv4 unicast

RP/0/0/CPU0:R2(config-bgp-nbr-af)#route-policy FILTER_LONG_PREFIXES in

RP/0/0/CPU0:R2(config-bgp-nbr-af)#end

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

1.59 CISC-RT-000580 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco BGP router must be configured to use its loopback address as the source address for iBGP peering sessions.

```
GROUP ID: V-216786
RULE ID: SV-216786r991935
```

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:

Step 1: Review the router configuration to verify that a loopback address has been configured.

interface Loopback0

ip address 10.1.1.1 255.255.255.255

Step 2: Verify that the loopback interface is used as the source address for all iBGP sessions.

router bgp xx

address-family ipv4 unicast

!

neighbor 10.1.23.3

remote-as xx

update-source Loopback0

If the router does not use its loopback address as the source address for all iBGP sessions, this is a finding.

Remediation:

Configure the router to use its loopback address as the source address for all iBGP peering.

RP/0/0/CPU0:R2(config)#router bgp 2

RP/0/0/CPU0:R2(config-bgp)#neighbor 10.1.24.4

RP/0/0/CPU0:R2(config-bgp-nbr)#update-source lo0

RP/0/0/CPU0:R2(config-bgp-nbr)#end

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

1.60 CISC-RT-000590 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco MPLS router must be configured to use its loopback address as the source address for LDP peering sessions.

```
GROUP ID: V-216787
RULE ID: SV-216787r991936
```

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 router configuration to determine if it is compliant with this requirement.

Verify that a loopback address has been configured as shown in the following example:

interface Loopback0

ip address 10.1.1.1 255.255.255.255

By default, routers will use its loopback address for LDP peering. If an address has not be configured on the loopback interface, it will use its physical interface connecting to the LDP peer. If the router-id command is specified that overrides this default behavior, verify that it is the IP address of the designated loopback interface as shown in the example below.

mpls ldp

router-id 10.1.1.1

If the router is not configured do use its loopback address for LDP peering, this is a finding.

Remediation:

Configure the router to use their loopback address as the source address for LDP peering sessions. As noted in the check content, the default behavior is to use its loopback address.

RP/0/0/CPU0:R3(config)#mpls ldp router-id 10.1.1.1

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

1.61 CISC-RT-000600 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216788
RULE ID: SV-216788r531087
```

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:

OSPF Example

router ospf 1

mpls ldp sync

IS-IS Example

router isis 1

net 49.0001.1234.1600.5531.00

interface GigabitEthernet0/0/0/1

address-family ipv4 unicast

mpls ldp sync

If the router is not configured to synchronize IGP and LDP, this is a finding.

Remediation:

Configure the MPLS router to synchronize IGP and LDP, minimizing packet loss when an IGP adjacency is established prior to LDP peers completing label exchange.

OSPF Example

RP/0/0/CPU0:R3(config)#router ospf 1

RP/0/0/CPU0:R3(config-ospf)#mpls ldp sync

RP/0/0/CPU0:R3(config-ospf)#end

IS-IS Example

RP/0/0/CPU0:R3(config)#router isis 1

RP/0/0/CPU0:R3(config-isis)#interface g0/0/0/1

RP/0/0/CPU0:R3(config-isis-if)#address-family ipv4 unicast

RP/0/0/CPU0:R3(config-isis-if-af)#mpls ldp sync

RP/0/0/CPU0:R3(config-isis-if-af)#end

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

1.62 CISC-RT-000610 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The MPLS router with RSVP-TE enabled must be configured with message pacing to adjust maximum burst and 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-216789
RULE ID: SV-216789r531087
```

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:

Review the router configuration to determine RSVP messages are rate limited.

Step 1: Determine if MPLS TE is enabled on any interface as shown in the example below.

mpls traffic-eng

interface GigabitEthernet0/0/0/1

Step 2: If MPLS TE is enabled, verify that RSVP messages are rate limited on each interface. The example allows 50 messages per 500 milliseconds.

rsvp

interface GigabitEthernet0/0/0/1

signaling rate-limit rate 50 interval 500

Note: The command rsvp msg-pacing has been deprecated by the command ip rsvp signaling rate-limit command.

If the router with RSVP-TE enabled does not have message pacing configured based on the link speed and input queue size of adjacent core routers, this is a finding.

Remediation:

Configure the router to rate limit RSVP messages as shown in the example.

RP/0/0/CPU0:R3(config)#rsvp interface g0/0/0/1

RP/0/0/CPU0:R3(config-rsvp-if)#signaling rate-limit rate 50 interval 500

RP/0/0/CPU0:R3(config-rsvp-if)#end

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)

1.63 CISC-RT-000620 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco MPLS router must be configured to have TTL Propagation disabled.

```
GROUP ID: V-216790
RULE ID: SV-216790r531087
```

Rationale:

The head end of the label-switched 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-switched 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 switched, 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 router configuration to verify that TTL propagation is disabled as shown in the example below.

mpls ip-ttl-propagate disable

If the MPLS router is not configured to disable TTL propagation, this is a finding.

Remediation:

Configure the MPLS router to disable TTL propagation as shown in the example below.

RP/0/0/CPU0:R3(config)#mpls ip-ttl-propagate disable

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

1.64 CISC-RT-000630 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco 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-216791
RULE ID: SV-216791r991937
```

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:

Step 1: Review the design plan for deploying L3VPN and VRF-lite.

Step 2: Review the design plan for deploying L3VPN and VRF-lite. Review all CE-facing interfaces and verify that the proper VRF is defined via the ip vrf forwarding command. In the example below, COI1 is bound to interface GigabitEthernet0/0/0/1, while COI2 is bound to GigabitEthernet0/0/0/2.

interface GigabitEthernet0/0/0/1

description link to COI1

vrf COI1

ipv4 address x.1.34.12 255.255.255.252

!

interface GigabitEthernet0/0/0/2

description link to COI2

vrf COI2

ipv4 address x.1.22.12 255.255.255.252

If any VRFs are not bound to the appropriate physical or logical interface, this is a finding.

Remediation:

Configure the PE router to have each VRF bound to the appropriate physical or logical interfaces to maintain traffic separation between all MPLS L3VPNs.

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

1.65 CISC-RT-000640 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco PE router must be configured to have each Virtual Routing and Forwarding (VRF) instance with the appropriate Route Target (RT).

```
GROUP ID: V-216792
RULE ID: SV-216792r991939
```

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:

Review the design plan for MPLS/L3VPN and VRF-lite to determine what RTs have been assigned for each VRF. Review the router configuration and verify that the correct RT is configured for each VRF. In the example below, route target 13:13 has been configured for COI1.

vrf COI1

address-family ipv4 unicast

import route-target

13:13

!

export route-target

13:13

!

!

!

If there are VRFs configured with the wrong RT, this is a finding.

Remediation:

Configure the router to have each VRF instance defined with the correct RT.

RP/0/0/CPU0:R3(config)#vrf COI1

RP/0/0/CPU0:R3(config-vrf)#address-family ipv4 unicast

RP/0/0/CPU0:R3(config-vrf-af)#import route-target 13:13

RP/0/0/CPU0:R3(config-vrf-af)#export route-target 13:13

RP/0/0/CPU0:R3(config-vrf-af)#end

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

1.66 CISC-RT-000650 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco PE router must be configured to have each VRF with the appropriate Route Distinguisher (RD).

```
GROUP ID: V-216793
RULE ID: SV-216793r991940
```

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 design plan for MPLS/L3VPN to determine what RD have been assigned for each VRF. Review the router configuration and verify that the correct RD is configured for each VRF. In the example below, route distinguisher 13:13 has been configured for COI1.

router bgp nn

address-family ipv4 unicast

!

address-family vpnv4 unicast

!

…

…

…

!

vrf COI1

rd 13:13

address-family ipv4 unicast

redistribute ospf 1

!

!

!

Note: This requirement is only applicable for MPLS VPN implementations.

If the wrong RD has been configured for any VRF, this is a finding.

Remediation:

Configure the correct RD for each VRF.

RP/0/0/CPU0:R3(config)#router bgp nn

RP/0/0/CPU0:R3(config-bgp)#vrf COI1

RP/0/0/CPU0:R3(config-bgp-vrf)#rd 13:13

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 CISC-RT-000660 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216794
RULE ID: SV-216794r878124
```

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:

The Cisco router is not compliant with this requirement; hence, it is a finding. However, the severity level can be downgraded to a category 3 if the router is configured to authenticate targeted LDP sessions using MD5 as shown in the configuration example below.

mpls ldp

router-id 10.1.1.2

neighbor 10.1.1.1

password encrypted xxxxxxxxxxxxxxx

neighbor 10.1.2.1

password encrypted xxxxxxxxxxxxxxx

If the router is not configured to authenticate targeted LDP sessions using MD5, the finding will remain as a CAT II.

Remediation:

The severity level can be downgraded to a category 3 if the router is configured to authenticate targeted LDP sessions using MD5 as shown in the example below.

RP/0/0/CPU0:R3(config)#mpls ldp

RP/0/0/CPU0:R3(config-ldp)#neighbor 10.1.1.1

RP/0/0/CPU0:R3(config-ldp)#neighbor password clear xxxxxxxx

RP/0/0/CPU0:R3(config-ldp)#neighbor 10.1.2.1

RP/0/0/CPU0:R3(config-ldp)#neighbor password clear xxxxxxxx

RP/0/0/CPU0:R3(config-ldp)#commit

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

1.68 CISC-RT-000670 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco PE router providing MPLS Virtual Private Wire Service (VPWS) must be configured to have the appropriate pseudowire ID for each attachment circuit.

```
GROUP ID: V-216795
RULE ID: SV-216795r991942
```

Rationale:

VPWS is an L2VPN technology that provides a virtual circuit (aka pseudowire) between two PE routers to forward Layer 2 frames between two customer-edge routers or switches 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 label bindings for the specified pseudowire ID. The pseudowire 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 pseudowire ID is configured for each attachment circuit.

Audit:

Verify that the correct pseudowire ID has been configured for the appropriate attachment circuit. In the example below GigabitEthernet0/0/0/1 is the CE-facing interface that is configured for VPWS.

l2vpn

pw-class ETHOM

encapsulation mpls

!

!

xconnect group COI1

p2p COI1-S1-S2

interface GigabitEthernet0/0/0/1

neighbor ipv4 10.1.12.4 pw-id 55

pw-class ETHOM

!

!

!

!

If the correct pseudowire ID has not been configured on both routers, this is a finding.

Remediation:

Assign globally unique pseudowire IDs for each virtual circuit and configure the attachment circuits with the appropriate pseudowire ID.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#xconnect group COI1

RP/0/0/CPU0:R3(config-l2vpn-xc)#p2p COI1-S1-S2

RP/0/0/CPU0:R3(config-l2vpn-xc-p2p)#interface g0/0/0/1

RP/0/0/CPU0:R3(config-l2vpn-xc-p2p)#neighbor 10.1.12.4 pw-id 55

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 CISC-RT-000680 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco PE router providing Virtual Private LAN Services (VPLS) must be configured to have all attachment circuits defined to the virtual forwarding instance (VFI) with the globally unique VPN ID assigned for each customer VLAN.

```
GROUP ID: V-216796
RULE ID: SV-216796r531087
```

Rationale:

VPLS defines an architecture that delivers Ethernet multipoint services over an MPLS network. Customer Layer 2 frames are forwarded across the MPLS core via pseudowires using IEEE 802.1q Ethernet bridging principles. A pseudowire is a virtual bidirectional connection between two attachment circuits (virtual connections between PE and CE routers). A pseudowire contains two unidirectional label-switched paths (LSP) between two PE routers. Each MAC virtual forwarding table instance (VFI) is interconnected using pseudowires provisioned for the bridge domain, thereby maintaining privacy and logical separation between each VPLS bridge domain.

The VFI specifies the pseudowires associated with connecting PE routers and the customer-facing attachment circuits belonging to a given VLAN. Resembling a Layer 2 switch, the VFI is responsible for learning MAC addresses and providing loop-free forwarding of customer traffic to the appropriate end nodes. Each VPLS domain is identified by a globally unique VPN ID; hence, VFIs of the same VPLS domain must be configured with the same VPN ID on all participating PE routers. To guarantee traffic separation for all customer VLANs and that all packets are forwarded to the correct destination, it is imperative that the correct attachment circuits are associated with the appropriate VFI and that each VFI is associated to the unique VPN ID assigned to the customer VLAN.

Audit:

Review the implementation plan and the VPN IDs assigned to customer VLANs for the VPLS deployment.

Review the PE router configuration to verify that customer attachment circuits are associated to the appropriate VFI. In the example below, the attached circuit at interface GigabitEthernet0/0/0/2 is associated to VPN ID 110.

interface GigabitEthernet0/0/0/2

l2transport

…

…

…

l2vpn

pw-class ETH_O_MPLS

encapsulation mpls

transport-mode ethernet

!

!

bridge group L2GROUP

bridge-domain L2_BRIDGE_COI1

interface GigabitEthernet0/0/0/2

!

vfi VFI_COI1

vpn-id 101

neighbor 10.1.1.1 pw-id 55

pw-class ETH_O_MPLS

!

neighbor 10.2.2.2 pw-id 55

pw-class ETH_O_MPLS

!

!

!

!

If the attachment circuits have not been bound to VFI configured with the assigned VPN ID for each VLAN, this is a finding.

Remediation:

Assign globally unique VPN IDs for each customer using VPLS for carrier Ethernet services between multiple sites, and configure the attachment circuits to the appropriate VFI.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#bridge group L2GROUP

RP/0/0/CPU0:R3(config-l2vpn-bg)#bridge-domain L2_BRIDGE_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#interface GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#exit

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#vfi VFI_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-vfi)#vpn-id 101

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

1.70 CISC-RT-000690 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco PE router must be configured to enforce the split-horizon rule for all pseudowires within a Virtual Private LAN Services (VPLS) bridge domain.

```
GROUP ID: V-216797
RULE ID: SV-216797r531087
```

Rationale:

A virtual forwarding instance (VFI) must be created on each participating PE router for each customer VLAN using VPLS for carrier Ethernet services. The VFI specifies the VPN ID of a VPLS domain, the addresses of other PE routers in the domain, and the type of tunnel signaling and encapsulation mechanism for each peer PE router. The set of VFIs formed by the interconnection of the emulated VCs is called a VPLS instance, which forms the logic bridge over the MPLS core network.

The PE routers use the VFI with a unique VPN ID to establish a full mesh of emulated virtual circuits or pseudowires to all the other PE routers in the VPLS instance. The full-mesh configuration allows the PE router to maintain a single broadcast domain. With a full-mesh configuration, signaling and packet replication requirements for each provisioned virtual circuit on a PE can be high. To avoid the problem of a packet looping in the provider core, thereby adding more overhead, the PE devices must enforce a split-horizon principle for the emulated virtual circuits; that is, if a packet is received on an emulated virtual circuit, it is not forwarded on any other virtual circuit.

Audit:

Review the PE router configuration to verify that split horizon is enabled at each attachment circuit within each bridge domain as shown in the example below.

bridge group L2GROUP

bridge-domain L2_BRIDGE_COI1

interface GigabitEthernet0/0/0/2

split-horizon group

!

If split horizon is not enabled, this is a finding.

Note: This requirement is only applicable to a mesh VPLS topology. VPLS solves the loop problem by using a split-horizon rule which states that member PE routers of a VPLS must forward VPLS traffic only to the local attachment circuits when they receive the traffic from the other PE routers. In a ring VPLS, split horizon must be disabled so that a PE router can forward a packet received from one pseudowire to another pseudowire. To prevent the consequential loop, at least one span in the ring would not have a pseudowire for any given VPLS instance.

Remediation:

Enable split horizon on all PE routers deploying VPLS in a full-mesh configuration.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#bridge group L2GROUP

RP/0/0/CPU0:R3(config-l2vpn-bg)#bridge-domain L2_BRIDGE_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#interface GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#split-horizon group

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#end

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

1.71 CISC-RT-000700 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco PE router providing Virtual Private LAN Services (VPLS) must be configured to have traffic storm control thresholds on CE-facing interfaces.

```
GROUP ID: V-216798
RULE ID: SV-216798r531087
```

Rationale:

A traffic storm occurs when packets flood a VPLS bridge, creating excessive traffic and degrading network performance. Traffic storm control prevents VPLS bridge disruption by suppressing traffic when the number of packets reaches configured threshold levels. Traffic storm control monitors incoming traffic levels on a port and drops traffic when the number of packets reaches the configured threshold level during any one-second interval.

Audit:

Review the router configuration to verify that storm control is enabled on CE-facing interfaces deploying VPLS as shown in the example below.

bridge group L2GROUP

bridge-domain L2_BRIDGE_COI1

interface GigabitEthernet0/0/0/2

storm-control unknown-unicast kbps 1200

storm-control multicast kbps 1200

storm-control broadcast kbps 1200

split-horizon group

!

If storm control is not enabled at a minimum for broadcast traffic, this is a finding.

Remediation:

Configure storm control for each CE-facing interface as shown in the example below.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#bridge group L2GROUP

RP/0/0/CPU0:R3(config-l2vpn-bg)# bridge-domain L2_BRIDGE_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#interface GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#storm-control broadcast kbps 1200

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#storm-control multicast kbps 1200

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#storm-control unknown-unicast kbps 1200

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#end

Note: The acceptable range is 10000000 -1000000000 for a gigabit ethernet interface, and 100000000-10000000000 for a ten gigabit interface. Storm control is not supported on most FastEthernet interfaces.

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)

1.72 CISC-RT-000710 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco PE router must be configured to implement Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) snooping for each Virtual Private LAN Services (VPLS) bridge domain.

```
GROUP ID: V-216799
RULE ID: SV-216799r856448
```

Rationale:

IGMP snooping provides a way to constrain multicast traffic at Layer 2. By monitoring the IGMP membership reports sent by hosts within the bridge domain, the snooping application can set up Layer 2 multicast forwarding tables to deliver traffic only to ports with at least one interested member within the VPLS bridge, thereby significantly reducing the volume of multicast traffic that would otherwise flood an entire VPLS bridge domain. The IGMP snooping operation applies to both access circuits and pseudowires within a VPLS bridge domain.

Audit:

Review the router configuration to verify that IGMP or MLD snooping has been configured for IPv4 and IPv6 multicast traffic respectively for each VPLS bridge domain.

l2vpn

bridge group L2GROUP

bridge-domain L2_BRIDGE_COI1

interface GigabitEthernet0/0/0/2

igmp snooping profile default

If the router is not configured to implement IGMP or MLD snooping for each VPLS bridge domain, this is a finding.

Remediation:

Configure IGMP or MLD snooping for IPv4 and IPv6 multicast traffic respectively for each VPLS bridge domain.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#bridge group L2GROUP

RP/0/0/CPU0:R3(config-l2vpn-bg)# bridge-domain L2_BRIDGE_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#interface GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#igmp snooping profile default

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#end

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

1.73 CISC-RT-000720 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco PE router must be configured to limit the number of MAC addresses it can learn for each Virtual Private LAN Services (VPLS) bridge domain.

```
GROUP ID: V-216800
RULE ID: SV-216800r531087
```

Rationale:

VPLS defines an architecture that delivers Ethernet multipoint services over an MPLS network. Customer Layer 2 frames are forwarded across the MPLS core via pseudowires using IEEE 802.1q Ethernet bridging principles. A pseudowire is a virtual bidirectional connection between two attachment circuits (virtual connections between PE and CE routers). A pseudowire contains two unidirectional label-switched paths (LSP). Each MAC forwarding table instance is interconnected using domain-specific LSPs, thereby maintaining privacy and logical separation between each VPLS domain.

When a frame arrives on a bridge port (pseudowire or attachment circuit) and the source MAC address is unknown to the receiving PE router, the source MAC address is associated with the pseudowire or attachment circuit and the forwarding table is updated accordingly. Frames are forwarded to the appropriate pseudowire or attachment circuit according to the forwarding table entry for the destination MAC address. Ethernet frames sent to broadcast and unknown destination addresses must be flooded out to all interfaces for the bridge domain; hence, a PE router must replicate packets across both attachment circuits and pseudowires.

A malicious attacker residing in a customer network could launch a source MAC address spoofing attack by flooding packets to a valid unicast destination, each with a different MAC source address. The PE router receiving this traffic would try to learn every new MAC address and would quickly run out of space for the VFI forwarding table. Older, valid MAC addresses would be removed from the table, and traffic sent to them would have to be flooded until the storm threshold limit is reached. Hence, it is essential that a limit is established to control the number of MAC addresses that will be learned and recorded into the forwarding table for each bridge domain.

Audit:

Review the PE router configuration to determine if a MAC address limit has been set for each VPLS bridge domain.

bridge group L2GROUP

bridge-domain L2_BRIDGE_COI1

interface GigabitEthernet0/0/0/2

mac

limit

maximum nnnn

!

!

If a limit has not been configured, this is a finding.

Remediation:

Configure a MAC address learning limit for each VPLS bridge domain.

RP/0/0/CPU0:R3(config)#l2vpn

RP/0/0/CPU0:R3(config-l2vpn)#bridge group L2GROUP

RP/0/0/CPU0:R3(config-l2vpn-bg)#bridge-domain L2_BRIDGE_COI1

RP/0/0/CPU0:R3(config-l2vpn-bg-bd)#interface GigabitEthernet0/0/0/2

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#mac limit maximum nnn

RP/0/0/CPU0:R3(config-l2vpn-bg-bd-ac)#end

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)

1.74 CISC-RT-000730 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

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

```
GROUP ID: V-216801
RULE ID: SV-216801r531087
```

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:

Step 1: Review the router configuration to verify that an ingress ACL is applied to all external or CE-facing interfaces.

interface GigabitEthernet1/1/0/0

ipv4 address x.1.12.2 255.255.255.252

ipv4 access-group BLOCK_TO_CORE ingress

Step 2: Verify that the ingress ACL discards and logs packets destined to the IP core address space.

Ipv4 access-list BLOCK_TO_CORE

10 deny ipv4 any 10.1.x.0 0.0.255.255 log-input

20 permit ipv4 any any

!

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

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

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 discard and log packets destined to the IP core address space.

RP/0/0/CPU0:R3(config)#Ipv4 access-list BLOCK_TO_CORE

RP/0/0/CPU0:R3(config-ipv4-acl)#deny tcp any any eq tacacs log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#deny ipv4 any 10.1.x.0 0.0.255.255 log-input

RP/0/0/CPU0:R3(config-ipv4-acl)#end

Step 2: Apply the ACL inbound to all external or CE-facing interfaces.

RP/0/0/CPU0:R3(config)#int g1/1/0/0

RP/0/0/CPU0:R3(config-if)#ipv4 access-group BLOCK_TO_CORE in

RP/0/0/CPU0:R3(config-if)#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

1.75 CISC-RT-000740 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216802
RULE ID: SV-216802r531087
```

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

interface GigabitEthernet1/1/0/0

ip address x.1.12.2 255.255.255.252

ipv4 verify unicast source reachable-via any

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

Remediation:

Configure uRPF loose mode on all CE-facing interfaces as shown in the example

RP/0/0/CPU0:R3(config)#int g1/1/0/0

RP/0/0/CPU0:R3(config-if)#ipv4 verify unicast source reachable-via any

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

1.76 CISC-RT-000750 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco PE router must be configured to ignore or block all packets with any IP options.

```
GROUP ID: V-217009
RULE ID: SV-217009r945860
```

Rationale:

Packets with IP options are not fast switched 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:

In Cisco IOS XR, all IPv4 packets with any header option other than the "source-route" header options are dropped. By default, ipv4 source routing is disabled.

Verify that the following command is not configured:

ipv4 source-route

If the router is not configured to drop all packets with IP options, this is a finding.

Remediation:

Configure the router to drop all packets with ipv4 source-route as shown below.

RP/0/0/CPU0:R3(config)#no ipv4 source-route

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)

1.77 CISC-RT-000760 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco PE router must be configured to enforce a Quality-of-Service (QoS) policy to provide preferred treatment for mission-critical applications.

```
GROUP ID: V-216804
RULE ID: SV-216804r917439
```

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.

Step 1: Verify that the class-maps are configured to match on DSCP values as shown in the configuration example below.

class-map match-all VIDEO

match dscp af41

end-class-map

!

class-map match-all VOICE

match dscp ef

end-class-map

!

class-map match-all C2_VOICE

match dscp 47

end-class-map

!

class-map match-all CONTROL_PLANE

match dscp cs6

end-class-map

!

class-map match-all PREFERRED_DATA

match dscp af33

end-class-map

!

Step 2: Verify that the policy map reserves the bandwidth for each traffic type as shown in the following example:

policy-map QOS_POLICY

class C2_VOICE

bandwidth percent 10

!

class VOICE

bandwidth percent 15

!

class VIDEO

bandwidth percent 25

!

class CONTROL_PLANE

bandwidth percent 10

!

class PREFERRED_DATA

bandwidth percent 25

!

class class-default

bandwidth percent 15

!

end-policy-map

!

Step 3: Verify that an output service policy is bound to all interface as shown in the configuration example below.

interface GigabitEthernet0/0/0/1

service-policy output QOS_POLICY

ipv4 address x.1.24.2 255.255.255.252

!

interface GigabitEthernet0/0/0/2

service-policy output QOS_POLICY

ipv4 address x.1.24.5 255.255.255.252

Note: Enclaves must mark or re-mark their traffic to be consistent with the DODIN backbone admission criteria to gain the appropriate level of service. A general DiffServ principle is to mark or trust traffic as close to the source as administratively and technically possible. However, certain traffic types might need to be re-marked before handoff to the DODIN backbone to gain admission to the correct class. If such re-marking is required, it is recommended that the re-marking be performed at the CE egress edge.

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

Remediation:

Configure to enforce a QoS policy to provide preferred treatment for mission-critical applications.

Step 1: Configure class-maps to match on DSCP values as shown in the configuration example below.

RP/0/0/CPU0:R2(config-cmap)#class-map match-all C2_VOICE

RP/0/0/CPU0:R2(config-cmap)#match dscp 47

RP/0/0/CPU0:R2(config-cmap)#class-map match-all VOICE

RP/0/0/CPU0:R2(config-cmap)#match dscp ef

RP/0/0/CPU0:R2(config-cmap)#class-map match-all VIDEO

RP/0/0/CPU0:R2(config-cmap)#match dscp af41

RP/0/0/CPU0:R2(config-cmap)#class-map match-all CONTROL_PLANE

RP/0/0/CPU0:R2(config-cmap)#match dscp cs6

RP/0/0/CPU0:R2(config-cmap)#class-map match-all PREFERRED_DATA

RP/0/0/CPU0:R2(config-cmap)#match dscp af33

RP/0/0/CPU0:R2(config-cmap)#exit

Step 2: Configure a policy map to be applied to the core-layer-facing interface that reserves the bandwidth for each traffic type as shown in the example below.

RP/0/0/CPU0:R2(config-pmap)#policy-map QOS_POLICY

RP/0/0/CPU0:R2(config-pmap)#class C2_VOICE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 10

RP/0/0/CPU0:R2(config-pmap-c)#class VOICE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 15

RP/0/0/CPU0:R2(config-pmap-c)#class VIDEO

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 25

RP/0/0/CPU0:R2(config-pmap-c)#class CONTROL_PLANE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 10

RP/0/0/CPU0:R2(config-pmap-c)#class PREFERRED_DATA

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 25

RP/0/0/CPU0:R2(config-pmap-c)#class class-default

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 15

RP/0/0/CPU0:R2(config-pmap-c)#exit

Step 3: Apply the output service policy to the core-layer-facing interface as shown in the configuration example below.

RP/0/0/CPU0:R2(config)#int g0/0/0/1

RP/0/0/CPU0:R2(config-if)#service-policy output QOS_POLICY

RP/0/0/CPU0:R2(config-if)#exit

RP/0/0/CPU0:R2(config)#int g0/0/0/2

RP/0/0/CPU0:R2(config-if)#service-policy output QOS_POLICY

RP/0/0/CPU0:R2(config-if)#end

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)

1.78 CISC-RT-000770 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco P router must be configured to enforce a Quality-of-Service (QoS) policy to provide preferred treatment for mission-critical applications.

```
GROUP ID: V-216805
RULE ID: SV-216805r917442
```

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.

Step 1: Verify that the class-maps are configured to match on DSCP values as shown in the configuration example below.

class-map match-all VIDEO

match dscp af41

end-class-map

!

class-map match-all VOICE

match dscp ef

end-class-map

!

class-map match-all C2_VOICE

match dscp 47

end-class-map

!

class-map match-all CONTROL_PLANE

match dscp cs6

end-class-map

!

class-map match-all PREFERRED_DATA

match dscp af33

end-class-map

!

Step 2: Verify that the policy map reserves the bandwidth for each traffic type as shown in the following example:

policy-map QOS_POLICY

class C2_VOICE

bandwidth percent 10

!

class VOICE

bandwidth percent 15

!

class VIDEO

bandwidth percent 25

!

class CONTROL_PLANE

bandwidth percent 10

!

class PREFERRED_DATA

bandwidth percent 25

!

class class-default

bandwidth percent 15

!

end-policy-map

!

Step 3: Verify that an output service policy is bound to all interfaces as shown in the configuration example below.

interface GigabitEthernet0/0/0/1

service-policy output QOS_POLICY

ipv4 address x.1.24.2 255.255.255.252

!

interface GigabitEthernet0/0/0/2

service-policy output QOS_POLICY

ipv4 address x.1.24.5 255.255.255.252

Note: Enclaves must mark or re-mark their traffic to be consistent with the DODIN backbone admission criteria to gain the appropriate level of service. A general DiffServ principle is to mark or trust traffic as close to the source as administratively and technically possible. However, certain traffic types might need to be re-marked before handoff to the DODIN backbone to gain admission to the correct class. If such re-marking is required, it is recommended that the re-marking be performed at the CE egress edge.

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

Remediation:

Configure to enforce a QoS policy to provide preferred treatment for mission-critical applications.

Step 1: Configure class-maps to match on DSCP values as shown in the configuration example below.

RP/0/0/CPU0:R2(config-cmap)#class-map match-all C2_VOICE

RP/0/0/CPU0:R2(config-cmap)#match dscp 47

RP/0/0/CPU0:R2(config-cmap)#class-map match-all VOICE

RP/0/0/CPU0:R2(config-cmap)#match dscp ef

RP/0/0/CPU0:R2(config-cmap)#class-map match-all VIDEO

RP/0/0/CPU0:R2(config-cmap)#match dscp af41

RP/0/0/CPU0:R2(config-cmap)#class-map match-all CONTROL_PLANE

RP/0/0/CPU0:R2(config-cmap)#match dscp cs6

RP/0/0/CPU0:R2(config-cmap)#class-map match-all PREFERRED_DATA

RP/0/0/CPU0:R2(config-cmap)#match dscp af33

RP/0/0/CPU0:R2(config-cmap)#exit

Step 2: Configure a policy map to be applied to the core-layer-facing interface that reserves the bandwidth for each traffic type as shown in the example below.

RP/0/0/CPU0:R2(config-pmap)#policy-map QOS_POLICY

RP/0/0/CPU0:R2(config-pmap)#class C2_VOICE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 10

RP/0/0/CPU0:R2(config-pmap-c)#class VOICE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 15

RP/0/0/CPU0:R2(config-pmap-c)#class VIDEO

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 25

RP/0/0/CPU0:R2(config-pmap-c)#class CONTROL_PLANE

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 10

RP/0/0/CPU0:R2(config-pmap-c)#class PREFERRED_DATA

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 25

RP/0/0/CPU0:R2(config-pmap-c)#class class-default

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 15

RP/0/0/CPU0:R2(config-pmap-c)#exit

Step 3: Apply the output service policy to the core-layer-facing interface as shown in the configuration example below.

RP/0/0/CPU0:R2(config)#int g0/0/0/1

RP/0/0/CPU0:R2(config-if)#service-policy output QOS_POLICY

RP/0/0/CPU0:R2(config-if)#exit

RP/0/0/CPU0:R2(config)#int g0/0/0/2

RP/0/0/CPU0:R2(config-if)#service-policy output QOS_POLICY

RP/0/0/CPU0:R2(config-if)#end

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)

1.79 CISC-RT-000780 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216806
RULE ID: SV-216806r531087
```

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:

Step 1: Verify that a class map has been configured for the Scavenger class as shown in the example below.

class-map match-all SCAVENGER

match dscp cs1

end-class-map

Step 2: Verify that the policy map includes the SCAVENGER class with low priority as shown in the following example below.

policy-map QOS_POLICY

class C2_VOICE

bandwidth percent 10

!

class VOICE

bandwidth percent 15

!

class VIDEO

bandwidth percent 25

!

class CONTROL_PLANE

bandwidth percent 10

!

class PREFERRED_DATA

bandwidth percent 25

!

class SCAVENGER

bandwidth percent 5

!

class class-default

bandwidth percent 10

!

end-policy-map

Note: Traffic out of profile must be marked at the customer access layer or CE egress edge.

If the policy map does not include the SCAVENGER class with low priority, this is a finding.

Remediation:

Step 1: Configure a class map for the SCAVENGER class.

RP/0/0/CPU0:R2(config)#class-map match-all SCAVENGER

RP/0/0/CPU0:R2(config-cmap)#match dscp cs1

RP/0/0/CPU0:R2(config-cmap)#exit

Step 2: Add the SCAVENGER class to the policy map as shown in the example below.

RP/0/0/CPU0:R2(config)#policy-map QOS_POLICY

RP/0/0/CPU0:R2(config-pmap)#no class class-default

RP/0/0/CPU0:R2(config-pmap)#class SCAVENGER

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 5

RP/0/0/CPU0:R2(config-pmap-c)#class class-default

RP/0/0/CPU0:R2(config-pmap-c)#bandwidth percent 10

RP/0/0/CPU0:R2(config-pmap-c)#end

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)

1.80 CISC-RT-000790 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216807
RULE ID: SV-216807r531087
```

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:

Step 1: Review the network's multicast topology diagram.

Step 2: Review the router configuration to verify that only the interfaces as shown in the multicast topology diagram are enabled as shown in the example below.

!

router pim

address-family ipv4

interface GigabitEthernet0/0/0/0

enable

!

interface GigabitEthernet0/0/0/1

enable

!

!

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.

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#int g0/0/0/1

RP/0/0/CPU0:R2(config-pim-ipv4-if)#disable

RP/0/0/CPU0:R2(config-pim-ipv4-if)#end

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

1.81 CISC-RT-000800 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216808
RULE ID: SV-216808r531087
```

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.

Step 1: Verify all interfaces enabled for PIM have a neighbor ACL bound to the interface as shown in the example below.

router pim

address-family ipv4

interface GigabitEthernet0/0/0/1

enable

neighbor-filter PIM_NEIGHBOR_1

!

interface GigabitEthernet0/0/0/2

enable

neighbor-filter PIM_NEIGHBOR_2

!

!

!

Step 2: Review the configured ACL for filtering PIM neighbors as shown in the example below.

ipv4 access-list PIM_NEIGHBOR_1

10 permit ipv4 host 10.1.1.2 any

!

ipv4 access-list PIM_NEIGHBOR_2

10 permit ipv4 host 10.1.2.8 any

!

If PIM neighbor ACLs 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 ACLs to only accept PIM control plane traffic from documented PIM neighbors. Bind neighbor ACLs to all PIM enabled interfaces.

Step 1: Configure ACL for PIM neighbors.

RP/0/0/CPU0:R2(config)#ipv4 access-list PIM_NEIGHBOR_1

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 host 10.1.1.2 any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

Step 2: Apply the ACL to all interfaces enabled for PIM.

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#int g0/0/0/1

RP/0/0/CPU0:R2(config-pim-ipv4-if)#neighbor-filter PIM_NEIGHBOR_1

RP/0/0/CPU0:R2(config-pim-ipv4-if)#exit

RP/0/0/CPU0:R2(config-pim-default-ipv4)#int g0/0/0/2

RP/0/0/CPU0:R2(config-pim-ipv4-if)#neighbor-filter PIM_NEIGHBOR_2

RP/0/0/CPU0:R2(config-pim-ipv4-if)#end

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

1.82 CISC-RT-000810 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216809
RULE ID: SV-216809r531087
```

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 router configuration and verify that admin-scope multicast traffic is blocked at the external edge as shown in the example below.

ipv4 access-list MULTICAST_SCOPE

10 deny ipv4 239.0.0.0 0.255.255.255 any

20 permit ipv4 any any

…

…

…

multicast-routing

address-family ipv4

interface GigabitEthernet0/0/0/1

enable

boundary MULTICAST_SCOPE

!

interface GigabitEthernet0/0/0/2

enable

!

!

!

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

Remediation:

Step 1: Configure the ACL to deny packets with multicast administratively scoped destination addresses as shown in the example below.

RP/0/0/CPU0:R2(config)#Ipv4 access-list MULTICAST_SCOPE

RP/0/0/CPU0:R2(config-ipv4-acl)#deny 239.0.0.0 0.255.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#permit any

Step 2: Apply the multicast boundary at the appropriate interfaces as shown in the example below.

RP/0/0/CPU0:R2(config)#multicast-routing

RP/0/0/CPU0:R2(config-mcast)#address-family ipv4

RP/0/0/CPU0:R2(config-mcast-default-ipv4)#int g0/0/0/1

RP/0/0/CPU0:R2(config-mcast-default-ipv4-if)#boundary MULTICAST_SCOPE

RP/0/0/CPU0:R2(config-mcast-default-ipv4-if)#end

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

1.83 CISC-RT-000820 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216810
RULE ID: SV-216810r856449
```

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:

The Cisco router does not have a mechanism to limit the multicast forwarding cache. However, the risk associated with this requirement can be fully mitigated by configuring the router to:

1. Filter PIM register messages.
2. Rate limiting the number of PIM register messages.
3. Accept MSDP packets only from known MSDP peers.

Step 1: Verify that the RP router is configured to filter PIM register messages for any undesirable multicast groups and sources. The example below will deny any multicast streams for groups 239.5.0.0/16 and allow from only sources 10.1.2.6 and 10.1.2.7.
ipv4 access-list PIM_REGISTER_FILTER
10 deny ipv4 any 239.5.0.0 0.0.255.255
20 permit ipv4 host 10.1.2.6 any
30 permit ipv4 host 10.1.2.7 any
40 deny ipv4 any any
…
…
…
router pim
address-family ipv4
rp-address 2.2.2.2
accept-register PIM_REGISTER_FILTER

Step 2: Verify that the router is configured to rate limiting the number of PIM register messages as shown in the example below.
router pim
address-family ipv4
allow-rp group-list FILTER_PIM_JOINS
rp-address 10.2.2.2
accept-register PIM_REGISTER_FILTER
maximum register-states 250
Note: The maximum register-states command is used to set an upper limit for PIM register states. When the limit is reached, PIM discontinues route creation from PIM register messages. If not configured, the default is 2000 which would be an overage for a small to average size multicast deployment.
Step 3: Review the router configuration to determine if there is a receive path or interface filter to only accept MSDP packets from known MSDP peers as shown in the example below.
Step 3a: Determine which interfaces would be peering MSDP with an external router by the configured peer addresses as shown in the example below.
router msdp
peer x.14.2.1
remote-as nn
!
peer x.15.3.5
remote-as nn
!
!
Step 3b: Verify that interfaces used for MSDP peering have an inbound ACL as shown in the example.
interface GigabitEthernet0/0/0/1
ipv4 address x.14.2.2 255.255.255.252
ipv4 access-group EXTERNAL_ACL_INBOUND ingress
Step 3c: Verify that the ACL restricts MSDP peering to only known sources.
ipv4 access-list EXTERNAL_ACL_INBOUND
10 permit tcp host x.1.28.2 host x.1.28.8 eq 639
20 deny tcp any host x.1.28.8 eq 639 log
30 permit tcp host x.1.28.2 host x.1.28.8 eq bgp
40 permit tcp host x.1.28.2 eq bgp host x.1.28.8
50 permit pim host x.1.28.2 host x.1.28.8
60 permit tcp any any established
…
…
…
140 deny ipv4 any any log
Note: MSDP connections is via TCP port 639
If the RP router is not configured to filter PIM register messages, rate limiting the number of PIM register messages, and accept MSDP packets only from known MSDP peers, this is a finding.

Remediation:

The risk associated with this requirement can be fully mitigated by configuring the router to filter PIM register messages, rate limiting the number of PIM register messages, and accept MSDP packets only from known MSDP peers.

Step 1: Configure the router to filter PIM register messages received from a multicast DR for any undesirable multicast groups and sources. The example below will deny any multicast streams for groups 239.5.0.0/16 and allow from only sources 10.1.2.6 and 10.1.2.7.

RP/0/0/CPU0:R2(config)#ipv4 access-list PIM_REGISTER_FILTER

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any 239.5.0.0 0.0.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 host 10.1.2.6 any

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 host 10.1.2.7 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#accept-register PIM_REGISTER_FILTER

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

Step 2: Configure the RP to rate limit the number of multicast register messages.

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#maximum register-states 250

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

Step 3: Configure the receive path or interface ACLs to only accepts MSDP packets from known MSDP peers.

RP/0/0/CPU0:R2(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 host x.1.28.8 eq 639

RP/0/0/CPU0:R2(config-ipv4-acl)#deny tcp any host x.1.28.8 eq 639 log

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 host x.1.28.8 eq bgp

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 eq bgp host x.1.28.8

RP/0/0/CPU0:R2(config-ipv4-acl)#permit pim host x.1.28.2 host x.1.28.8

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any any established

…

…

…

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log

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

1.84 CISC-RT-000830 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216811
RULE ID: SV-216811r531087
```

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.

Audit:

Verify that the RP router is configured to filter PIM register messages. The example below will deny any multicast streams for groups 239.5.0.0/16 and allow from only sources 10.1.2.6 and 10.1.2.7.

ipv4 access-list PIM_REGISTER_FILTER

10 deny ipv4 any 239.5.0.0 0.0.255.255

20 permit ipv4 host 10.1.2.6 any

30 permit ipv4 host 10.1.2.7 any

40 deny ipv4 any any

…

…

…

router pim

address-family ipv4

rp-address 2.2.2.2

accept-register PIM_REGISTER_FILTER

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

Remediation:

Configure the router to filter PIM register messages received from a multicast DR for any undesirable multicast groups and sources. The example below will deny any multicast streams for groups 239.5.0.0/16 and allow from only sources 10.1.2.6 and 10.1.2.7.

RP/0/0/CPU0:R2(config)#ipv4 access-list PIM_REGISTER_FILTER

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any 239.5.0.0 0.0.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 host 10.1.2.6 any

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 host 10.1.2.7 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#accept-register PIM_REGISTER_FILTER

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

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

1.85 CISC-RT-000840 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

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

```
GROUP ID: V-216812
RULE ID: SV-216812r531087
```

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 join messages are only accepted for authorized multicast groups.

Audit:

Verify that the RP router is configured to filter PIM join messages for any undesirable multicast groups. In the example below, groups from 239.8.0.0/16 are no allowed.

ipv4 access-list FILTER_PIM_JOINS

10 deny ipv4 239.8.0.0 0.0.255.255 any

20 permit ipv4 any any

…

…

…

router pim

address-family ipv4

allow-rp group-list FILTER_PIM_JOINS

If the RP is not configured to filter join messages received from the DR for any undesirable multicast groups, this is a finding.

Remediation:

Configure the RP to filter PIM join messages for any undesirable multicast groups as shown in the example below.

RP/0/0/CPU0:R2(config)#ipv4 access-list FILTER_PIM_JOINS

RP/0/0/CPU0:R2(config-ipv4-acl)#deny 239.8.0.0 0.0.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#permit any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#allow-rp group-list FILTER_PIM_JOINS

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

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

1.86 CISC-RT-000850 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco multicast Rendezvous Point (RP) must be configured to rate limit the number of Protocol Independent Multicast (PIM) Register messages.

```
GROUP ID: V-216813
RULE ID: SV-216813r856450
```

Rationale:

When a new source starts transmitting in a PIM Sparse Mode network, the Designated Router (DR) will encapsulate the multicast packets into register messages and forward them to the RP using unicast.

This process can be taxing on the CPU for both the DR and the RP if the source is running at a high data rate and there are many new sources starting at the same time. This scenario can potentially occur immediately after a network failover.

The rate limit for the number of register messages should be set to a relatively low value based on the known number of multicast sources within the multicast domain.

Audit:

Review the configuration of the RP to verify that it is limiting the number of PIM register states as shown in the example.

router pim

address-family ipv4

allow-rp group-list FILTER_PIM_JOINS

rp-address 10.2.2.2

accept-register PIM_REGISTER_FILTER

maximum register-states 250

Note: The maximum register-states command is used to set an upper limit for PIM register states. When the limit is reached, PIM discontinues route creation from PIM register messages. If not configured, the default is 2000 which would be an overage for a small to average size multicast deployment.

If the RP is not limiting PIM register states, this is a finding.

Remediation:

Configure the RP to rate limit the number of multicast register states.

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#maximum register-states 250

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

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

1.87 CISC-RT-000860 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216814
RULE ID: SV-216814r878125
```

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 DR to verify that it is filtering IGMP or MLD Membership Report messages, allowing hosts to join only those groups that have been approved.

Step 1: Verify that all host facing interfaces are configured to filter IGMP Membership Report messages (IGMP joins) as shown in the example below.

router igmp

interface GigabitEthernet0/0/1/0

access-group IGMP_JOIN_FILTER

!

interface GigabitEthernet0/0/1/1

access-group IGMP_JOIN_FILTER

!

Step 2: Verify that the ACL denies unauthorized groups or permits only authorized groups. The example below denies all groups from 239.8.0.0/16 range.

ipv4 access-list IGMP_JOIN_FILTER

10 deny ipv4 239.8.0.0 0.0.255.255 any

20 permit ipv4 any any

!

Note: This requirement is only applicable to Source Specific Multicast (SSM) implementation. This requirement is not applicable to Any Source Multicast (ASM) since the filtering is being performed by the Rendezvous Point router.

If the DR is not filtering IGMP or MLD Membership Report messages, this is a finding.

Remediation:

Configure the DR to filter the IGMP or MLD Membership Report messages to allow hosts to join only those multicast groups that have been approved.

Step 1: Configure the ACL to filter IGMP Membership Report messages as shown in the example.

RP/0/0/CPU0:R5(config)#ipv4 access-list IGMP_JOIN_FILTER

RP/0/0/CPU0:R5(config-ipv4-acl)#deny 239.8.0.0 0.0.255.255

RP/0/0/CPU0:R5(config-ipv4-acl)#exit

Step 2: Apply the filter to all host facing interfaces.

RP/0/0/CPU0:R5(config)#router igmp

RP/0/0/CPU0:R5(config-igmp)#interface g0/0/1/0

RP/0/0/CPU0:R5(config-igmp-default-if)#access-group IGMP_JOIN_FILTER

RP/0/0/CPU0:R5(config-igmp-default-if)#exit

RP/0/0/CPU0:R5(config-igmp)#interface g0/0/1/1

RP/0/0/CPU0:R5(config-igmp-default-if)#end

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)

1.88 CISC-RT-000870 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco 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-216815
RULE ID: SV-216815r878126
```

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 DR to verify that it is filtering IGMP or MLD report messages, allowing hosts to only join multicast groups from sources that have been approved.

Step 1: Verify that all host facing interfaces are configured to filter IGMP Membership Report messages (IGMP joins) as shown in the example below.

router igmp

interface GigabitEthernet0/0/1/0

access-group IGMP_JOIN_FILTER

!

interface GigabitEthernet0/0/1/1

access-group IGMP_JOIN_FILTER

!

Step 2: Verify that the ACL denies unauthorized sources or allows only authorized sources. The example below denies all groups from 232.8.0.0/16 range and permits sources only from the x.0.0.0/8 network.

ipv4 access-list IGMP_JOIN_FILTER

10 deny ipv4 any 232.8.0.0 0.0.255.255

20 permit ipv4 x.0.0.0 0.255.255.255 any

30 deny ipv4 any any

!

Note: This requirement is only applicable to Source Specific Multicast (SSM) implementation.

If the DR is not filtering IGMP or MLD report messages, this is a finding.

Remediation:

Configure the DR to filter the IGMP and MLD report messages to allow hosts to join only those multicast groups from sources that have been approved as shown in the example.

Step 1: Configure the access list to filter multicast joins as shown in the example below.

RP/0/0/CPU0:R2(config)#ipv4 access-list IGMP_JOIN_FILTER

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any 232.8.0.0 0.0.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip x.0.0.0 0.255.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any

Step 2: Apply the filter to all host facing interfaces.

RP/0/0/CPU0:R5(config)#router igmp

RP/0/0/CPU0:R5(config-igmp)#interface g0/0/1/0

RP/0/0/CPU0:R5(config-igmp-default-if)#access-group IGMP_JOIN_FILTER

RP/0/0/CPU0:R5(config-igmp-default-if)#exit

RP/0/0/CPU0:R5(config-igmp)#interface g0/0/1/1

RP/0/0/CPU0:R5(config-igmp-default-if)#end

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)

1.89 CISC-RT-000880 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco multicast Designated Router (DR) must be configured to limit the number of mroute states resulting from Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Host Membership Reports.

```
GROUP ID: V-216816
RULE ID: SV-216816r856453
```

Rationale:

The current multicast paradigm can let any host join any multicast group at any time by sending an IGMP or MLD membership report to the DR. In a Protocol Independent Multicast (PIM) Sparse Mode network, the DR will send a PIM Join message for the group to the RP. Without any form of admission control, this can pose a security risk to the entire multicast domain - specifically the multicast routers along the shared tree from the DR to the RP that must maintain the mroute state information for each group join request. Hence, it is imperative that the DR is configured to limit the number of mroute state information that must be maintained to mitigate the risk of IGMP or MLD flooding.

Audit:

Review the DR configuration to verify that it is limiting the number of mroute states via IGMP or MLD.

Verify IGMP limits have been configured globally or on each host-facing interface via the ip igmp limit command as shown in the example.

router igmp

interface GigabitEthernet0/0/0/0

access-group IGMP_JOIN_FILTER

!

interface GigabitEthernet0/0/0/1

access-group IGMP_JOIN_FILTER

!

maximum groups 200

!Note: After the maximum groups value is met, all additional memberships learned are ignored. If not configured, the default is 5000 which would be an overage for a small to average size multicast deployment.

If the DR is not limiting multicast join requests via IGMP or MLD on a global or interfaces basis, this is a finding.

Remediation:

Configure the DR on a global or interface basis to limit the number of mroute states resulting from IGMP or MLD membership reports.

RP/0/0/CPU0:R5(config)#router igmp

RP/0/0/CPU0:R2(config-igmp)#maximum groups 200

RP/0/0/CPU0:R5(config-igmp)#end

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

1.90 CISC-RT-000890 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco multicast Designated Router (DR) must be configured to set the shortest-path tree (SPT) threshold to infinity to minimalize source-group (S, G) state within the multicast topology where Any Source Multicast (ASM) is deployed.

```
GROUP ID: V-216817
RULE ID: SV-216817r945856
```

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 switch 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 switchover 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 switchover occurs.

Audit:

Review the DR configuration to verify that the SPT switchover threshold is increased (default is "0") or set to infinity (never switch over).

router pim

address-family ipv4

…

…

…

spt-threshold infinity

If the DR is not configured to increase the SPT threshold or set to infinity to minimalize (S, G) state, this is a finding.

Remediation:

Configure the DR to increase the SPT threshold or set it to infinity to minimalize (S, G) state within the multicast topology where ASM is deployed.

RP/0/0/CPU0:R2(config)#router pim

RP/0/0/CPU0:R2(config-pim)#address-family ipv4

RP/0/0/CPU0:R2(config-pim-default-ipv4)#spt-threshold infinity

RP/0/0/CPU0:R2(config-pim-default-ipv4)#end

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

1.91 CISC-RT-000900 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco Multicast Source Discovery Protocol (MSDP) router must be configured to only accept MSDP packets from known MSDP peers.

```
GROUP ID: V-216818
RULE ID: SV-216818r856455
```

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 router configuration to determine if there is a receive path or interface filter to only accept MSDP packets from known MSDP peers.

Step 1: Determine which interfaces would be peering MSDP with an external router by the configured peer addresses as shown in the example below.

router msdp

peer x.14.2.1

remote-as nn

!

peer x.15.3.5

remote-as nn

!

!

Step 2: Verify that interfaces used for MSDP peering have an inbound ACL as shown in the example.

interface GigabitEthernet0/0/0/1

ipv4 address x.14.2.2 255.255.255.252

ipv4 access-group EXTERNAL_ACL_INBOUND ingress

Step 3: Verify that the ACL restricts MSDP peering to only known sources.

ipv4 access-list EXTERNAL_ACL_INBOUND

10 permit tcp host x.1.28.2 host x.1.28.8 eq 639

20 deny tcp any host x.1.28.8 eq 639 log

30 permit tcp host x.1.28.2 host x.1.28.8 eq bgp

40 permit tcp host x.1.28.2 eq bgp host x.1.28.8

50 permit pim host x.1.28.2 host x.1.28.8

60 permit tcp any any established

…

…

…

140 deny ipv4 any any log

Note: MSDP connections is via TCP port 639

If the router is not configured to only accept MSDP packets from known MSDP peers, this is a finding.

Remediation:

Configure the interface ACLs to only accept MSDP packets from known MSDP peers.

RP/0/0/CPU0:R2(config)#ipv4 access-list EXTERNAL_ACL_INBOUND

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 host x.1.28.8 eq 639

RP/0/0/CPU0:R2(config-ipv4-acl)#deny tcp any host x.1.28.8 eq 639 log

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 host x.1.28.8 eq bgp

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp host x.1.28.2 eq bgp host x.1.28.8

RP/0/0/CPU0:R2(config-ipv4-acl)#permit pim host x.1.28.2 host x.1.28.8

RP/0/0/CPU0:R2(config-ipv4-acl)#permit tcp any any established

…

…

…

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any any log

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)

1.92 CISC-RT-000910 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco Multicast Source Discovery Protocol (MSDP) router must be configured to authenticate all received MSDP packets.

```
GROUP ID: V-216819
RULE ID: SV-216819r856456
```

Rationale:

MSDP peering with customer network routers presents additional risks to the core, whether from a rogue or misconfigured MSDP-enabled router.

MSDP password authentication is used to validate each segment sent on the TCP connection between MSDP peers, protecting the MSDP session against the threat of spoofed packets being injected into the TCP connection stream.

Audit:

Review the router configuration to determine if received MSDP packets are authenticated.

router msdp

peer x.14.2.1

password encrypted 094E410B1B1C

remote-as nn

!

peer x.15.3.5

password encrypted 04500A140A2F

remote-as nn

!

!

If the router does not require MSDP authentication, this is a finding.

Remediation:

Configure the router to authenticate MSDP messages as shown in the following example:

RP/0/0/CPU0:R2(config)#router msdp

RP/0/0/CPU0:R2(config-msdp)#peer x.14.2.1

RP/0/0/CPU0:R2(config-msdp-peer)#password clear xxxxxxxxxxx

RP/0/0/CPU0:R2(config-msdp-peer)#exit

RP/0/0/CPU0:R2(config-msdp)#peer x.15.3.5

RP/0/0/CPU0:R2(config-msdp-peer)#password clear xxxxxxxxxxx

RP/0/0/CPU0:R2(config-msdp-peer)#end

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

1.93 CISC-RT-000920 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216820
RULE ID: SV-216820r531087
```

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 router configuration to determine if there is 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.

router msdp

sa-filter in list INBOUND_MSDP_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.

ipv4 access-list INBOUND_MSDP_SA_FILTER

10 deny ipv4 any host 224.0.1.3

20 deny ipv4 any host 224.0.1.24

30 deny ipv4 any host 224.0.1.22

40 deny ipv4 any host 224.0.1.2

50 deny ipv4 any host 224.0.1.35

60 deny ipv4 any host 224.0.1.60

70 deny ipv4 any host 224.0.1.39

80 deny ipv4 any host 224.0.1.40

90 deny ipv4 any 232.0.0.0 0.255.255.255

100 deny ipv4 any 239.0.0.0 0.255.255.255

110 deny ipv4 10.0.0.0 0.255.255.255 any

120 deny ipv4 127.0.0.0 0.255.255.255 any

130 deny ipv4 172.16.0.0 0.15.255.255 any

140 deny ipv4 192.168.0.0 0.0.255.255 any

150 permit ipv4 any any

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

Remediation:

Configure the MSDP router to filter received source-active multicast advertisements for any undesirable multicast groups and sources as shown in the example below.

RP/0/0/CPU0:R2(config)#ipv4 access-list INBOUND_MSDP_SA_FILTER

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ip any host 224.0.1.3

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.24

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.22

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.2

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.35

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.60

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.39

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any host 224.0.1.40

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any 232.0.0.0 0.255.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 any 239.0.0.0 0.255.255.255

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 10.0.0.0 0.255.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 127.0.0.0 0.255.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 172.16.0.0 0.15.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 192.168.0.0 0.0.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ipv4 any any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#router msdp

RP/0/0/CPU0:R2(config-msdp)#sa-filter in list INBOUND_MSDP_SA_FILTER

RP/0/0/CPU0:R2(config-msdp)#end

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

1.94 CISC-RT-000930 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco 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-216821
RULE ID: SV-216821r531087
```

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 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 as shown in the example below.

router msdp

sa-filter in list INBOUND_MSDP_SA_FILTER

sa-filter out list OUTBOUND_MSDP_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.

ipv4 access-list OUTBOUND_MSDP_SA_FILTER

10 deny ipv4 10.0.0.0 0.255.255.255 any

20 permit ipv4 any any

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

Remediation:

Configure the router with an export policy to avoid global visibility of local multicast (S, G) states. The example below will prevent exporting multicast active sources belonging to the private network.

RP/0/0/CPU0:R2(config)#ipv4 access-list OUTBOUND_MSDP_SA_FILTER

RP/0/0/CPU0:R2(config-ipv4-acl)#deny ipv4 10.0.0.0 0.255.255.255 any

RP/0/0/CPU0:R2(config-ipv4-acl)#permit ip any any

RP/0/0/CPU0:R2(config-ipv4-acl)#exit

RP/0/0/CPU0:R2(config)#router msdp

RP/0/0/CPU0:R2(config-msdp)#sa-filter out list OUTBOUND_MSDP_SA_FILTER

RP/0/0/CPU0:R2(config-msdp)#end

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

1.95 CISC-RT-000940 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco Multicast Source Discovery Protocol (MSDP) router must be configured to limit the amount of source-active messages it accepts on a per-peer basis.

```
GROUP ID: V-216822
RULE ID: SV-216822r531087
```

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:

Review the router configuration to determine if it is configured to limit the amount of source-active messages it accepts on a per-peer basis.

router msdp

…

…

…

peer 4.4.4.4

remote-as 33

maximum external-sa 555

!

peer 5.5.5.5

remote-as 44

maximum external-sa 555

!

!

If the router is not configured to limit the source-active messages it accepts, this is a finding.

Remediation:

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

RP/0/0/CPU0:R2(config)#router msdp

RP/0/0/CPU0:R2(config-msdp)#peer x.14.2.1

RP/0/0/CPU0:R2(config-msdp-peer)#maximum external-sa nnn

RP/0/0/CPU0:R2(config-msdp-peer)#exit

RP/0/0/CPU0:R2(config-msdp)#peer x.15.3.5

RP/0/0/CPU0:R2(config-msdp-peer)#maximum external-sa nnn

RP/0/0/CPU0:R2(config-msdp-peer)#end

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

1.96 CISC-RT-000950 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco Multicast Source Discovery Protocol (MSDP) router must be configured to use a loopback address as the source address when originating MSDP traffic.

```
GROUP ID: V-216823
RULE ID: SV-216823r991943
```

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:

Verify that the loopback interface is used as the source address for all MSDP packets generated by the router.

router msdp

connect-source Loopback0

If the router does not use its loopback address as the source address when originating MSDP traffic, this is a finding.

Remediation:

Configure the router to use its loopback address is used as the source address when sending MSDP packets.

RP/0/0/CPU0:R2(config)#router msdp

RP/0/0/CPU0:R2(config-msdp)#connect-source lo0

RP/0/0/CPU0:R2(config-msdp)#end

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

Appendix: Summary Table

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1	**STIG RULES**		
1.1	CISC-RT-000010 (Manual)	☐	☐
1.2	CISC-RT-000050 (Manual)	☐	☐
1.3	CISC-RT-000060 (Manual)	☐	☐
1.4	CISC-RT-000070 (Manual)	☐	☐
1.5	CISC-RT-000130 (Manual)	☐	☐
1.6	CISC-RT-000140 (Manual)	☐	☐
1.7	CISC-RT-000160 (Manual)	☐	☐
1.8	CISC-RT-000170 (Manual)	☐	☐
1.9	CISC-RT-000180 (Manual)	☐	☐
1.10	CISC-RT-000190 (Manual)	☐	☐
1.11	CISC-RT-000200 (Manual)	☐	☐
1.12	CISC-RT-000210 (Manual)	☐	☐
1.13	CISC-RT-000220 (Manual)	☐	☐
1.14	CISC-RT-000235 (Manual)	☐	☐
1.15	CISC-RT-000236 (Manual)	☐	☐
1.16	CISC-RT-000237 (Manual)	☐	☐
1.17	CISC-RT-000240 (Manual)	☐	☐
1.18	CISC-RT-000250 (Manual)	☐	☐
1.19	CISC-RT-000260 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.20	CISC-RT-000270 (Manual)	☐	☐
1.21	CISC-RT-000280 (Manual)	☐	☐
1.22	CISC-RT-000290 (Manual)	☐	☐
1.23	CISC-RT-000300 (Manual)	☐	☐
1.24	CISC-RT-000310 (Manual)	☐	☐
1.25	CISC-RT-000320 (Manual)	☐	☐
1.26	CISC-RT-000330 (Manual)	☐	☐
1.27	CISC-RT-000340 (Manual)	☐	☐
1.28	CISC-RT-000350 (Manual)	☐	☐
1.29	CISC-RT-000360 (Manual)	☐	☐
1.30	CISC-RT-000370 (Manual)	☐	☐
1.31	CISC-RT-000380 (Manual)	☐	☐
1.32	CISC-RT-000390 (Manual)	☐	☐
1.33	CISC-RT-000391 (Manual)	☐	☐
1.34	CISC-RT-000392 (Manual)	☐	☐
1.35	CISC-RT-000393 (Manual)	☐	☐
1.36	CISC-RT-000394 (Manual)	☐	☐
1.37	CISC-RT-000395 (Manual)	☐	☐
1.38	CISC-RT-000396 (Manual)	☐	☐
1.39	CISC-RT-000397 (Manual)	☐	☐
1.40	CISC-RT-000398 (Manual)	☐	☐
1.41	CISC-RT-000400 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.42	CISC-RT-000410 (Manual)	☐	☐
1.43	CISC-RT-000420 (Manual)	☐	☐
1.44	CISC-RT-000430 (Manual)	☐	☐
1.45	CISC-RT-000440 (Manual)	☐	☐
1.46	CISC-RT-000450 (Manual)	☐	☐
1.47	CISC-RT-000460 (Manual)	☐	☐
1.48	CISC-RT-000470 (Manual)	☐	☐
1.49	CISC-RT-000480 (Manual)	☐	☐
1.50	CISC-RT-000490 (Manual)	☐	☐
1.51	CISC-RT-000500 (Manual)	☐	☐
1.52	CISC-RT-000510 (Manual)	☐	☐
1.53	CISC-RT-000520 (Manual)	☐	☐
1.54	CISC-RT-000530 (Manual)	☐	☐
1.55	CISC-RT-000540 (Manual)	☐	☐
1.56	CISC-RT-000550 (Manual)	☐	☐
1.57	CISC-RT-000560 (Manual)	☐	☐
1.58	CISC-RT-000570 (Manual)	☐	☐
1.59	CISC-RT-000580 (Manual)	☐	☐
1.60	CISC-RT-000590 (Manual)	☐	☐
1.61	CISC-RT-000600 (Manual)	☐	☐
1.62	CISC-RT-000610 (Manual)	☐	☐
1.63	CISC-RT-000620 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.64	CISC-RT-000630 (Manual)	☐	☐
1.65	CISC-RT-000640 (Manual)	☐	☐
1.66	CISC-RT-000650 (Manual)	☐	☐
1.67	CISC-RT-000660 (Manual)	☐	☐
1.68	CISC-RT-000670 (Manual)	☐	☐
1.69	CISC-RT-000680 (Manual)	☐	☐
1.70	CISC-RT-000690 (Manual)	☐	☐
1.71	CISC-RT-000700 (Manual)	☐	☐
1.72	CISC-RT-000710 (Manual)	☐	☐
1.73	CISC-RT-000720 (Manual)	☐	☐
1.74	CISC-RT-000730 (Manual)	☐	☐
1.75	CISC-RT-000740 (Manual)	☐	☐
1.76	CISC-RT-000750 (Manual)	☐	☐
1.77	CISC-RT-000760 (Manual)	☐	☐
1.78	CISC-RT-000770 (Manual)	☐	☐
1.79	CISC-RT-000780 (Manual)	☐	☐
1.80	CISC-RT-000790 (Manual)	☐	☐
1.81	CISC-RT-000800 (Manual)	☐	☐
1.82	CISC-RT-000810 (Manual)	☐	☐
1.83	CISC-RT-000820 (Manual)	☐	☐
1.84	CISC-RT-000830 (Manual)	☐	☐
1.85	CISC-RT-000840 (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.86	CISC-RT-000850 (Manual)	☐	☐
1.87	CISC-RT-000860 (Manual)	☐	☐
1.88	CISC-RT-000870 (Manual)	☐	☐
1.89	CISC-RT-000880 (Manual)	☐	☐
1.90	CISC-RT-000890 (Manual)	☐	☐
1.91	CISC-RT-000900 (Manual)	☐	☐
1.92	CISC-RT-000910 (Manual)	☐	☐
1.93	CISC-RT-000920 (Manual)	☐	☐
1.94	CISC-RT-000930 (Manual)	☐	☐
1.95	CISC-RT-000940 (Manual)	☐	☐
1.96	CISC-RT-000950 (Manual)	☐	☐

Appendix: Change History

Date	Version	Changes for this version
Sep 3, 2025	1.0.0	Initial CIS Release