CIS Cisco IOS Router RTR STIG Benchmark

v1.0.0 – 10-25-2023

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.

Internal Only - General

Table of Contents

Terms of Use *1*

Table of Contents *2*

Overview *5*

Target Technology Details **5**

Intended Audience **5**

Recommendation Definitions *6*

Title **6**

Assessment Status **6**

Automated **6**

Manual **6**

Profile **6**

Description **6**

Rationale Statement **6**

Audit Procedure **6**

Remediation Procedure **7**

Additional Information **7**

Profile Definitions **8**

Acknowledgements **9**

Recommendations *10*

1 STIG RULES **10**

1.1 CISC-RT-000010 (Manual) 11
1.2 CISC-RT-000050 (Manual) 14
1.3 CISC-RT-000060 (Manual) 19
1.4 CISC-RT-000090 (Manual) 21
1.5 CISC-RT-000120 (Manual) 23
1.6 CISC-RT-000150 (Manual) 29
1.7 CISC-RT-000160 (Manual) 30
1.8 CISC-RT-000170 (Manual) 32
1.9 CISC-RT-000180 (Manual) 34
1.10 CISC-RT-000190 (Manual) 35
1.11 CISC-RT-000200 (Manual) 36
1.12 CISC-RT-000210 (Manual) 38
1.13 CISC-RT-000220 (Manual) 40
1.14 CISC-RT-000230 (Manual) 42
1.15 CISC-RT-000235 (Manual) 44
1.16 CISC-RT-000236 (Manual) 45
1.17 CISC-RT-000237 (Manual) 46
1.18 CISC-RT-000240 (Manual) 47
1.19 CISC-RT-000250 (Manual) 50
1.20 CISC-RT-000260 (Manual) 52

Internal Only - General

1.21 CISC-RT-000270 (Manual) 54
1.22 CISC-RT-000280 (Manual) 57
1.23 CISC-RT-000290 (Manual) 60
1.24 CISC-RT-000300 (Manual) 62
1.25 CISC-RT-000310 (Manual) 65
1.26 CISC-RT-000320 (Manual) 68
1.27 CISC-RT-000330 (Manual) 71
1.28 CISC-RT-000340 (Manual) 73
1.29 CISC-RT-000350 (Manual) 75
1.30 CISC-RT-000360 (Manual) 76
1.31 CISC-RT-000370 (Manual) 78
1.32 CISC-RT-000380 (Manual) 79
1.33 CISC-RT-000390 (Manual) 81
1.34 CISC-RT-000391 (Manual) 83
1.35 CISC-RT-000392 (Manual) 84
1.36 CISC-RT-000393 (Manual) 86
1.37 CISC-RT-000394 (Manual) 88
1.38 CISC-RT-000395 (Manual) 90
1.39 CISC-RT-000396 (Manual) 92
1.40 CISC-RT-000397 (Manual) 94
1.41 CISC-RT-000398 (Manual) 96
1.42 CISC-RT-000400 (Manual) 99
1.43 CISC-RT-000410 (Manual) 102
1.44 CISC-RT-000420 (Manual) 106
1.45 CISC-RT-000430 (Manual) 108
1.46 CISC-RT-000440 (Manual) 111
1.47 CISC-RT-000450 (Manual) 114
1.48 CISC-RT-000460 (Manual) 117
1.49 CISC-RT-000470 (Manual) 120
1.50 CISC-RT-000480 (Manual) 122
1.51 CISC-RT-000490 (Manual) 124
1.52 CISC-RT-000500 (Manual) 127
1.53 CISC-RT-000510 (Manual) 129
1.54 CISC-RT-000520 (Manual) 131
1.55 CISC-RT-000530 (Manual) 133
1.56 CISC-RT-000540 (Manual) 135
1.57 CISC-RT-000550 (Manual) 137
1.58 CISC-RT-000570 (Manual) 139
1.59 CISC-RT-000560 (Manual) 141
1.60 CISC-RT-000580 (Manual) 143
1.61 CISC-RT-000590 (Manual) 145
1.62 CISC-RT-000600 (Manual) 147
1.63 CISC-RT-000610 (Manual) 149
1.64 CISC-RT-000620 (Manual) 151
1.65 CISC-RT-000630 (Manual) 153
1.66 CISC-RT-000640 (Manual) 155
1.67 CISC-RT-000650 (Manual) 157
1.68 CISC-RT-000660 (Manual) 159
1.69 CISC-RT-000670 (Manual) 160
1.70 CISC-RT-000730 (Manual) 162
1.71 CISC-RT-000740 (Manual) 164
1.72 CISC-RT-000750 (Manual) 165
1.73 CISC-RT-000760 (Manual) 166
1.74 CISC-RT-000770 (Manual) 171
1.75 CISC-RT-000780 (Manual) 175
1.76 CISC-RT-000790 (Manual) 178

Internal Only - General

1.77 CISC-RT-000800 (Manual) 180
1.78 CISC-RT-000810 (Manual) 182
1.79 CISC-RT-000820 (Manual) 184
1.80 CISC-RT-000830 (Manual) 187
1.81 CISC-RT-000840 (Manual) 189
1.82 CISC-RT-000850 (Manual) 191
1.83 CISC-RT-000860 (Manual) 192
1.84 CISC-RT-000870 (Manual) 194
1.85 CISC-RT-000880 (Manual) 196
1.86 CISC-RT-000890 (Manual) 198
1.87 CISC-RT-000900 (Manual) 200
1.88 CISC-RT-000910 (Manual) 202
1.89 CISC-RT-000920 (Manual) 203
1.90 CISC-RT-000930 (Manual) 206
1.91 CISC-RT-000940 (Manual) 208
1.92 CISC-RT-000950 (Manual) 209

Appendix: Summary Table 210

Appendix: Change History 215

Internal Only - General

Overview

Target Technology Details

Cisco IOS Router RTR Secure Technical Implementation Guide (STIG)
Version: 2 Release: 6 Benchmark
Date: 25 Oct 2023

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 Router RTR and are looking to comply with the STIG guidance

Internal Only - General

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 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.

Internal Only - General

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.

Internal Only - General

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.

Internal Only - General

Acknowledgements

The Recommendations in this Benchmark are a representation of the Rules in the unclassified DISA STIG for Cisco IOS Router RTR.

Internal Only - General

Recommendations

1 STIG RULES

Cisco IOS XR router

Cisco IOS Router RTR Secure Technical Implementation Guide (STIG)

Version: 2 Release: 6 Benchmark

Date: 25 Oct 2023

CLASSIFICATION unclassified

Internal Only - General

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-216551
RULE ID: SV-216551r531085
```

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).

Internal Only - General

Audit:

This requirement is not applicable for the DODIN Backbone.
Review the router configuration to verify that Access Control Lists (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 GigabitEthernet1/1
description link to core
ip address 10.1.12.2 255.255.255.0
ip access-group FILTER_SERVER_TRAFFIC in
…
…
…
ip access-list extended FILTER_SERVER_TRAFFIC
permit tcp any 10.1.23.0 0.0.0.255 eq lpd 631 9100
permit tcp any 10.1.24.0 0.0.0.255 eq 1433 1434 4022
permit icmp any any
permit ospf any any
deny ip 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.

```
R5(config)#ip access-list extended FILTER_SERVER_TRAFFIC
R5(config-ext-nacl)#permit tcp any 10.1.23.0 0.0.0.255 eq 515 631 9100
R5(config-ext-nacl)#permit tcp any 10.1.24.0 0.0.0.255 eq 1433 1434 4022
R5(config-ext-nacl)#permit icmp any any
R5(config-ext-nacl)#permit ospf any any
R5(config-ext-nacl)#deny ip any any
R5(config-ext-nacl)#exit
R5(config)#interface GigabitEthernet1/1
R5(config-if)#ip access-group FILTER_SERVER_TRAFFIC in
```


Internal Only - General

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

Internal Only - General

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-216555
RULE ID: SV-216555r929040
```

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.

Internal Only - General

Audit:

Review the router configuration using the configuration examples below for BGP and OSPF.
EIGRP, RIP, and IS-IS only support MD5 and will incur a permanent finding for those protocols.
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> tcp
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
!
!
router bgp <ASN>
no synchronization
bgp log-neighbor-changes
neighbor x.x.x.x remote-as <ASN>
neighbor x.x.x.x ao <KEY-CHAIN-NAME>
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

Internal Only - General

```
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
…
…
…
interface GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
ip ospf authentication key-chain OSPF_KEY_CHAIN
```


Internal Only - General

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> tcp
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.
router bgp <ASN>
no synchronization
bgp log-neighbor-changes
neighbor x.x.x.x remote-as <ASN>
neighbor x.x.x.x ao <KEY-CHAIN-NAME>
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.
interface GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
ip ospf authentication key-chain OSPF_KEY_CHAIN

Internal Only - General

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)

Internal Only - General

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-216556
RULE ID: SV-216556r531085
```

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 GigabitEthernet3
shutdown
!
interface GigabitEthernet4
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

Internal Only - General

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

Internal Only - General

1.4 CISC-RT-000090 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216559
RULE ID: SV-216559r856180
```

Rationale:

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

Audit:

Review the device configuration to determine if auto-configuration or zero-touch deployment via Cisco Networking Services (CNS) is enabled.
Auto-configuration example
version 15.0
service config
…
…
…
boot-start-marker
boot network tftp://x.x.x.x/R5-config
boot-end-marker
CNS Zero-Touch Example
cns trusted-server config x.x.x.x
cns trusted-server image x.x.x.x
cns config initial x.x.x.x 80
cns exec 80
cns image
If a configuration auto-loading feature or zero-touch deployment feature is enabled, this is a finding.
Note: Auto-configuration or zero-touch deployment features can be enabled when the router is offline for the purpose of image loading or building out the configuration. In addition, this would not be applicable to the provisioning of virtual routers via a software-defined network (SDN) orchestration system.

Internal Only - General

Remediation:

Disable configuration auto-loading if enabled using the following commands.
R8(config)#no boot network
R8(config)#no service config
Disable CNS zero-touch deployment if enabled as shown in the example below.
R2(config)#no cns config initial
R2(config)#no cns exec
R2(config)#no cns image
R2(config)#no cns trusted-server config x.x.x.x
R2(config)#no cns trusted-server image x.x.x.x

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

Internal Only - General

1.5 CISC-RT-000120 (Manual)

Profile Applicability:

- SEVERITY: CAT I

Description:

The Cisco router must be configured to protect against or limit the effects of denial-of-service (DoS) attacks by employing control plane protection.

```
GROUP ID: V-216560
RULE ID: SV-216560r929043
```

Rationale:

The Route Processor (RP) is critical to all network operations because it is the component used to build all forwarding paths for the data plane via control plane processes. It is also instrumental with ongoing network management functions that keep the routers and links available for providing network services. Any disruption to the RP or the control and management planes can result in mission-critical network outages.

A DoS attack targeting the RP can result in excessive CPU and memory utilization. To maintain network stability and RP security, the router must be able to handle specific control plane and management plane traffic that is destined to the RP. In the past, one method of filtering was to use ingress filters on forwarding interfaces to filter both forwarding path and receiving path traffic, as well as limiting traffic destined to the device. However, this method does not scale well as the number of interfaces grows and the size of the ingress filters grows. Control plane policing increases the security of routers and multilayer switches by protecting the RP from unnecessary or malicious traffic. Filtering and rate limiting the traffic flow of control plane packets can be implemented to protect routers against reconnaissance and DoS attacks, allowing the control plane to maintain packet forwarding and protocol states despite an attack or heavy load on the router or multilayer switch.

Internal Only - General

Audit:

Review the Cisco router configuration to verify that it is compliant with this requirement.
Step 1: Verify traffic types have been classified based on importance levels. The following is an example configuration:
class-map match-all CoPP_CRITICAL
match access-group name CoPP_CRITICAL
class-map match-any CoPP_IMPORTANT
match access-group name CoPP_IMPORTANT
match protocol arp
class-map match-all CoPP_NORMAL
match access-group name CoPP_NORMAL
class-map match-any CoPP_UNDESIRABLE
match access-group name CoPP_UNDESIRABLE
class-map match-all CoPP_DEFAULT
match access-group name CoPP_DEFAULT
Step 2: Review the Access Control Lists (ACLs) referenced by the class maps to determine if the traffic is being classified appropriately. The following is an example configuration:
ip access-list extended CoPP_CRITICAL
remark our control plane adjacencies are critical
permit ospf host [OSPF neighbor A] any
permit ospf host [OSPF neighbor B] any
permit pim host [PIM neighbor A] any
permit pim host [PIM neighbor B] any
permit pim host [RP addr] any
permit igmp any 224.0.0.0 15.255.255.255
permit tcp host [BGP neighbor] eq bgp host [local BGP addr]
permit tcp host [BGP neighbor] host [local BGP addr] eq bgp
deny ip any any
ip access-list extended CoPP_IMPORTANT
permit tcp host [TACACS server] eq tacacs any
permit tcp [management subnet] 0.0.0.255 any eq 22
permit udp host [SNMP manager] any eq snmp
permit udp host [NTP server] eq ntp any
deny ip any any
ip access-list extended CoPP_NORMAL

Internal Only - General

remark we will want to rate limit ICMP traffic
deny icmp any host x.x.x.x fragments
permit icmp any any echo
permit icmp any any echo-reply
permit icmp any any time-exceeded
permit icmp any any unreachable
deny ip any any
ip access-list extended CoPP_UNDESIRABLE
remark other management plane traffic that should not be received
permit udp any any eq ntp
permit udp any any eq snmp
permit tcp any any eq 22
permit tcp any any eq 23
remark other control plane traffic not configured on router
permit eigrp any any
permit udp any any eq rip
deny ip any any
ip access-list extended CoPP_DEFAULT
permit ip any any
Note: Explicitly defining undesirable traffic with ACL entries enables the network operator to collect statistics. Excessive ARP packets can potentially monopolize Route Processor resources, starving other important processes. Currently, ARP is the only Layer 2 protocol that can be specifically classified using the match protocol command.
Step 3: Review the policy-map to determine if the traffic is being policed appropriately for each classification. The following is an example configuration:
policy-map CONTROL_PLANE_POLICY
class CoPP_CRITICAL
police 512000 8000 conform-action transmit exceed-action transmit
class CoPP_IMPORTANT
police 256000 4000 conform-action transmit exceed-action drop
class CoPP_NORMAL
police 128000 2000 conform-action transmit exceed-action drop
class CoPP_UNDESIRABLE
police 8000 1000 conform-action drop exceed-action drop
class CoPP_DEFAULT
police 64000 1000 conform-action transmit exceed-action drop
Step 4: Verify that the CoPP policy is enabled. The following is an example configuration:
control-plane
service-policy input CONTROL_PLANE_POLICY
Note: Control Plane Protection (CPPr) can be used to filter as well as police control plane traffic destined to the RP. CPPr is very similar to CoPP and has the ability to filter and police traffic using finer granularity by dividing the aggregate control plane into three separate categories: (1) host, (2) transit, and (3) CEF-exception. Hence, a separate policy-map could be configured for each traffic category.

Internal Only - General

If the Cisco router is not configured to protect against known types of DoS attacks by employing organization-defined security safeguards, this is a finding.

Remediation:

Configure the Cisco router to protect against known types of DoS attacks on the route processor. Implementing a CoPP policy as shown in the example below is a best practice method.
Step 1: Configure ACLs specific traffic types.

```
R1(config)#ip access-list extended CoPP_CRITICAL
R1(config-ext-nacl)#remark our control plane adjacencies are critical
R1(config-ext-nacl)#permit ospf host x.x.x.x any
R1(config-ext-nacl)#permit ospf host x.x.x.x any
R1(config-ext-nacl)#permit pim host x.x.x.x any
R1(config-ext-nacl)#permit pim host x.x.x.x any
R1(config-ext-nacl)#permit igmp any 224.0.0.0 15.255.255.255
R1(config-ext-nacl)#permit tcp host x.x.x.x eq bgp host x.x.x.x
R1(config-ext-nacl)#deny ip any any
R1(config-ext-nacl)#exit
R1(config)#ip access-list extended CoPP_IMPORTANT
R1(config-ext-nacl)#permit tcp host x.x.x.x eq tacacs any
R1(config-ext-nacl)#permit tcp x.x.x.x 0.0.0.255 any eq 22
R1(config-ext-nacl)#permit udp host x.x.x.x any eq snmp
R1(config-ext-nacl)#permit udp host x.x.x.x eq ntp any
R1(config-ext-nacl)#deny ip any any
R1(config-ext-nacl)#exit
R1(config)#ip access-list extended CoPP_NORMAL
R1(config-ext-nacl)#remark we will want to rate limit ICMP traffic
R1(config-ext-nacl)#deny icmp any host x.x.x.x fragments
R1(config-ext-nacl)#permit icmp any any echo
R1(config-ext-nacl)#permit icmp any any echo-reply
R1(config-ext-nacl)#permit icmp any any time-exceeded
R1(config-ext-nacl)#permit icmp any any unreachable
R1(config-ext-nacl)#deny ip any any
R1(config-ext-nacl)#exit
R1(config)#ip access-list extended CoPP_UNDESIRABLE
R1(config-ext-nacl)#remark management plane traffic that should not be received
R1(config-ext-nacl)#permit udp any any eq ntp
R1(config-ext-nacl)#permit udp any any eq snmp
R1(config-ext-nacl)#permit tcp any any eq 22
R1(config-ext-nacl)#permit tcp any any eq 23
R1(config-ext-nacl)#remark control plane traffic not configured on router
```


Internal Only - General

```
R1(config-ext-nacl)#permit eigrp any any
R1(config-ext-nacl)#permit udp any any eq rip
R1(config-ext-nacl)#deny ip any any
R1(config-ext-nacl)#exit
R1(config)#ip access-list extended CoPP_DEFAULT
R1(config-ext-nacl)#permit ip any any
R1(config-ext-nacl)#exit
```

Step 2: Configure class maps referencing each of the ACLs.

```
R1(config)#class-map match-all CoPP_CRITICAL
R1(config-cmap)#match access-group name CoPP_CRITICAL
R1(config-cmap)#class-map match-any CoPP_IMPORTANT
R1(config-cmap)#match access-group name CoPP_IMPORTANT
R1(config-cmap)#match protocol arp
R1(config-cmap)#class-map match-all CoPP_NORMAL
R1(config-cmap)#match access-group name CoPP_NORMAL
R1(config-cmap)#class-map match-any CoPP_UNDESIRABLE
R1(config-cmap)#match access-group name CoPP_UNDESIRABLE
R1(config-cmap)#class-map match-all CoPP_DEFAULT
R1(config-cmap)#match access-group name CoPP_DEFAULT
R1(config-cmap)#exit
```

Step 3: Configure a policy map referencing the configured class maps and apply appropriate bandwidth allowance and policing attributes.

```
R1(config)#policy-map CONTROL_PLANE_POLICY
R1(config-pmap)#class CoPP_CRITICAL
R1(config-pmap-c)#police 512000 8000 conform-action transmit exceed-action transmit
R1(config-pmap-c-police)#class CoPP_IMPORTANT
R1(config-pmap-c)#police 256000 4000 conform-action transmit exceed-action drop
R1(config-pmap-c-police)#class CoPP_NORMAL
R1(config-pmap-c)#police 128000 2000 conform-action transmit exceed-action drop
R1(config-pmap-c-police)#class CoPP_UNDESIRABLE
R1(config-pmap-c)#police 8000 1000 conform-action drop exceed-action drop
R1(config-pmap-c-police)#class CoPP_DEFAULT
R1(config-pmap-c)#police 64000 1000 conform-action transmit exceed-action drop
R1(config-pmap-c-police)#exit
R1(config-pmap-c)#exit
R1(config-pmap)#exit
```

Step 4: Apply the policy map to the control plane.

```
R1(config)#control-plane
R1(config-cp)#service-policy input CONTROL_PLANE_POLICY
R1(config-cp)#end
```


Internal Only - General

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

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

Internal Only - General

1.6 CISC-RT-000150 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

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

```
GROUP ID: V-216563
RULE ID: SV-216563r856182
```

Rationale:

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

Audit:

Review the configuration to determine if gratuitous ARP is disabled. The following command should not be found in the router configuration:
ip gratuitous-arps
Note: With Cisco IOS, Gratuitous ARP is enabled and disabled globally.
If gratuitous ARP is enabled on any external interface, this is a finding.

Remediation:

Disable gratuitous ARP as shown in the example below:
R5(config)#no ip gratuitous-arps

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

Internal Only - General

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-216564
RULE ID: SV-216564r856183
```

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).

Internal Only - General

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 GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
ip 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.
R4(config)#int g0/1
R4(config-if)#no ip directed-broadcast

Additional Information:

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

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

Internal Only - General

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-216565
RULE ID: SV-216565r856184
```

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 no ip unreachables command has been configured on all external interfaces as shown in the configuration example below.
interface GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
no ip unreachables
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 rate-limit unreachable 60000
ip icmp 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 DoS attack.
If the PE router is not configured to rate limit ICMP unreachable messages, this is a finding.

Internal Only - General

Remediation:

Step 1: Disable ip unreachables on all external interfaces.
R4(config)#int g0/1
R4(config-if)#no ip unreachables
Step 2: Disable ip unreachables on the Null0 interface if it is used to backhole packets.
R4(config-if)#int null 0
R4(config-if)#no ip unreachables
Alternative – DODIN Backbone:
Configure the PE router to rate limit ICMP unreachable messages as shown in the example below:
R4(config)#ip icmp rate-limit unreachable df 100
R4(config)#ip icmp rate-limit unreachable 100000
R4(config)#end
Alternative – Non DODIN Backbone.
An alternative for non-backbone networks (i.e. enclave, base, camp, etc.) is to filter messages generated by the router and silently drop ICMP Administratively Prohibited and Host Unreachable messages using the following configuration steps:
Step 1: Configure ACL to include ICMP Type 3 Code 1 (Host Unreachable) and Code 13 (Administratively Prohibited) as shown in the example below:
R2(config)#ip access-list ext ICMP_T3C1C13
R2(config-ext-nacl)#permit icmp any any host-unreachable
R2(config-ext-nacl)#permit icmp any any administratively-prohibited
R2(config-ext-nacl)#exit
Step 2: Create a route map to forward these ICMP messages to the Null0 interface.
R2(config)#route-map LOCAL_POLICY
R2(config-route-map)#match ip address ICMP_T3C1C13
R2(config-route-map)#set interface Null0
R2(config-route-map)#exit
Step 3: Configure no ip unreachables on the Null0 interface.
R2(config)#int null 0
R2(config-if)#no ip unreachables
R2(config-if)#exit
Step 4: Apply the policy to filter messages generated by the router.
R2(config)#ip local policy route-map LOCAL_POLICY
R2(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

Internal Only - General

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-216566
RULE ID: SV-216566r856185
```

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 ip mask-reply command is not enabled on any external interfaces as shown in the example below.
interface GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
ip mask-reply
If the ip mask-reply command is configured on any external interface, this is a finding.

Remediation:

Disable ip mask-reply on all external interfaces as shown below.
R4(config)#int g0/1
R4(config-if)#no ip 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

Internal Only - General

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-216567
RULE ID: SV-216567r856186
```

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 the no ip redirects command has been configured on all external interfaces as shown in the example below.
interface GigabitEthernet0/1
ip address x.x.x.x 255.255.255.0
no ip 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.
R4(config)#int g0/1
R4(config-if)#no ip 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

Internal Only - General

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 an ACL.

```
GROUP ID: V-216568
RULE ID: SV-216568r531085
```

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 Access Control Lists(ACLs) used to filter traffic and verify that packets being dropped are logged as shown in the configuration below.
ip access-list extended INGRESS_FILTER
permit tcp any any established
permit tcp host x.11.1.1 eq bgp host x.11.1.2
permit tcp host x.11.1.1 host x.11.1.2 eq bgp
permit tcp any host x.11.1.5 eq www
permit icmp host x.11.1.1 host x.11.1.2 echo
permit icmp any any echo-reply
…
…
…
deny ip 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.
R5(config)#ip access-list extended INGRESS_FILTER
…
…
…
R5(config-ext-nacl)#deny ip any any log

Internal Only - General

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

Internal Only - General

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-216569
RULE ID: SV-216569r531085
```

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.
ip access-list extended INGRESS_FILTER
permit tcp any any established
permit tcp host x.11.1.1 eq bgp host x.11.1.2
permit tcp host x.11.1.1 host x.11.1.2 eq bgp
permit tcp any host x.11.1.5 eq www
permit icmp host x.11.1.1 host x.11.1.2 echo
permit icmp any any echo-reply
…
…
…
deny ip 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.

Internal Only - General

Remediation:

Configure the router to log events containing information to establish where the events occurred as shown in the example below.
R5(config)#ip access-list extended INGRESS_FILTER
…
…
…
R5(config-ext-nacl)#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

Internal Only - General

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-216570
RULE ID: SV-216570r531085
```

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.
ip access-list extended INGRESS_FILTER
permit tcp any any established
permit tcp host x.11.1.1 eq bgp host x.11.1.2
permit tcp host x.11.1.1 host x.11.1.2 eq bgp
permit tcp any host x.11.1.5 eq www
permit icmp host x.11.1.1 host x.11.1.2 echo
permit icmp any any echo-reply
…
…
…
deny ip 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.

Internal Only - General

Remediation:

Configure the router to log events containing information to establish where the events occurred as shown in the example below.
R5(config)#ip access-list extended INGRESS_FILTER
…
…
…
R5(config-ext-nacl)#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

Internal Only - General

1.14 CISC-RT-000230 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco router must be configured to disable the auxiliary port unless it is connected to a secured modem providing encryption and authentication.

```
GROUP ID: V-216571
RULE ID: SV-216571r531085
```

Rationale:

The use of POTS lines to modems connecting to network devices provides clear text of authentication traffic over commercial circuits that could be captured and used to compromise the network. Additional war dial attacks on the device could degrade the device and the production network.

Secured modem devices must be able to authenticate users and must negotiate a key exchange before full encryption takes place. The modem will provide full encryption capability (Triple DES) or stronger. The technician who manages these devices will be authenticated using a key fob and granted access to the appropriate maintenance port; thus, the technician will gain access to the managed device (router, switch, etc.). The token provides a method of strong (two-factor) user authentication. The token works in conjunction with a server to generate one-time user passwords that will change values at second intervals. The user must know a personal identification number (PIN) and possess the token to be allowed access to the device.

Audit:

Review the configuration and verify that the auxiliary port is disabled unless a secured modem providing encryption and authentication is connected to it.
line aux 0
no exec
Note: transport input none is the default, hence it will not be shown in the configuration.
If the auxiliary port is not disabled or is not connected to a secured modem when it is enabled, this is a finding.

Remediation:

Disable the auxiliary port.
R2(config)#line aux 0
R2(config-line)#no exec
R2(config-line)#transport input none

Internal Only - General

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

Internal Only - General

1.15 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-229030
RULE ID: SV-229030r878127
```

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

Internal Only - General

1.16 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-230038
RULE ID: SV-230038r531386
```

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

Internal Only - General

1.17 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-230041
RULE ID: SV-230041r532998
```

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

Internal Only - General

1.18 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-216572
RULE ID: SV-216572r531085
```

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.

Internal Only - General

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 GigabitEthernet0/2
ip address x.11.1.2 255.255.255.254
ip access-group EXTERNAL_ACL in
Step 2: Review inbound ACL to verify that it is configured to deny all other traffic that is not explicitly allowed.
ip access-list extended EXTERNAL_ACL
permit tcp any any established
permit tcp host x.11.1.1 eq bgp host x.11.1.2
permit tcp host x.11.1.1 host x.11.1.2 eq bgp
permit icmp host x.11.1.1 host x.11.1.2 echo
permit icmp host x.11.1.1 host x.11.1.2 echo-reply
…
…
…
deny ip 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.
R1(config)#ip access-list extended EXTERNAL_ACL
R1(config-ext-nacl)#permit tcp any any established
R1(config-ext-nacl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

R1(config-ext-nacl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp
R1(config-ext-nacl)#permit icmp host x.11.1.1 host x.11.1.2 echo
R1(config-ext-nacl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply
…
…
…
R1(config-ext-nacl)#deny ip any any log-input
Step 2: Apply the ingress filter to all external interfaces
R1(config)#int g0/2
R1(config-if)#ip access-group EXTERNAL_ACL in

Internal Only - General

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)

Internal Only - General

1.19 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-216573
RULE ID: SV-216573r531085
```

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).
interface GigabitEthernet0/1
description Link to DISN
ip address x.12.1.10 255.255.255.0
ip access-group FILTER_PERIMETER in
…
…
…
ip access-list extended FILTER_PERIMETER
permit tcp any any established
permit tcp host x.12.1.9 host x.12.1.10 eq bgp
permit tcp host x.12.1.9 eq bgp host x.12.1.10
permit icmp host x.12.1.9 host x.12.1.10 echo
permit icmp host x.12.1.9 host x.12.1.10 echo-reply
permit tcp any host x.12.1.22 eq www
deny ip any any log-input
If the router is not configured to enforce approved authorizations for controlling the flow of information between interconnected networks, this is a finding.

Internal Only - General

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.
R1(config)#ip access-list extended FILTER_PERIMETER
R1(config-ext-nacl)#permit tcp any any established
R1(config-ext-nacl)#permit tcp host x.12.1.9 host x.12.1.10 eq bgp
R1(config-ext-nacl)#permit tcp host x.12.1.9 eq bgp host x.12.1.10
R1(config-ext-nacl)#permit icmp host x.12.1.9 host x.12.1.10 echo
R1(config-ext-nacl)#permit icmp host x.12.1.9 host x.12.1.10 echo-reply
R1(config-ext-nacl)#permit tcp any host x.12.1.22 eq www
R1(config-ext-nacl)#deny ip any any log-input
R1(config-ext-nacl)#exit
Step 2: Apply the ACL inbound on all external interfaces.
R2(config)#int g0/0
R1(config-if)#ip access-group FILTER_PERIMETER in

Additional Information:

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

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

Internal Only - General

1.20 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-216574
RULE ID: SV-216574r856187
```

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.

Internal Only - General

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.
ip access-list extended FILTER_PERIMETER
permit tcp any any established
…
…
…
permit udp host x.12.1.9 host x.12.1.21 eq ntp
deny ip 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.
R1(config)#ip access-list extended FILTER_PERIMETER
R1(config-ext-nacl)#nn permit udp host x.12.1.9 host x.12.1.21 eq ntp
R1(config-ext-nacl)#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)

Internal Only - General

1.21 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-216575
RULE ID: SV-216575r863237
```

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.

Internal Only - General

Audit:

This requirement is not applicable for the DODIN Backbone.
Review the router configuration to verify that an ingress Access Control List (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.

```
ip access-list extended FILTER_PERIMETER
deny ip 0.0.0.0 0.255.255.255 any log-input
deny ip 10.0.0.0 0.255.255.255 any log-input
deny ip 100.64.0.0 0.63.255.255 any log-input
deny ip 127.0.0.0 0.255.255.255 any log-input
deny ip 169.254.0.0 0.0.255.255 any log-input
deny ip 172.16.0.0 0.15.255.255 any log-input
deny ip 192.0.0.0 0.0.0.255 any log-input
deny ip 192.0.2.0 0.0.0.255 any log-input
deny ip 192.168.0.0 0.0.255.255 any log-input
deny ip 198.18.0.0 0.1.255.255 any log-input
deny ip 198.51.100.0 0.0.0.255 any log-input
deny ip 203.0.113.0 0.0.0.255 any log-input
deny ip 224.0.0.0 31.255.255.255 any log-input
deny ip 240.0.0.0 15.255.255.255 any log-input
permit tcp any any established
permit tcp host x.12.1.9 host x.12.1.10 eq bgp
permit tcp host x.12.1.9 eq bgp host x.12.1.10
permit icmp host x.12.1.9 host x.12.1.10 echo
permit icmp host x.12.1.9 host x.12.1.10 echo-reply
…
…
…
deny ip any any log-input
```

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

```
interface GigabitEthernet0/1
description Link to DISN
ip address x.12.1.10 255.255.255.254
ip access-group FILTER_PERIMETER in
```

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

Internal Only - General

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.

```
R5(config)#ip access-list extended FILTER_PERIMETER
R5(config-ext-nacl)#deny ip 0.0.0.0 0.255.255.255 any log-input
R5(config-ext-nacl)#deny ip 10.0.0.0 0.255.255.255 any log-input
R5(config-ext-nacl)#deny ip 100.64.0.0 0.63.255.255 any log-input
R5(config-ext-nacl)#deny ip 127.0.0.0 0.255.255.255 any log-input
R5(config-ext-nacl)#deny ip 169.254.0.0 0.0.255.255 any log-input
R5(config-ext-nacl)#deny ip 172.16.0.0 0.15.255.255 any log-input
R5(config-ext-nacl)#deny ip 192.0.0.0 0.0.0.255 any log-input
R5(config-ext-nacl)#deny ip 192.0.2.0 0.0.0.255 any log-input
R5(config-ext-nacl)#deny ip 192.168.0.0 0.0.255.255 any log-input
R5(config-ext-nacl)#deny ip 198.18.0.0 0.1.255.255 any log-input
R5(config-ext-nacl)#deny ip 198.51.100.0 0.0.0.255 any log-input
R5(config-ext-nacl)#deny ip 203.0.113.0 0.0.0.255 any log-input
R5(config-ext-nacl)#deny ip 224.0.0.0 31.255.255.255 any log-input
R5(config-ext-nacl)#deny ip 240.0.0.0 15.255.255.255 any log-input
R5(config-ext-nacl)#permit tcp any any established
R5(config-ext-nacl)#permit tcp host x.12.1.9 host x.12.1.10 eq bgp
R5(config-ext-nacl)#permit tcp host x.12.1.9 eq bgp host x.12.1.10
R5(config-ext-nacl)#permit icmp host x.12.1.9 host x.12.1.10 echo
R5(config-ext-nacl)#permit icmp host x.12.1.9 host x.12.1.10 echo-reply
…
…
…
R5(config-ext-nacl)#deny ip any any log-input
R5(config-ext-nacl)#end
```

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

```
R2(config)#int g0/0
R1(config-if)#ip access-group FILTER_PERIMETER in
R1(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)

Internal Only - General

1.22 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-216576
RULE ID: SV-216576r531085
```

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.

Internal Only - General

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 GigabitEthernet0/2
description Link to ISP
ip address x.22.1.15 255.255.255.240
ip access-group FILTER_ISP in
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.
ip access-list extended FILTER_ISP
permit tcp any any established
permit icmp host x.12.1.16 host x.12.1.17 echo
permit icmp host x.12.1.16 host x.12.1.17 echo-reply
permit tcp any host x.12.1.22 eq www
permit tcp any host x.12.1.23 eq www
permit 50 any host x.12.1.24
permit 51 any host x.12.1.24
deny ip 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.

Internal Only - General

Remediation:

This requirement is not applicable for the DODIN Backbone.
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.

```
R5(config)#ip access-list extended FILTER_ISP
R5(config-ext-nacl)#permit tcp any any established
R5(config-ext-nacl)#permit icmp host x.12.1.16 host x.12.1.17 echo
R5(config-ext-nacl)#permit icmp host x.12.1.16 host x.12.1.17 echo-reply
R5(config-ext-nacl)#permit tcp any host x.12.1.22 eq www
R5(config-ext-nacl)#permit tcp any host x.12.1.23 eq www
R5(config-ext-nacl)#permit 50 any host x.12.1.24
R5(config-ext-nacl)#permit 51 any host x.12.1.24
R5(config-ext-nacl)#deny ip any any log-input
R5(config-ext-nacl)#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

Internal Only - General

1.23 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-216577
RULE ID: SV-216577r531085
```

Rationale:

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

Audit:

This requirement is not applicable for the DODIN Backbone.
Review the router configuration and verify that it is not BGP peering with an alternate gateway service provider.
Step 1: Determine the ip address of the ISP router
interface GigabitEthernet0/2
description Link to ISP
ip address x.22.1.15 255.255.255.240
Step 2: Verify that the router is not BGP peering with this router.
router bgp nn
no synchronization
bgp log-neighbor-changes
neighbor x.11.1.7 remote-as nn
neighbor x.11.1.7 password xxxxxxx
no auto-summary
In the example above, the router is not peering with the ISP.
If the router is BGP peering with an alternate gateway service provider, 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

Internal Only - General

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

Internal Only - General

1.24 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-216578
RULE ID: SV-216578r531085
```

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.

Internal Only - General

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
log-adjacency-changes
redistribute static subnets
network 0.0.0.0 255.255.255.255 area 0
EIGRP example
router eigrp 1
network 10.1.15.0 0.0.0.255
redistribute static
RIP example
router rip
version 2
redistribute static
network 10.0.0.0
BGP example
router bgp nn
no synchronization
bgp log-neighbor-changes
redistribute static
neighbor x.11.1.7 remote-as nn
neighbor x.11.1.7 password xxxxxxx
no auto-summary
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.

Internal Only - General

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-map 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
R5(config)#ip prefix-list ISP_PREFIX permit x.x.x.0/24
Step 2: Configure a route map that will deny the state routes to the ISP
R5(config)#route-map FILTER_ISP_STATIC deny 10
R5(config-route-map)#match ip address prefix-list ISP_PREFIX
R5(config-route-map)#exit
R5(config)#route-map FILTER_ISP_STATIC permit 20
R5(config-route-map)#exit
Step 3: Apply the route-map to the IGP and BGP redistribute static commands as shown in the EIGRP example.
R5(config)#router eigrp 1
R5(config-router)#redistribute static route-map FILTER_ISP_STATIC

Additional Information:

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

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

Internal Only - General

1.25 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-216989
RULE ID: SV-216989r531085
```

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.

Internal Only - General

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 GigabitEthernet0/1
description downstream link to LAN
ip address 10.1.25.5 255.255.255.0
ip verify unicast source reachable-via rx
Egress ACL example
interface GigabitEthernet0/1
description downstream link to LAN
ip address 10.1.25.5 255.255.255.0
ip access-group EGRESS_FILTER in
…
…
…
ip access-list extended EGRESS_FILTER
permit udp 10.1.15.0 0.0.0.255 any eq domain
permit tcp 10.1.15.0 0.0.0.255 any eq ftp
permit tcp 10.1.15.0 0.0.0.255 any eq ftp-data
permit tcp 10.1.15.0 0.0.0.255 any eq www
permit icmp 10.1.15.0 0.0.0.255 any
permit icmp 10.1.15.0 0.0.0.255 any echo
deny ip any any
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.
R5(config)#int g0/1
R5(config-if)#ip verify unicast source reachable-via rx

Internal Only - General

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)

Internal Only - General

1.26 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-216580
RULE ID: SV-216580r531085
```

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.

If the perimeter is in a Deny-by-Default posture and what is allowed through the filter is in accordance with DoD Instruction 8551.1, and if the permit rule is explicitly defined with explicit ports and protocols allowed, then all requirements related to PPS being blocked would be satisfied.

Internal Only - General

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 GigabitEthernet0/2
ip address x.11.1.2 255.255.255.254
ip access-group EXTERNAL_ACL_INBOUND in
Step 2. Review the inbound ACL to verify that it is filtering traffic in accordance with DoD 8551.1.
ip access-list extended EXTERNAL_ACL_INBOUND
permit tcp any any established
permit tcp host x.11.1.1 eq bgp host x.11.1.2
permit tcp host x.11.1.1 host x.11.1.2 eq bgp
permit icmp host x.11.1.1 host x.11.1.2 echo
permit icmp host x.11.1.1 host x.11.1.2 echo-reply
…
… < must be in accordance with DoD Instruction 8551.1>
…
deny ip 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.
R1(config)#ip access-list extended EXTERNAL_ACL_INBOUND
R1(config-ext-nacl)#permit tcp any any established
R1(config-ext-nacl)#permit tcp host x.11.1.1 eq bgp host x.11.1.2

R1(config-ext-nacl)#permit tcp host x.11.1.1 host x.11.1.2 eq bgp
R1(config-ext-nacl)#permit icmp host x.11.1.1 host x.11.1.2 echo
R1(config-ext-nacl)#permit icmp host x.11.1.1 host x.11.1.2 echo-reply
…
… < must be in accordance with DoD Instruction 8551.1>
…
R1(config-ext-nacl)#deny ip any any log-input
R1(config-ext-nacl)#exit
R1(config)#int g0/2
R1(config-if)#ip access-group EXTERNAL_ACL_INBOUND in

Internal Only - General

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

Internal Only - General

1.27 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-216581
RULE ID: SV-216581r531085
```

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 GigabitEthernet0/2
ip address x.11.1.2 255.255.255.254
ip access-group EXTERNAL_ACL_INBOUND in
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.
R1(config)#int g0/2
R1(config-if)#ip access-group EXTERNAL_ACL_INBOUND in

Internal Only - General

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

Internal Only - General

1.28 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-216582
RULE ID: SV-216582r531085
```

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 interface GigabitEthernet0/2
description downstream link to LAN
ip address 10.1.25.5 255.255.255.0
ip access-group EGRESS_FILTER in
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.
R5(config)#int g0/2
R5(config-if)#ip access-group EGRESS_FILTER in

Internal Only - General

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

Internal Only - General

1.29 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-216990
RULE ID: SV-216990r856207
```

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.
Review the router configuration to determine if it will block all packets with IP options.
ip access-list extended EXTERNAL_ACL
permit tcp any any established
deny ip any any option any-options
permit …
…
…

…
deny ip any any log-input
If the router is not configured to drop all packets with IP options, this is a finding.

Remediation:

This requirement is not applicable for the DODIN Backbone.
Configure the router to drop all packets with IP options.
R1(config)#ip access-list extended EXTERNAL_ACL
R1(config-ext-nacl)#15 deny ip any any option any-options

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)

Internal Only - General

1.30 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-216584
RULE ID: SV-216584r856189
```

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 run
By default LLDP is not enabled globally. If LLDP is enabled, proceed to step 2.
Step 2: Verify LLDP is not enabled on any external interface as shown in the example below.
interface GigabitEthernet0/1
ip address x.1.12.1 255.255.255.252
no lldp transmit
Note: LLDP is enabled by default on all interfaces once it is enabled globally; hence the command lldp transmit 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.
R5(config)#int g0/1
R5(config-if)#no lldp transmit

Internal Only - General

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)

Internal Only - General

1.31 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-216585
RULE ID: SV-216585r856190
```

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 CDP is not enabled globally via the command no cdp run
By default CDP is enabled 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 GigabitEthernet0/1
ip address x.1.23.2 255.255.255.252
no cdp enable
Note: By default CDP is enabled on all interfaces if CDP is enabled globally.
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 enable command or disable CDP globally via no cdp run 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)

Internal Only - General

1.32 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-216586
RULE ID: SV-216586r856191
```

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 disabled on all external interfaces as shown in the example below.
interface GigabitEthernet0/1
description link to DISN
ip address x.1.12.2 255.255.255.252
no ip proxy-arp
Note: By default Proxy ARP is enabled on all interfaces; hence, if enabled, it will not be shown in the configuration.
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.
R2(config)#int g0/1

R2(config-if)#no ip proxy-arp

Internal Only - General

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)

Internal Only - General

1.33 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-216587
RULE ID: SV-216587r531085
```

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 have been configured with an outbound ACL as shown in the example below.
interface GigabitEthernet0/2
description link to DISN
ip address x.11.1.2 255.255.255.254
ip access-group EXTERNAL_ACL_OUTBOUND out
Step 2: Verify that the outbound ACL discards management traffic as shown in the example below.
ip access-list extended EXTERNAL_ACL_OUTBOUND
deny tcp any any eq tacacs log-input
deny tcp any any eq 22 log-input
deny udp any any eq snmp log-input
deny udp any any eq snmptrap log-input
deny udp any any eq syslog log-input
permit tcp any any eq www log-input
deny ip any any log-input
If management traffic is not blocked at the perimeter, this is a finding.

Internal Only - General

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.
R5(config)#ip access-list extended EXTERNAL_ACL_OUTBOUND
R5(config-ext-nacl)#deny tcp any any eq tacacs log-input
R5(config-ext-nacl)#deny tcp any any eq 22 log-input
R5(config-ext-nacl)#deny udp any any eq snmp log-input
R5(config-ext-nacl)#deny udp any any eq snmptrap log-input
R5(config-ext-nacl)#deny udp any any eq syslog log-input
R5(config-ext-nacl)#permit tcp any any eq www
R5(config-ext-nacl)#deny ip any any log-input
R5(config-ext-nacl)#exit
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.
R1(config)#int g0/2
R1(config-if)#ip access-group EXTERNAL_ACL_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

Internal Only - General

1.34 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-230044
RULE ID: SV-230044r533005
```

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

Internal Only - General

1.35 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-230047
RULE ID: SV-230047r856663
```

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
Step 2: Verify that the ACL drops undetermined transport packets as shown in the example below.
ipv6 access-list FILTER_IPV6
deny ipv6 any any log undetermined-transport
permit ipv6 …
…
…
…
deny ipv6 any any log
If the router is not configured to drop IPv6 undetermined transport packets, this is a finding.

Internal Only - General

Remediation:

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

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny ipv6 any any undetermined-transport log
R1(config-ipv6-acl)#permit ipv6 …
…
…
…
R1(config-ipv6-acl)#deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6 in
```

Additional Information:

CCI-002403 Only allow incoming communications from organization-defined authorized sources routed to organization-defined authorized destinations.

- NIST SP 800-53 Revision 4::SC-7 (11)
- NIST SP 800-53 Revision 5::SC-7 (11)

Internal Only - General

1.36 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-230050
RULE ID: SV-230050r856665
```

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.

Internal Only - General

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
permit ipv6 any host 2001:DB8::1:1:1234 routing-type 2
deny ipv6 any any log routing
permit ipv6 …
…
…
…
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.
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#permit ipv6 any host 2001:DB8::0:1:1:1234 routing-type 2
R1(config-ipv6-acl)#deny ipv6 any any routing log
R1(config-ipv6-acl)#permit …
…
…
…
R1(config-ipv6-acl)#deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6

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)

Internal Only - General

1.37 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-230145
RULE ID: SV-230145r856667
```

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
deny hbh any any dest-option-type 4 log
deny hbh any any dest-option-type 195 log
deny hbh any any dest-option-type home-address log
permit ipv6 …
…
…
…
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.

Internal Only - General

Remediation:

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

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny hbh any any dest-option-type 4 log
R1(config-ipv6-acl)#deny hbh any any dest-option-type 195 log
R1(config-ipv6-acl)#deny hbh any any dest-option-type home-address log
R1(config-ipv6-acl)# permit ipv6 …
…
…
…
R1(config-ipv6-acl)#deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6
R1(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)

Internal Only - General

1.38 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-230149
RULE ID: SV-230149r856669
```

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
deny 60 any any dest-option-type 5 log
deny 60 any any dest-option-type 194 log
permit ipv6 …
…
…
…
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.

Internal Only - General

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.

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny 60 any any dest-option-type 5 log
R1(config-ipv6-acl)#deny 60 any any dest-option-type 194 log
R1(config-ipv6-acl)#permit …
…
…
…
R1(config-ipv6-acl)#deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6
R1(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)

Internal Only - General

1.39 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-230152
RULE ID: SV-230152r856671
```

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
deny any any dest-option-type 138 log
permit ipv6 …
…
…
…
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.

Internal Only - General

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.

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny any any dest-option-type 138 log
R1(config-ipv6-acl)#permit ipv6 …
…
…
…
R1(config-ipv6-acl)# deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6
R1(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)

Internal Only - General

1.40 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-230155
RULE ID: SV-230155r856673
```

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 a Destination Option header with option type value of 0xC3 (NSAP address).
Step 1: Verify that an inbound IPv6 ACL has been configured on the external interface.
interface gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
deny 60 any any dest-option-type 195 log
permit ipv6 …
…
…
…
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.

Internal Only - General

Remediation:

Configure the router to to drop IPv6 packets containing the NSAP address option within Destination Option header as shown in the example below.

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny 60 any any dest-option-type 195 log
R1(config-ipv6-acl)#permit …
…
…
…
R1(config-ipv6-acl)# deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6
R1(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)

Internal Only - General

1.41 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-230158
RULE ID: SV-230158r856675
```

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.

Internal Only - General

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 gigabitethernet1/0
ipv6 address 2001::1:0:22/64
ipv6 traffic-filter FILTER_IPV6 in
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
deny any any dest-option-type 2
deny any any dest-option-type 3
deny any any dest-option-type 6
deny any any dest-option-type 9
deny any any dest-option-type 10
deny any any dest-option-type 11
deny any any dest-option-type 12
deny any any dest-option-type 13
deny any any dest-option-type 14
deny any any dest-option-type 16
…
deny any any dest-option-type 34
deny any any dest-option-type 36
deny any any dest-option-type 37
deny any any dest-option-type 39
…
deny any any dest-option-type 47
deny any any dest-option-type 49
…
deny any any dest-option-type 255
permit …
…
…
…
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.

Internal Only - General

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.

```
R1(config)#ipv6 access-list FILTER_IPV6
R1(config-ipv6-acl)#deny any any dest-option-type 2
R1(config-ipv6-acl)#deny any any dest-option-type 3
R1(config-ipv6-acl)#deny any any dest-option-type 6
R1(config-ipv6-acl)#deny any any dest-option-type 9
R1(config-ipv6-acl)#deny any any dest-option-type 10
R1(config-ipv6-acl)#deny any any dest-option-type 11
R1(config-ipv6-acl)#deny any any dest-option-type 12
R1(config-ipv6-acl)#deny any any dest-option-type 13
R1(config-ipv6-acl)#deny any any dest-option-type 14
R1(config-ipv6-acl)#deny any any dest-option-type 16
…
R1(config-ipv6-acl)#deny any any dest-option-type 34
R1(config-ipv6-acl)#deny any any dest-option-type 36
R1(config-ipv6-acl)#deny any any dest-option-type 37
R1(config-ipv6-acl)#deny any any dest-option-type 39
…
R1(config-ipv6-acl)#deny any any dest-option-type 47
R1(config-ipv6-acl)#deny any any dest-option-type 49
…
R1(config-ipv6-acl)#deny any any dest-option-type 255
R1(config-ipv6-acl)#permit …
…
…
…
R1(config-ipv6-acl)#deny ipv6 any any log
R1(config-ipv6-acl)#exit
R1(config)#int g1/0
R1(config-if)#ipv6 traffic-filter FILTER_IPV6
```

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)

Internal Only - General

1.42 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-216588
RULE ID: SV-216588r531085
```

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.

Internal Only - General

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 crypto map applied to the external interface
interface interface GigabitEthernet0/2
description link to DISN
ip address x.1.24.4 255.255.255.0
crypto map IPSEC_MGMT_MAP
Step 2: Review the ISAKMP policy for Phase 1 negotiations and Phase 2 policy for data encryption
crypto isakmp policy 10
authentication pre-share
hash sha256
crypto isakmp key xxxxxx address x.1.12.1
!
!
crypto IPsec transform-set TRANS_SET ah-sha256-hmac esp-aes
Step 3: Review the crypto map that was bound to the external interface and note the ACL defined that identifies the interesting traffic for the IPsec tunnel.
crypto map IPSEC_MGMT_MAP 10 ipsec-isakmp
set peer x.1.12.1
set transform-set TRANS_SET
match address MGMT_TRAFFIC_ACL
Step 4: Review the ACL defined in the crypto map and verify that the destination is the management network.
ip access-list extended MGMT_TRAFFIC_ACL
permit ip 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.

Internal Only - General

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.
R4(config)#ip access-list extended MGMT_TRAFFIC_ACL
R4(config-ext-nacl)#permit ip 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255
R4(config-ext-nacl)#exit
Step 2: Create an ISAKMP policy for Phase 1 negotiations.
R4(config)#crypto isakmp policy 10
R4(config-isakmp)#hash sha256
R4(config-isakmp)#authentication pre-share
R4(config-isakmp)#exit
Step 3: Specify the pre-shared key and the remote peer address.
R4(config)#crypto isakmp key 0 xxxxxx address x.1.12.1
Note: Digital certificates can be utilized as an alternative.
Step 4: Create the IPSec transform set for the data encryption
R4(config)#crypto ipsec transform-set TRANS_SET ah-sha256-hmac esp-aes
R4(cfg-crypto-trans)#mode tunnel
R4(cfg-crypto-trans)#exit
Step 5: Create the crypto map.
R4(config)#crypto map IPSEC_MGMT_MAP 10 ipsec-isakmp
R4(config-crypto-map)#set peer x.1.12.1
R4(config-crypto-map)#match address MGMT_TRAFFIC_ACL
R4(config-crypto-map)#set transform-set TRANS_SET
R4(config-crypto-map)#end
Step 6: Apply the crypto map to the external interface.
R4(config)#int g0/2
R4(config-if)#crypto map IPSEC_MGMT_MAP

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

Internal Only - General

1.43 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-216589
RULE ID: SV-216589r531085
```

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.

Internal Only - General

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 GigabitEthernet0/2
description OOB link to NOC
ip address 10.11.1.8 255.255.255.0
ip access-group MGMT_TRAFFIC_ACL out
Step 2: Review the outbound ACL and verify only management traffic is forwarded to the NOC.
ip access-list extended MGMT_TRAFFIC_ACL
permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs
permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq 22
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp-trap
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog
permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255
deny ip 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 crypto map applied to the external interface.
interface interface GigabitEthernet0/2
description link to DISN
ip address x.1.24.4 255.255.255.0
crypto map IPSEC_MGMT_MAP
Step 2: Review the crypto map that was bound to the external interface and note the ACL defined that identifies the interesting traffic for the IPsec tunnel.
crypto map IPSEC_MGMT_MAP 10 ipsec-isakmp
set peer x.1.12.1
set transform-set TRANS_SET
match address MGMT_TRAFFIC_ACL

Internal Only - General

Step 3: Review the ACL defined in the crypto map and verify only management traffic is forwarded to the NOC.
ip access-list extended MGMT_TRAFFIC_ACL
permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs
permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq 22
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp-trap
permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog
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
R4(config)#ip access-list extended MGMT_TRAFFIC_ACL
R4(config-ext-nacl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs
R4(config-ext-nacl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq 22
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp-trap
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog
R4(config-ext-nacl)#permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255 echo
R4(config-ext-nacl)#permit icmp 10.1.34.0 0.0.0.255 10.22.22.0 0.0.0.255 echo-reply
R4(config-ext-nacl)#deny ip any any log-input
R4(config-ext-nacl)#exit
IPsec Tunnel
R4(config)#ip access-list extended MGMT_TRAFFIC_ACL
R4(config-ext-nacl)#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq tacacs
R4(config-ext-nacl#permit tcp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq 22
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq snmp-trap
R4(config-ext-nacl)#permit udp 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255 eq syslog
R4(config-ext-nacl)#permit icmp 10.1.34.0 0.0.0.255 22.22.22.0 0.0.0.255
R4(config-ext-nacl)#exit

Internal Only - General

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

Internal Only - General

1.44 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-216590
RULE ID: SV-216590r531085
```

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 1 vrf MGMT
log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0
!
router ospf 2 vrf PROD
log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0
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.
R3(config)#router ospf 1 vrf MGMT
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#exit
R3(config)#router ospf 2 vrf PROD
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#end

Internal Only - General

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

Internal Only - General

1.45 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-216591
RULE ID: SV-216591r531085
```

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.

Internal Only - General

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.

```
ip vrf MGMT
rd 4:4
route-target export 4:4
route-target import 4:4
route-target import 8:8
!
ip vrf PROD
rd 8:8
route-target import 8:8
route-target export 8:8
…
…
…
router ospf 1 vrf MGMT
log-adjacency-changes
redistribute bgp 64512 subnets
network 0.0.0.0 255.255.255.255 area 0
!
router ospf 2 vrf PROD
log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0
!
router bgp 64512
no synchronization
bgp log-neighbor-changes
no auto-summary
!
address-family ipv4 vrf MGMT
no synchronization
redistribute ospf 1 vrf MGMT
exit-address-family
!
address-family ipv4 vrf PROD
no synchronization
redistribute ospf 2 vrf PROD
exit-address-family
```

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.

Internal Only - General

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.
R1(config)#ip vrf MGMT
R1(config-vrf)#no route-target import 8:8

Additional Information:

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

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

Internal Only - General

1.46 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-216592
RULE ID: SV-216592r531085
```

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.

Internal Only - General

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 GigabitEthernet0/2
description OOB link to NOC
ip address 10.11.1.8 255.255.255.0
ip access-group TRAFFIC_FROM_NOC in
!
interface GigabitEthernet0/3
description link to OOBM LAN access switch
ip address 10.13.1.1 255.255.255.0
ip access-group TRAFFIC_TO_NOC in
```

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.

```
ip access-list extended TRAFFIC_FROM_NOC
permit ip 10.11.1.0 0.255.255.255 host 10.11.1.8
permit ip 10.12.1.0 0.255.255.255 host 10.11.1.8
permit ip 10.11.1.0 0.255.255.255 host 10.13.1.1
permit ip 10.12.1.0 0.255.255.255 host 10.13.1.1
deny ip any host 10.11.1.8 log-input
deny ip any host 10.13.1.1 log-input
permit ip 10.11.1.0 0.0.0.255 10.13.1.0 0.0.0.255
permit ip 10.12.1.0 0.0.0.255 10.13.1.0 0.0.0.255
deny ip 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.

```
ip access-list extended TRAFFIC_TO_NOC
permit ip 10.13.1.0 0.255.255.255 host 10.13.1.1
permit ip 10.13.1.0 0.255.255.255 host 10.11.1.8
deny ip any host 10.13.1.1 log-input
deny ip any host 10.11.1.8 log-input
permit ip 10.13.1.0 0.255.255.255 10.11.1.0 0.0.0.255
permit ip 10.13.1.0 0.255.255.255 10.12.1.0 0.0.0.255
deny ip 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.

Internal Only - General

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.

```
R4(config)#ip access-list extended TRAFFIC_FROM_NOC
R4(config-ext-nacl)#permit ip 10.11.1.0 0.255.255.255 host 10.11.1.8
R4(config-ext-nacl)#permit ip 10.12.1.0 0.255.255.255 host 10.11.1.8
R4(config-ext-nacl)#permit ip 10.11.1.0 0.255.255.255 host 10.13.1.1
R4(config-ext-nacl)#permit ip 10.12.1.0 0.255.255.255 host 10.13.1.1
R4(config-ext-nacl)#deny ip any host 10.11.1.8 log-input
R4(config-ext-nacl)#deny ip any host 10.13.1.1 log-input
R4(config-ext-nacl)#permit ip 10.11.1.0 0.0.0.255 10.13.1.0 0.0.0.255
R4(config-ext-nacl)#permit ip 10.12.1.0 0.0.0.255 10.13.1.0 0.0.0.255
R4(config-ext-nacl)#deny ip any any log-input
```

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

```
R4(config)#ip access-list extended TRAFFIC_TO_NOC
R4(config-ext-nacl)#permit ip 10.13.1.0 0.255.255.255 host 10.13.1.1
R4(config-ext-nacl)#permit ip 10.13.1.0 0.255.255.255 host 10.11.1.8
R4(config-ext-nacl)#deny ip any host 10.13.1.1 log-input
R4(config-ext-nacl)#deny ip any host 10.11.1.8 log-input
R4(config-ext-nacl)#permit ip 10.13.1.0 0.255.255.255 10.11.1.0 0.0.0.255
R4(config-ext-nacl)#permit ip 10.13.1.0 0.255.255.255 10.12.1.0 0.0.0.255
R4(config-ext-nacl)#deny ip any any log-input
R4(config-ext-nacl)#exit
```

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

```
R4(config)#int g0/2
R4(config-if)#ip access-group TRAFFIC_FROM_NOC in
R4(config)#int g0/3
R4(config-if)#ip access-group TRAFFIC_TO_NOC in
R4(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

Internal Only - General

1.47 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-216593
RULE ID: SV-216593r531085
```

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.

Internal Only - General

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 GigabitEthernet0/7
description link to OOBM access switch
ip address 10.11.1.22 255.255.255.0
ip access-group INGRESS_MANAGEMENT_ACL in
ip access-group EGRESS_MANAGEMENT_ACL out
Step 2: Verify that the ingress ACL only allows management and ICMP traffic.
ip access-list extended INGRESS_MANAGEMENT_ACL
permit tcp any host 10.11.1.22 eq tacacs
permit tcp any host 10.11.1.22 eq 22
permit udp any host 10.11.1.22 eq snmp
permit udp any host 10.11.1.22 eq snmptrap
permit udp any host 10.11.1.22 eq ntp
permit icmp any host 10.11.1.22
deny ip any any log-input
Step 3: Verify that the egress ACL blocks any transit traffic.
ip access-list extended EGRESS_MANAGEMENT_ACL
deny ip 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.

Internal Only - General

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.

```
R5(config)#ip access-list extended INGRESS_MANAGEMENT_ACL
R5(config-ext-nacl)#permit tcp any host 10.11.1.22 eq tacacs
R5(config-ext-nacl)#permit tcp any host 10.11.1.22 eq 22
R5(config-ext-nacl)#permit udp any host 10.11.1.22 eq snmp
R5(config-ext-nacl)#permit udp any host 10.11.1.22 eq snmptrap
R5(config-ext-nacl)#permit udp any host 10.11.1.22 eq ntp
R5(config-ext-nacl)#permit icmp any host 10.11.1.22
R5(config-ext-nacl)#deny ip any any log-input
R5(config-ext-nacl)#exit
```

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

```
R5(config)#ip access-list extended EGRESS_MANAGEMENT_ACL
R5(config-ext-nacl)#deny ip any any log-input
R5(config-ext-nacl)#exit
```

Step 3: Apply the ACLs to the OOBM interfaces.

```
R4(config)#int g0/7
R4(config-if)#ip access-group INGRESS_MANAGEMENT_ACL in
R4(config-if)#ip access-group EGRESS_MANAGEMENT_ACL 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

Internal Only - General

1.48 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-216594
RULE ID: SV-216594r531085
```

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.

Internal Only - General

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 crypto map applied to the external interface.
interface GigabitEthernet0/2
description link to DISN
ip address x.1.24.4 255.255.255.0
crypto map IPSEC_MGMT_MAP
Step 2: Review the ISAKMP policy for Phase 1 negotiations and Phase 2 policy for data encryption.
crypto isakmp policy 10
authentication pre-share
hash sha256
crypto isakmp key xxxxxx address x.1.12.1
!
!
crypto ipsec transform-set TRANS_SET ah-sha256-hmac esp-aes
Step 3: Review the crypto map that was bound to the external interface and note the ACL defined that identifies the interesting traffic for the IPsec tunnel.
crypto map IPSEC_MGMT_MAP 10 ipsec-isakmp
set peer x.1.12.1
set transform-set TRANS_SET
match address MGMT_TRAFFIC_ACL
Step 4: Review the ACL defined in the crypto map and verify that the destination is the management network.
ip access-list extended MGMT_TRAFFIC_ACL
permit ip 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.

Internal Only - General

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.
R4(config)#ip access-list extended MGMT_TRAFFIC_ACL
R4(config-ext-nacl)#permit ip 10.1.34.0 0.0.0.255 10.22.2.0 0.0.0.255
R4(config-ext-nacl)#exit
Step 2: Create an ISAKMP policy for Phase 1 negotiations
R4(config)#crypto isakmp policy 10
R4(config-isakmp)#hash sha256
R4(config-isakmp)#authentication pre-share
R4(config-isakmp)#exit
Step 3: Specify the pre-shared key and the remote peer address
R4(config)#crypto isakmp key 0 xxxxxx address 10.1.12.1
Note: Digital certificates can be utilized as an alternative.
Step 4: Create the Phase 2 policy for the data encryption
R4(config)#crypto ipsec transform-set TRANS_SET ah-sha256-hmac esp-aes
R4(cfg-crypto-trans)#mode tunnel
R4(cfg-crypto-trans)#exit
Step 5: Create the crypto map
R4(config)#crypto map IPSEC_MGMT_MAP 10 ipsec-isakmp
R4(config-crypto-map)#set peer 10.1.12.1
R4(config-crypto-map)#match address MGMT_TRAFFIC_ACL
R4(config-crypto-map)#set transform-set TRANS_SET
R4(config-crypto-map)#end
Step 6: Apply the crypto map to the external interface
R4(config)#int g0/2
R4(config-if)#crypto map IPSEC_MGMT_MAP

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

Internal Only - General

1.49 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-216991
RULE ID: SV-216991r856208
```

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 xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 password xxxxxxxx
neighbor x.1.1.9 ttl-security hops 1
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 password xxxxxxxx
neighbor x.2.1.7 ttl-security hops 1
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.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 ttl-security hops 1
R1(config-router)#neighbor x.2.1.7 ttl-security hops 1

Internal Only - General

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

Internal Only - General

1.50 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-216992
RULE ID: SV-216992r877980
```

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:

This check is Not Applicable for JRSS internal EBGP use.
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 xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 password yyyyyyyy
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 password zzzzzzzzz
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.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 password yyyyyyyy
R1(config-router)#neighbor x.2.1.7 password zzzzzzzzz

Internal Only - General

Additional Information:

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

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

Internal Only - General

1.51 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-216597
RULE ID: SV-216597r877976
```

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.

Internal Only - General

Audit:

This check is Not Applicable for JRSS internal EBGP use.
Review the router configuration to verify that it will reject BGP routes for any Bogon prefixes.
Step 1: Verify a prefix list has been configured containing the current Bogon prefixes as shown in the example below.
ip prefix-list PREFIX_FILTER seq 5 deny 0.0.0.0/8 le 32
ip prefix-list PREFIX_FILTER seq 10 deny 10.0.0.0/8 le 32
ip prefix-list PREFIX_FILTER seq 15 deny 100.64.0.0/10 le 32
ip prefix-list PREFIX_FILTER seq 20 deny 127.0.0.0/8 le 32
ip prefix-list PREFIX_FILTER seq 25 deny 169.254.0.0/16 le 32
ip prefix-list PREFIX_FILTER seq 30 deny 172.16.0.0/12 le 32
ip prefix-list PREFIX_FILTER seq 35 deny 192.0.2.0/24 le 32
ip prefix-list PREFIX_FILTER seq 40 deny 192.88.99.0/24 le 32
ip prefix-list PREFIX_FILTER seq 45 deny 192.168.0.0/16 le 32
ip prefix-list PREFIX_FILTER seq 50 deny 198.18.0.0/15 le 32
ip prefix-list PREFIX_FILTER seq 55 deny 198.51.100.0/24 le 32
ip prefix-list PREFIX_FILTER seq 60 deny 203.0.113.0/24 le 32
ip prefix-list PREFIX_FILTER seq 65 deny 224.0.0.0/4 le 32
ip prefix-list PREFIX_FILTER seq 70 deny 240.0.0.0/4 le 32
ip prefix-list PREFIX_FILTER seq 75 permit 0.0.0.0/0 ge 8
Step 2: Verify that the prefix list has been applied to all external BGP peers as shown in the example below.
router bgp xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 prefix-list PREFIX_FILTER in
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 prefix-list PREFIX_FILTER in
Route Map Alternative
Verify that the route map applied to the external neighbors references the configured Bogon prefix list shown above.
router bgp xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 route-map FILTER_PREFIX_MAP
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 route-map FILTER_PREFIX_MAP
…
route-map FILTER_PREFIX_MAP permit 10
match ip address prefix-list PREFIX_FILTER
If the router is not configured to reject inbound route advertisements for any Bogon prefixes, this is a finding.

Internal Only - General

Remediation:

Configure the router to reject inbound route advertisements for any Bogon prefixes.
Step 1: Configure a prefix list containing the current Bogon prefixes as shown below.
R1(config)#ip prefix-list PREFIX_FILTER deny 0.0.0.0/8 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 10.0.0.0/8 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 100.64.0.0/10 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 127.0.0.0/8 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 169.254.0.0/16 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 172.16.0.0/12 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 192.0.2.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 192.88.99.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 192.168.0.0/16 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 198.18.0.0/15 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 198.51.100.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 203.0.113.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 224.0.0.0/4 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 240.0.0.0/4 le 32
R1(config)#ip prefix-list PREFIX_FILTER deny 240.0.0.0/4 le 32
R1(config)#ip prefix-list PREFIX_FILTER permit 0.0.0.0/0 ge 8
Step 2: Apply the prefix list filter inbound to each external BGP neighbor as shown in the example.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 prefix-list PREFIX_FILTER in
R1(config-router)#neighbor x.2.1.7 prefix-list PREFIX_FILTER in
Route Map Alternative
Step 1: Configure the route map referencing the configured prefix list above.
R1(config)#route-map FILTER_PREFIX_MAP 10
R1(config-route-map)#match ip address prefix-list PREFIX_FILTER
R1(config-route-map)#exit
Step 2: Apply the route-map inbound to each external BGP neighbor as shown in the example.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 route-map FILTER_PREFIX_MAP in
R1(config-router)#neighbor x.2.1.7 route-map FILTER_PREFIX_MAP in
R1(config-router)#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

Internal Only - General

1.52 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-216598
RULE ID: SV-216598r531085
```

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 a prefix list has been configured containing prefixes belonging to the local AS. In the example below x.13.1.0/24 is the global address space allocated to the local AS.
ip prefix-list PREFIX_FILTER seq 5 deny 0.0.0.0/8 le 32
…
…
…
ip prefix-list PREFIX_FILTER seq 74 deny x.13.1.0/24 le 32
ip prefix-list PREFIX_FILTER seq 75 permit 0.0.0.0/0 ge 8
Step 2: Verify that the prefix list has been applied to all external BGP peers as shown in the example below.
router bgp xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 prefix-list PREFIX_FILTER in
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 prefix-list PREFIX_FILTER in
If the router is not configured to reject inbound route advertisements belonging to the local AS, this is a finding.

Internal Only - General

Remediation:

Configure the router to reject inbound route advertisements for any prefixes belonging to the local AS.
Step 1: Add to the prefix filter list those prefixes belonging to the local autonomous system.
R1(config)#ip prefix-list PREFIX_FILTER seq 74 deny x.13.1.0/24 le 32
Step 2: If not already completed to be compliant with previous requirement, apply the prefix list filter inbound to each external BGP neighbor as shown in the example.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 prefix-list PREFIX_FILTER in
R1(config-router)#neighbor x.2.1.7 prefix-list 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

Internal Only - General

1.53 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-216599
RULE ID: SV-216599r917412
```

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 prefix list has been configured for each customer containing prefixes belonging to each customer as shown in the example below.
ip prefix-list PREFIX_FILTER_CUST1 seq 5 permit x.13.1.0/24 le 32
ip prefix-list PREFIX_FILTER_CUST1 seq 10 deny 0.0.0.0/0 ge 8
ip prefix-list PREFIX_FILTER_CUST2 seq 5 permit x.13.2.0/24 le 32
ip prefix-list PREFIX_FILTER_CUST2 seq 10 deny 0.0.0.0/0 ge 8
Step 2: Verify that the prefix lists has been applied to all to the applicable CE peers as shown in the example below.
router bgp xx
no synchronization
bgp log-neighbor-changes
neighbor x.12.4.14 remote-as 64514
neighbor x.12.4.14 prefix-list FILTER_PREFIXES_CUST1 in
neighbor x.12.4.16 remote-as 64516
neighbor x.12.4.16 prefix-list FILTER_PREFIXES_CUST2 in
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.
NOTE: This check is NA for JRSS systems.
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.

Internal Only - General

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 list for each customer containing prefixes belonging to each.
R1(config)#ip prefix-list PREFIX_FILTER_CUST1 permit x.13.1.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER_CUST1 deny 0.0.0.0/0 ge 8
R1(config)#ip prefix-list PREFIX_FILTER_CUST2 permit x.13.2.0/24 le 32
R1(config)#ip prefix-list PREFIX_FILTER_CUST2 deny 0.0.0.0/0 ge 8
Step 2: Apply the prefix list filter inbound to each CE neighbor as shown in the example.
R1(config)#router bgp xx
R1(config-router)#neighbor x.12.4.14 prefix-list FILTER_PREFIXES_CUST1 in
R1(config-router)#neighbor x.12.4.16 prefix-list FILTER_PREFIXES_CUST2 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

Internal Only - General

1.54 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-216600
RULE ID: SV-216600r917414
```

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 a prefix list has been configured containing prefixes belonging to customers as well as the local AS as shown in the example below.
ip prefix-list CE_PREFIX_ADVERTISEMENTS seq 5 permit x.13.1.0/24 le 32
ip prefix-list CE_PREFIX_ADVERTISEMENTS seq 10 permit x.13.2.0/24 le 32
ip prefix-list CE_PREFIX_ADVERTISEMENTS seq 15 permit x.13.3.0/24 le 32
ip prefix-list CE_PREFIX_ADVERTISEMENTS seq 20 permit x.13.4.0/24 le 32
…
…
…
ip prefix-list CE_PREFIX_ADVERTISEMENTS seq 80 deny 0.0.0.0/0 ge 8
Step 2: Verify that the prefix lists has been applied to all CE peers as shown in the example below.
router bgp 64512
no synchronization
bgp log-neighbor-changes
neighbor x.12.4.14 remote-as 64514
neighbor x.12.4.14 prefix-list CE_PREFIX_ADVERTISEMENTS out
neighbor x.12.4.16 remote-as 64516
neighbor x.12.4.16 prefix-list CE_PREFIX_ADVERTISEMENTS out
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.
Note: This check is NA for JRSS systems.

Internal Only - General

Remediation:

This requirement is not applicable for the DODIN Backbone.
Step 1: Configure a prefix list for containing all customer and local AS prefixes as shown in the example below.
R1(config)#ip prefix-list CE_PREFIX_ADVERTISEMENTS permit x.13.1.0/24 le 32
R1(config)#ip prefix-list CE_PREFIX_ADVERTISEMENTS permit x.13.2.0/24 le 32
R1(config)#ip prefix-list CE_PREFIX_ADVERTISEMENTS permit x.13.3.0/24 le 32
R1(config)#ip prefix-list CE_PREFIX_ADVERTISEMENTS permit x.13.4.0/24 le 32
…
…
…
R1(config)#ip prefix-list CE_PREFIX_ADVERTISEMENTS deny 0.0.0.0/0 ge 8

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

Internal Only - General

1.55 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-216601
RULE ID: SV-216601r531085
```

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 a prefix list has been configured containing prefixes belonging to the IP core.
ip prefix-list FILTER_CORE_PREFIXES seq 5 deny x.1.1.0/24 le 32
ip prefix-list FILTER _CORE_PREFIXES seq 10 deny x.1.2.0/24 le 32
ip prefix-list FILTER _CORE_PREFIXES seq 15 permit 0.0.0.0/0 ge 8
Step 2: Verify that the prefix lists has been applied to all external BGP peers as shown in the example below.
router bgp xx
no synchronization
bgp log-neighbor-changes
neighbor x.1.4.12 remote-as yy
neighbor x.1.4.12 prefix-list FILTER _CORE_PREFIXES out
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 list for containing all customer and local AS prefixes as shown in the example below.
R1(config)#ip prefix-list FILTER_CORE_PREFIXES deny x.1.1.0/24 le 32
R1(config)#ip prefix-list FILTER _CORE_PREFIXES deny x.1.2.0/24 le 32
R1(config)#ip prefix-list FILTER _CORE_PREFIXES permit 0.0.0.0/0 ge 8
Step 2: Apply the prefix list filter outbound to each CE neighbor as shown in the example.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.4.12 prefix-list FILTER _CORE_PREFIXES out

Internal Only - General

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

Internal Only - General

1.56 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-216602
RULE ID: SV-216602r531085
```

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 no bgp enforce-first-as is not configured.
router bgp xx
no synchronization
no bgp enforce-first-as
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.
R1(config)#router bgp xx
R1(config-router)#bgp enforce-first-as

Internal Only - General

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)

Internal Only - General

1.57 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-216603
RULE ID: SV-216603r531085
```

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: Review router configuration and verify that there is an as-path access-list statement defined to only accept routes from a CE router whose AS did not originate the route. The configuration should look similar to the following:
ip as-path access-list 10 permit ^yy$
ip as-path access-list 10 deny .*
Note: the characters “^” and “$” representing the beginning and the end of the expression respectively are optional and are implicitly defined if omitted.
Step 2: Verify that the as-path access-list is referenced by the filter-list inbound for the appropriate BGP neighbors as shown in the example below:
router bgp xx
neighbor x.1.4.12 remote-as yy
neighbor x.1.4.12 filter-list 10 in
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.

Internal Only - General

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 the as-path ACL as shown in the example below:
R1(config)#ip as-path access-list 10 permit ^yy$
R1(config)#ip as-path access-list 10 deny .*
Step 2: Apply the as-path filter inbound as shown in the example below:
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.4.12 filter-list 10 in

Additional Information:

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

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

Internal Only - General

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-216605
RULE ID: SV-216605r856193
```

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 a route filter has been configured to reject prefixes longer than /24, or the least significant prefixes issued to the customers as shown in the example below:
ip prefix-list FILTER_PREFIX_LENGTH seq 5 permit 0.0.0.0/0 ge 8 le 24
ip prefix-list FILTER_PREFIX_LENGTH seq 10 deny 0.0.0.0/0 le 32
Step 2: Verify that prefix filtering has been applied to each eBGP peer as shown in the example:
router bgp xx
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 prefix-list FILTER_PREFIX_LENGTH in
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 prefix-list FILTER_PREFIX_LENGTH in
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.

Internal Only - General

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 list to reject any prefix that is longer than /24.
R1(config)#ip prefix-list FILTER_PREFIX_LENGTH permit 0.0.0.0/0 ge 8 le 24
R1(config)#ip prefix-list FILTER_PREFIX_LENGTH deny 0.0.0.0/0 le 32
Step 2: Apply the prefix list to all eBGP peers as shown in the example below.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 prefix-list FILTER_PREFIX_LENGTH in
R1(config-router)#neighbor x.2.1.7 prefix-list FILTER_PREFIX_LENGTH in

Additional Information:

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

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

Internal Only - General

1.59 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-216604
RULE ID: SV-216604r856192
```

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.

Internal Only - General

Audit:

Review the router configuration to verify that the number of received prefixes from each eBGP neighbor is controlled.
router bgp xx
neighbor x.1.1.9 remote-as yy
neighbor x.1.1.9 maximum-prefix nnnnnnn
neighbor x.2.1.7 remote-as zz
neighbor x.2.1.7 maximum-prefix nnnnnnn
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.
R1(config)#router bgp xx
R1(config-router)#neighbor x.1.1.9 maximum-prefix nnnnnnn
R1(config-router)#neighbor x.2.1.7 maximum-prefix nnnnnnn

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

Internal Only - General

1.60 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-216606
RULE ID: SV-216606r531085
```

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.

Internal Only - General

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
no synchronization
no bgp enforce-first-as
bgp log-neighbor-changes
redistribute static
neighbor 10.1.1.1 remote-as xx
neighbor 10.1.1.1 password xxxxxxxx
neighbor 10.1.1.1 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.
R1(config)#router bgp xx
R1(config-router)#neighbor 10.1.1.1 update-source Loopback0

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

Internal Only - General

1.61 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-216607
RULE ID: SV-216607r531085
```

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 a loopback interface as shown in the example below.
mpls ldp router-id Loopback0
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.
R4(config)#mpls ldp router-id lo0

Internal Only - General

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

Internal Only - General

1.62 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-216608
RULE ID: SV-216608r531085
```

Rationale:

Packet loss can occur when an IGP adjacency is established and the router begins forwarding packets using the new adjacency before the LDP label exchange completes between the peers on that link. Packet loss can also occur if an LDP session closes and the router continues to forward traffic using the link associated with the LDP peer rather than an alternate pathway with a fully synchronized LDP session. The MPLS LDP-IGP Synchronization feature provides a means to synchronize LDP with OSPF or IS-IS to minimize MPLS packet loss. When an IGP adjacency is established on a link but LDP-IGP synchronization is not yet achieved or is lost, the IGP will advertise the max-metric on that link.

Audit:

Review the router OSPF or IS-IS configuration and verify that LDP will synchronize with the link-state routing protocol as shown in the example below.
OSPF Example
router ospf 1
mpls ldp sync
IS-IS Example
router isis
mpls ldp sync
net 49.0001.1234.1600.5531.00
If the router is not configured to synchronize IGP and LDP, this is a finding.

Internal Only - General

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
R2(config)#router ospf 1
R2(config-router)#mpls ldp sync
IS-IS Example
R5(config)#router isis
R5(config-router)#mpls ldp sync

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

Internal Only - General

1.63 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-216609
RULE ID: SV-216609r531085
```

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.

Internal Only - General

Audit:

Review the router configuration to determine RSVP messages are rate limited.
Step 1: Determine if MPLS TE is enabled globally and at least one interface as shown in the example below.
mpls traffic-eng tunnels
…
…
…
interface GigabitEthernet0/2
ip address x.x.x.x 255.255.255.0
mpls traffic-eng tunnels
mpls ip
Step 2: If MPLS TE is enabled, verify that message pacing is enabled.
ip rsvp signaling rate-limit period 30 burst 9 maxsize 2100 limit 50
Note: The command ip rsvp msg-pacing has been deprecated by the command ip rsvp signaling rate-limit.
If the router with RSVP-TE enabled does not rate limit RSVP messages 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 per interface as shown in the example.
R2(config)#ip rsvp signaling rate-limit burst 9 maxsize 2100 period 30 limit 50

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)

Internal Only - General

1.64 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-216610
RULE ID: SV-216610r531085
```

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.
no mpls ip propagate-ttl
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.
R5(config)#no mpls ip propagate-ttl

Internal Only - General

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

Internal Only - General

1.65 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-216611
RULE ID: SV-216611r531085
```

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 MPLS/L3VPN.
Step 2: 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 GigabitEthernet0/1, while COI2 is bound to GigabitEthernet0/2.
interface GigabitEthernet0/1
description link to COI1
ip vrf forwarding COI1
ip address x.1.0.1 255.255.255.0
!
interface GigabitEthernet0/2
description link to COI2
ip vrf forwarding COI2
ip address x.2.0.2 255.255.255.0
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.

Internal Only - General

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

Internal Only - General

1.66 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-216612
RULE ID: SV-216612r531085
```

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 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 customer 1.
ip vrf CUST1
rd 13:13
route-target export 13:13
route-target import 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.
R5(config)#ip vrf CUST1
R5(config-vrf)#route-target import 13:13
R5(config-vrf)#route-target export 13:13
R5(config-vrf)#end

Internal Only - General

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

Internal Only - General

1.67 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-216613
RULE ID: SV-216613r531085
```

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 customer 1.
ip vrf CUST1
rd 13:13
Note: This requirement is only applicable for MPLS L3VPN implementations.
If the wrong RD has been configured for any VRF, this is a finding.

Remediation:

Configure the correct RD for each VRF.
R5(config)#ip vrf CUST1
R5(config-vrf)#rd 13:13
R5(config-vrf)#end

Internal Only - General

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

Internal Only - General

1.68 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-216614
RULE ID: SV-216614r864155
```

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 mitigated 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 neighbor 10.1.1.2 password xxxxxxx
mpls label protocol ldp
If the router is not configured to authenticate targeted LDP sessions using MD5, the finding will remain as a category 2.

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.
R5(config)#mpls ldp neighbor 10.1.1.2 password xxxxxxxx

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

Internal Only - General

1.69 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 virtual circuit identification (VC ID) for each attachment circuit.

```
GROUP ID: V-216615
RULE ID: SV-216615r531085
```

Rationale:

VPWS is an L2VPN technology that provides a virtual circuit between two PE routers to forward Layer 2 frames between two customer-edge routers or 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 VC label bindings for the specified VC ID. The VC ID specifies a pseudowire associated with an ingress and egress PE router and the customer-facing attachment circuits.

To guarantee that all frames are forwarded onto the correct pseudowire and to the correct customer and attachment circuits, it is imperative that the correct VC ID is configured for each attachment circuit.

Audit:

Verify that the correct and unique VCID has been configured for the appropriate attachment circuit. In the example below GigabitEthernet0/1 is the CE-facing interface that is configured for VPWS with the VCID of 55.
interface GigabitEthernet0/1
xconnect x.2.2.12 55 encapsulation mpls
If the correct VC ID has not been configured on both routers, this is a finding.

Remediation:

Assign globally unique VC IDs for each virtual circuit and configure the attachment circuits with the appropriate VC ID.
R5(config)#int g0/1
R5(config-if)#xconnect x.2.2.12 55 encapsulation mpls

Internal Only - General

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

Internal Only - General

1.70 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-216616
RULE ID: SV-216616r531085
```

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 GigabitEthernet0/2
ip address x.1.12.2 255.255.255.252
ip access-group BLOCK_TO_CORE in
Step 2: Verify that the ingress ACL discards and logs packets destined to the IP core address space.
ip access-list extended BLOCK_TO_CORE
deny ip any 10.1.x.0 0.0.255.255 log-input
permit ip 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.

Internal Only - General

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.
R2(config)#ip access-list extended BLOCK_TO_CORE
R2(config-ext-nacl)#deny ip any 10.1.x.0 0.0.255.255 log-input
R2(config-ext-nacl)#exit
Step 2: Apply the ACL inbound to all external or CE-facing interfaces.
R2(config)#int R4(config)#int g0/2
R2(config-if)#ip access-group BLOCK_TO_CORE in
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

Internal Only - General

1.71 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-216617
RULE ID: SV-216617r531085
```

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 GigabitEthernet0/2
ip address x.1.12.2 255.255.255.252
ip access-group BLOCK_TO_CORE in
ip 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 below.
R2(config)#int R4(config)#int g0/2
R2(config-if)#ip verify unicast source reachable-via any
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

Internal Only - General

1.72 CISC-RT-000750 (Manual)

Profile Applicability:

- SEVERITY: CAT II

Description:

The Cisco PE router must be configured to drop all packets with any IP options.

```
GROUP ID: V-216993
RULE ID: SV-216993r856210
```

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:

Review the router configuration to determine if it will drop all packets with IP options as shown below.
ip options drop
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 IP options as shown below.
R4(config)#ip options drop

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)

Internal Only - General

1.73 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-216619
RULE ID: SV-216619r917417
```

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.

Internal Only - General

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 C2_VOICE
match ip dscp af47
class-map match-all VOICE
match ip dscp ef
class-map match-all VIDEO
match ip dscp af41
class-map match-all CONTROL_PLANE
match ip dscp cs6
class-map match-all PREFERRED_DATA
match ip dscp af33
```

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
priority percent 10
class VOICE
priority percent 15
class VIDEO
bandwidth percent 25
class CONTROL_PLANE
priority percent 10
class PREFERRED_DATA
bandwidth percent 25
class class-default
bandwidth percent 15
```

Step 3: Verify that an output service policy is bound to all interfaces as shown in the configuration example below.

```
interface GigabitEthernet1/1
ip address 10.1.15.1 255.255.255.252
service-policy output QOS_POLICY
!
```


Internal Only - General

interface GigabitEthernet1/2
ip address 10.1.15.4 255.255.255.252
service-policy output QOS_POLICY
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.

Internal Only - General

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.

```
R5(config-cmap)#class-map match-all C2_VOICE
R5(config-cmap)#match ip dscp 47
R5(config-cmap)#class-map match-all VOICE
R5(config-cmap)#match ip dscp ef
R5(config-cmap)#class-map match-all VIDEO
R5(config-cmap)#match ip dscp af41
R5(config-cmap)#class-map match-all CONTROL_PLANE
R5(config-cmap)#match ip dscp cs6
R5(config)#class-map match-all PREFERRED_DATA
R5(config-cmap)#match ip dscp af33
R5(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.

```
R5(config)#policy-map QOS_POLICY
R5(config-pmap-c)#class C2_VOICE
R5(config-pmap-c)#priority percent 10
R5(config-pmap-c)#class VOICE
R5(config-pmap-c)#priority percent 15
R5(config-pmap-c)#class VIDEO
R5(config-pmap-c)#bandwidth percent 25
R5(config-pmap)#class CONTROL_PLANE
R5(config-pmap-c)#priority percent 10
R5(config-pmap-c)#class PREFERRED_DATA
R5(config-pmap-c)#bandwidth percent 25
R5(config-pmap-c)#class class-default
R5(config-pmap-c)#bandwidth percent 15
R5(config-pmap-c)#exit
R5(config-pmap)#exit
```

Step 3: Apply the output service policy to the core-layer-facing interface as shown in the configuration example below.

```
R5(config)#int g1/1
R5(config-if)#service-policy output QOS_POLICY
R5(config-if)#exit
R5(config)#int g1/2
R5(config-if)#service-policy output QOS_POLICY
R5(config-if)#end
```


Internal Only - General

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)

Internal Only - General

1.74 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-216620
RULE ID: SV-216620r917420
```

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.

Internal Only - General

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 PREFERRED_DATA
match ip dscp af33
class-map match-all CONTROL_PLANE
match ip dscp cs6
class-map match-all VIDEO
match ip dscp af41
class-map match-all VOICE
match ip dscp ef
class-map match-all C2_VOICE
match ip dscp 47
```

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 CONTROL_PLANE
priority percent 10
class C2_VOICE
priority percent 10
class VOICE
priority percent 15
class VIDEO
bandwidth percent 25
class PREFERRED_DATA
bandwidth percent 25
class class-default
bandwidth percent 15
```

Step 3: Verify that an output service policy is bound to all interfaces as shown in the configuration example below.

```
interface GigabitEthernet1/1
ip address 10.1.15.5 255.255.255.252
service-policy output QOS_POLICY
!
interface GigabitEthernet1/2
ip address 10.1.15.8 255.255.255.252
service-policy output QOS_POLICY
```

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

Internal Only - General

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.

```
R5(config)#class-map match-all PREFERRED_DATA
R5(config-cmap)#match ip dscp af33
R5(config-cmap)#class-map match-all CONTROL_PLANE
R5(config-cmap)#match ip dscp cs6
R5(config-cmap)#class-map match-all VIDEO
R5(config-cmap)#match ip dscp af41
R5(config-cmap)#class-map match-all VOICE
R5(config-cmap)#match ip dscp ef
R5(config-cmap)#class-map match-all C2_VOICE
R5(config-cmap)#match ip dscp 47
R5(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.

```
R5(config)#policy-map QOS_POLICY
R5(config-pmap)#class CONTROL_PLANE
R5(config-pmap-c)#priority percent 10
R5(config-pmap-c)#class C2_VOICE
R5(config-pmap-c)#priority percent 10
R5(config-pmap-c)#class VOICE
R5(config-pmap-c)#priority percent 15
R5(config-pmap-c)#class VIDEO
R5(config-pmap-c)#bandwidth percent 25
R5(config-pmap-c)#class PREFERRED_DATA
R5(config-pmap-c)#bandwidth percent 25
R5(config-pmap-c)#class class-default
R5(config-pmap-c)#bandwidth percent 15
R5(config-pmap-c)#exit
R5(config-pmap)#exit
```

Step 3: Apply the output service policy to all interfaces as shown in the configuration example below.

```
R5(config)#int g1/1
R5(config-if)#service-policy output QOS_POLICY
R5(config-if)#exit
R5(config)#int g1/2
R5(config-if)#service-policy output QOS_POLICY
R5(config-if)#end
```


Internal Only - General

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)

Internal Only - General

1.75 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-216621
RULE ID: SV-216621r531085
```

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).

Internal Only - General

Audit:

Review the router configuration to determine if it is configured to enforce a QoS policy to limit the effects of packet flooding DoS attacks.
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 ip dscp cs1
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 CONTROL_PLANE
priority percent 10
class C2_VOICE
priority percent 10
class VOICE
priority percent 15
class VIDEO
bandwidth percent 25
class PREFERRED_DATA
bandwidth percent 25
class SCAVENGER
bandwidth percent 5
class class-default
bandwidth percent 10
Note: Traffic out of profile must be marked at the customer access layer or CE egress edge.
If the router is not configured to enforce a QoS policy to limit the effects of packet flooding DoS attacks, this is a finding.

Remediation:

Step 1: Configure a class map for the SCAVENGER class.
R5(config)#class-map match-all SCAVENGER
R5(config-cmap)#match ip dscp cs1
Step 2: Add the SCAVENGER class to the policy map as shown in the example below.
R5(config)#policy-map QOS_POLICY
R5(config-pmap-c)#no class class-default
R5(config-pmap)#class SCAVENGER
R5(config-pmap-c)#bandwidth percent 5
R5(config-pmap-c)#class class-default
R5(config-pmap-c)#bandwidth percent 10
R5(config-pmap-c)#end

Internal Only - General

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)

Internal Only - General

1.76 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-216622
RULE ID: SV-216622r531085
```

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 PIM interfaces as shown in the multicast topology diagram are enabled for PIM as shown in the example below.
interface GigabitEthernet1/1
ip address 10.1.3.3 255.255.255.0
ip pim sparse-mode
If an interface is not required to support multicast routing and it is enabled, this is a finding.

Internal Only - General

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.
R5(config)#int g1/1
R5(config-if)#no ip pim sparse-mode

Additional Information:

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

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

Internal Only - General

1.77 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-216623
RULE ID: SV-216623r531085
```

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.
interface GigabitEthernet1/1
ip address 10.1.2.2 255.255.255.0
ip pim neighbor-filter PIM_NEIGHBORS
ip pim sparse-mode
Step 2: Review the configured ACL for filtering PIM neighbors as shown in the example below.
ip access-list standard PIM_NEIGHBORS
permit 10.1.2.6
If PIM neighbor ACLs are not bound to all interfaces that have PIM enabled, this is a finding.

Internal Only - General

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
R2(config)#ip access-list standard PIM_NEIGHBORS
R2(config-std-nacl)#permit 10.1.2.6
R2(config-std-nacl)#exit
Step 2: Apply the ACL to all interfaces enabled for PIM
R2(config)#int g1/1
R2(config-if)#ip pim neighbor-filter PIM_NEIGHBORS

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

Internal Only - General

1.78 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-216624
RULE ID: SV-216624r531085
```

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.
interface GigabitEthernet1/2
ip address x.1.12.2 255.255.255.252
ip pim sparse-mode
ip multicast boundary MULTICAST_SCOPE
…
…
…
ip access-list standard MULTICAST_SCOPE
deny 239.0.0.0 0.255.255.255
permit any
If the router is not configured to establish boundaries for administratively scoped multicast traffic, this is a finding.

Internal Only - General

Remediation:

Step 1: Configure the ACL to deny packets with multicast administratively scoped destination addresses as shown in the example below.
R2(config)#ip access-list standard MULTICAST_SCOPE
R2(config-std-nacl)#deny 239.0.0.0 0.255.255.255
R2(config-std-nacl)#permit any
R2(config-std-nacl)#exit
Step 2: Apply the multicast boundary at the appropriate interfaces as shown in the example below.
R2(config)#int g1/2
R2(config-if)#ip multicast boundary MULTICAST_SCOPE
R2(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

Internal Only - General

1.79 CISC-RT-000820 (Manual)

Profile Applicability:

- SEVERITY: CAT III

Description:

The Cisco multicast Rendezvous Point (RP) router must be configured to limit the multicast forwarding cache so that its resources are not saturated by managing an overwhelming number of Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) source-active entries.

```
GROUP ID: V-216625
RULE ID: SV-216625r864156
```

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.

Internal Only - General

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.
ip pim rp-address 10.1.12.3
ip pim accept-register list PIM_REGISTER_FILTER
…
…
…
ip access-list extended PIM_REGISTER_FILTER
deny ip any 239.5.0.0 0.0.255.255
permit ip host 10.1.2.6 any
permit ip host 10.1.2.7 any
deny ip any any
Step 2: Verify that the router is configured to rate limiting the number of PIM register messages as shown in the example below.
ip pim rp-address 10.2.2.2
ip pim register-rate-limit nn
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: Verify that interfaces used for MSDP peering have an inbound ACL as shown in the example.
interface GigabitEthernet1/1
ip address x.1.28.8 255.255.255.0
ip access-group EXTERNAL_ACL_INBOUND in
ip pim sparse-mode
Step 3b: Verify that the ACL restricts MSDP peering to only known sources.
ip access-list extended EXTERNAL_ACL_INBOUND
permit tcp any any established
permit tcp host x.1.28.2 host x.1.28.8 eq 639
deny tcp any host x.1.28.8 eq 639 log
permit tcp host x.1.28.2 host 10.1.28.8 eq bgp
permit tcp host x.1.28.2 eq bgp host x.1.28.8
permit pim host x.1.28.2 pim host x.1.28.8
…
…
…
deny ip 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.

Internal Only - General

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.

```
R2(config)#ip access-list extended PIM_REGISTER_FILTER
R2(config-ext-nacl)#deny ip any 239.5.0.0 0.0.255.255
R2(config-ext-nacl)#permit ip host 10.1.2.6 any
R2(config-ext-nacl)#permit ip host 10.1.2.7 any
R2(config-ext-nacl)#deny ip any any
R2(config-ext-nacl)#exit
R2(config)#ip pim accept-register list PIM_REGISTER_FILTER
R2(config)#end
```

Step 2: Configure the RP to rate limit the number of multicast register messages.

```
R2(config)#ip pim register-rate-limit nn
```

Step 3: Configure the receive path or interface ACLs to only accepts MSDP packets from known MSDP peers.

```
R8(config)#ip access-list extended EXTERNAL_ACL_INBOUND
R8(config-ext-nacl)#permit tcp any any established
R8(config-ext-nacl)#permit tcp host x.1.28.2 host x.1.28.8 eq 639
R8(config-ext-nacl)#deny tcp any host x.1.28.8 eq 639
R8(config-ext-nacl)#permit tcp host x.1.28.2 host x.1.28.8 eq bgp
R8(config-ext-nacl)#permit tcp host x.1.28.2 eq bgp host x.1.28.8
R8(config-ext-nacl)#permit pim host x.1.28.2 host x.1.28.8
…
…
…
R8(config-ext-nacl)#deny ip any any
```

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

Internal Only - General

1.80 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-216626
RULE ID: SV-216626r531085
```

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.
ip pim rp-address 10.1.12.3
ip pim accept-register list PIM_REGISTER_FILTER
…
…
…
ip access-list extended PIM_REGISTER_FILTER
deny ip any 239.5.0.0 0.0.255.255
permit ip host 10.1.2.6 any
permit ip host 10.1.2.7 any
deny ip any any
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.

Internal Only - General

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.

```
R2(config)#ip access-list extended PIM_REGISTER_FILTER
R2(config-ext-nacl)#deny ip any 239.5.0.0 0.0.255.255
R2(config-ext-nacl)#permit ip host 10.1.2.6 any
R2(config-ext-nacl)#permit ip host 10.1.2.7 any
R2(config-ext-nacl)#deny ip any any
R2(config-ext-nacl)#exit
R2(config)#ip pim accept-register list PIM_REGISTER_FILTER
R2(config)#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

Internal Only - General

1.81 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-216627
RULE ID: SV-216627r531085
```

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 not allowed.
ip pim rp-address 10.2.2.2
ip pim accept-rp 10.2.2.2 FILTER_PIM_JOINS
…
…
…
ip access-list standard FILTER_PIM_JOINS
deny 239.8.0.0 0.0.255.255
permit any
!
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.
R2(config)#ip access-list standard PIM_JOIN_FILTER
R2(config-std-nacl)#deny 239.8.0.0 0.0.255.255
R2(config-std-nacl)#permit any
R2(config-std-nacl)#exit
R2(config)#ip pim accept-rp 10.2.2.2 PIM_JOIN_FILTER
R2(config)#end

Internal Only - General

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

Internal Only - General

1.82 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-216628
RULE ID: SV-216628r856196
```

Rationale:

When a new source starts transmitting in a PIM Sparse Mode network, the 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 rate limiting the number of PIM register messages.
ip pim rp-address 10.2.2.2
ip pim register-rate-limit nn
If the RP is not limiting PIM register messages, this is a finding.

Remediation:

Configure the RP to rate limit the number of multicast register messages.
R2(config)#ip pim register-rate-limit nn

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

Internal Only - General

1.83 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-216629
RULE ID: SV-216629r864157
```

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.
interface GigabitEthernet0/0
ip address 10.3.3.3 255.255.255.0
ip pim sparse-mode
ip igmp access-group IGMP_JOIN_FILTER
ip igmp version 3
Step 2: Verify that the Access Control List (ACL) denies unauthorized groups or permits only authorized groups. The example below denies all groups from 239.8.0.0/16 range.
ip access-list standard IGMP_JOIN_FILTER
deny 239.8.0.0 0.0.255.255
permit 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.

Internal Only - General

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.

```
R3(config)#ip access-list standard IGMP_JOIN_FILTER
R3(config-std-nacl)#deny 239.8.0.0 0.0.255.255
R3(config-std-nacl)#permit any
R3(config-std-nacl)#exit
```

Step 2: Apply the filter to all host facing interfaces.

```
R3(config)#int g0/0
R3(config-if)#ip igmp access-group IGMP_JOIN_FILTER
```

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)

Internal Only - General

1.84 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-216630
RULE ID: SV-216630r864158
```

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.
interface GigabitEthernet0/0
ip address 10.3.3.3 255.255.255.0
ip pim sparse-mode
ip igmp access-group IGMP_JOIN_FILTER
ip igmp version 3
Step 2: Verify that the Access Control List (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.
ip access-list extended IGMP_JOIN_FILTER
deny ip any 232.8.0.0 0.0.255.255
permit ip x.0.0.0 0.255.255.255 any
deny ip 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.

Internal Only - General

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.
R3(config)#ip access-list extended IGMP_JOIN_FILTER
R3(config-ext-nacl)#deny ip any 232.8.0.0 0.0.255.255
R3(config-ext-nacl)#permit ip x.0.0.0 0.255.255.255 any
R3(config-ext-nacl)#deny ip any any
R3(config-ext-nacl)#exit
Step 2: Apply the filter to all host facing interfaces.
R3(config)#int g0/0
R3(config-if)#ip igmp access-group IGMP_JOIN_FILTER

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)

Internal Only - General

1.85 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-216631
RULE ID: SV-216631r856199
```

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.
interface GigabitEthernet0/0
ip address 10.3.3.3 255.255.255.0
…
…
…
ip igmp limit nn
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.
R3(config)#int g0/0
R3(config-if)#ip igmp limit 2

Internal Only - General

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

Internal Only - General

1.86 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-216632
RULE ID: SV-216632r856200
```

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).
ip pim rp-address 10.2.2.2
ip pim 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.
R3(config)#ip pim spt-threshold infinity

Internal Only - General

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

Internal Only - General

1.87 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-216633
RULE ID: SV-216633r856201
```

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: Verify that interfaces used for MSDP peering have an inbound ACL as shown in the example.
interface GigabitEthernet1/1
ip address x.1.28.8 255.255.255.0
ip access-group EXTERNAL_ACL_INBOUND in
ip pim sparse-mode
Step 2: Verify that the ACL restricts MSDP peering to only known sources.
ip access-list extended EXTERNAL_ACL_INBOUND
permit tcp any any established
permit tcp host x.1.28.2 host x.1.28.8 eq 639
deny tcp any host x.1.28.8 eq 639 log
permit tcp host x.1.28.2 host 10.1.28.8 eq bgp
permit tcp host x.1.28.2 eq bgp host x.1.28.8
permit pim host x.1.28.2 pim host x.1.28.8
…
…
…
deny ip 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.

Internal Only - General

Remediation:

Configure the receive path or interface ACLs to only accept MSDP packets from known MSDP peers.
R8(config)#ip access-list extended EXTERNAL_ACL_INBOUND
R8(config-ext-nacl)#permit tcp any any established
R8(config-ext-nacl)#permit tcp host x.1.28.2 host x.1.28.8 eq 639
R8(config-ext-nacl)#deny tcp any host x1.28.8 eq 639
R8(config-ext-nacl)#permit tcp host x.1.28.2 host x.1.28.8 eq bgp
R8(config-ext-nacl)#permit tcp host x.1.28.2 eq bgp host x.1.28.8
R8(config-ext-nacl)#permit pim host x.1.28.2 host x.1.28.8
…
…
…
R8(config-ext-nacl)#deny ip any any

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)

Internal Only - General

1.88 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-216634
RULE ID: SV-216634r856202
```

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.
ip msdp peer x.1.28.8 remote-as 8
ip msdp password peer x.1.28.8 xxxxxxxxxxxx
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:
R2(config)#ip msdp password peer x.1.28.8 xxxxxxxxxxxx

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

Internal Only - General

1.89 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-216635
RULE ID: SV-216635r531085
```

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.

Internal Only - General

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.
ip msdp peer x.1.28.2 remote-as 2
ip msdp sa-filter in x.1.28.2 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.
ip access-list extended INBOUND_MSDP_SA_FILTER
deny ip any host 224.0.1.3
deny ip any host 224.0.1.24
deny ip any host 224.0.1.22
deny ip any host 224.0.1.2
deny ip any host 224.0.1.35
deny ip any host 224.0.1.60
deny ip any host 224.0.1.39
deny ip any host 224.0.1.40
deny ip any 232.0.0.0 0.255.255.255
deny ip any 239.0.0.0 0.255.255.255
deny ip 10.0.0.0 0.255.255.255 any
deny ip 127.0.0.0 0.255.255.255 any
deny ip 172.16.0.0 0.15.255.255 any
deny ip 192.168.0.0 0.0.255.255 any
permit ip any any
If the router is not configured with an import policy to filter undesirable SA multicast advertisements, this is a finding.

Internal Only - General

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.

```
R8(config)#ip access-list extended INBOUND_MSDP_SA_FILTER
R8(config-ext-nacl)#deny ip any host 224.0.1.3 ! Rwhod
R8(config-ext-nacl)#deny ip any host 224.0.1.24 ! Microsoft-ds
R8(config-ext-nacl)#deny ip any host 224.0.1.22 ! SVRLOC
R8(config-ext-nacl)#deny ip any host 224.0.1.2 ! SGI-Dogfight
R8(config-ext-nacl)#deny ip any host 224.0.1.35 ! SVRLOC-DA
R8(config-ext-nacl)#deny ip any host 224.0.1.60 ! hp-device-disc
R8(config-ext-nacl)#deny ip any host 224.0.1.39 ! Auto-RP
R8(config-ext-nacl)#deny ip any host 224.0.1.40 ! Auto-RP
R8(config-ext-nacl)#deny ip any 232.0.0.0 0.255.255.255 ! SSM range
R8(config-ext-nacl)#deny ip any 239.0.0.0 0.255.255.255 ! Admin scoped range
R8(config-ext-nacl)#deny ip 10.0.0.0 0.255.255.255 any ! RFC 1918 address range
R8(config-ext-nacl)#deny ip 127.0.0.0 0.255.255.255 any ! RFC 1918 address range
R8(config-ext-nacl)#deny ip 172.16.0.0 0.15.255.255 any ! RFC 1918 address range
R8(config-ext-nacl)#deny ip 192.168.0.0 0.0.255.255 any ! RFC 1918 address range
R8(config-ext-nacl)#permit ip any any
R8(config-ext-nacl)#exit
R8(config)#ip msdp sa-filter in x.1.28.2 list INBOUND_MSDP_SA_FILTER
```

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

Internal Only - General

1.90 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-216636
RULE ID: SV-216636r531085
```

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.
ip msdp peer 10.1.28.8 remote-as 8
ip msdp sa-filter out 10.1.28.8 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.
ip access-list extended OUTBOUND_MSDP_SA_FILTER
deny ip 10.0.0.0 0.255.255.255 any
permit ip any any
If the router is not configured with an export policy to filter local source-active multicast advertisements, this is a finding.

Internal Only - General

Remediation:

Configure the router with an export policy avoid global visibility of local multicast (S, G) states. The example below will prevent exporting multicast active sources belonging to the private network.
R8(config)#ip access-list extended OUTBOUND_MSDP_SA_FILTER
R8(config-ext-nacl)#deny ip 10.0.0.0 0.255.255.255 any
R8(config-ext-nacl)#permit ip any any
R8(config-ext-nacl)#exit
R8(config)#ip msdp sa-filter in x.1.28.2 list OUTBOUND_MSDP_SA_FILTER

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

Internal Only - General

1.91 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-216637
RULE ID: SV-216637r531085
```

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.
ip msdp peer x.1.28.2 remote-as nn
ip msdp sa-filter in 10.1.28.2 list MSDP_SA_FILTER
ip msdp sa-limit X.1.28.2 nnn
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.
R8(config)#ip msdp sa-limit x.1.28.2 nnn

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

Internal Only - General

1.92 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-216638
RULE ID: SV-216638r531085
```

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.
ip msdp peer x.44.2.34 connect-source Loopback12 remote-as nn
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.
R2(config)#ip msdp peer x.44.2.34 connect-source lo12 remote-as nn

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

Internal Only - General

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-000090 (Manual)	☐	☐
1.5	CISC-RT-000120 (Manual)	☐	☐
1.6	CISC-RT-000150 (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-000230 (Manual)	☐	☐
1.15	CISC-RT-000235 (Manual)	☐	☐
1.16	CISC-RT-000236 (Manual)	☐	☐
1.17	CISC-RT-000237 (Manual)	☐	☐
1.18	CISC-RT-000240 (Manual)	☐	☐
1.19	CISC-RT-000250 (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.20	CISC-RT-000260 (Manual)	☐	☐
1.21	CISC-RT-000270 (Manual)	☐	☐
1.22	CISC-RT-000280 (Manual)	☐	☐
1.23	CISC-RT-000290 (Manual)	☐	☐
1.24	CISC-RT-000300 (Manual)	☐	☐
1.25	CISC-RT-000310 (Manual)	☐	☐
1.26	CISC-RT-000320 (Manual)	☐	☐
1.27	CISC-RT-000330 (Manual)	☐	☐
1.28	CISC-RT-000340 (Manual)	☐	☐
1.29	CISC-RT-000350 (Manual)	☐	☐
1.30	CISC-RT-000360 (Manual)	☐	☐
1.31	CISC-RT-000370 (Manual)	☐	☐
1.32	CISC-RT-000380 (Manual)	☐	☐
1.33	CISC-RT-000390 (Manual)	☐	☐
1.34	CISC-RT-000391 (Manual)	☐	☐
1.35	CISC-RT-000392 (Manual)	☐	☐
1.36	CISC-RT-000393 (Manual)	☐	☐
1.37	CISC-RT-000394 (Manual)	☐	☐
1.38	CISC-RT-000395 (Manual)	☐	☐
1.39	CISC-RT-000396 (Manual)	☐	☐
1.40	CISC-RT-000397 (Manual)	☐	☐
1.41	CISC-RT-000398 (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.42	CISC-RT-000400 (Manual)	☐	☐
1.43	CISC-RT-000410 (Manual)	☐	☐
1.44	CISC-RT-000420 (Manual)	☐	☐
1.45	CISC-RT-000430 (Manual)	☐	☐
1.46	CISC-RT-000440 (Manual)	☐	☐
1.47	CISC-RT-000450 (Manual)	☐	☐
1.48	CISC-RT-000460 (Manual)	☐	☐
1.49	CISC-RT-000470 (Manual)	☐	☐
1.50	CISC-RT-000480 (Manual)	☐	☐
1.51	CISC-RT-000490 (Manual)	☐	☐
1.52	CISC-RT-000500 (Manual)	☐	☐
1.53	CISC-RT-000510 (Manual)	☐	☐
1.54	CISC-RT-000520 (Manual)	☐	☐
1.55	CISC-RT-000530 (Manual)	☐	☐
1.56	CISC-RT-000540 (Manual)	☐	☐
1.57	CISC-RT-000550 (Manual)	☐	☐
1.58	CISC-RT-000570 (Manual)	☐	☐
1.59	CISC-RT-000560 (Manual)	☐	☐
1.60	CISC-RT-000580 (Manual)	☐	☐
1.61	CISC-RT-000590 (Manual)	☐	☐
1.62	CISC-RT-000600 (Manual)	☐	☐
1.63	CISC-RT-000610 (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.64	CISC-RT-000620 (Manual)	☐	☐
1.65	CISC-RT-000630 (Manual)	☐	☐
1.66	CISC-RT-000640 (Manual)	☐	☐
1.67	CISC-RT-000650 (Manual)	☐	☐
1.68	CISC-RT-000660 (Manual)	☐	☐
1.69	CISC-RT-000670 (Manual)	☐	☐
1.70	CISC-RT-000730 (Manual)	☐	☐
1.71	CISC-RT-000740 (Manual)	☐	☐
1.72	CISC-RT-000750 (Manual)	☐	☐
1.73	CISC-RT-000760 (Manual)	☐	☐
1.74	CISC-RT-000770 (Manual)	☐	☐
1.75	CISC-RT-000780 (Manual)	☐	☐
1.76	CISC-RT-000790 (Manual)	☐	☐
1.77	CISC-RT-000800 (Manual)	☐	☐
1.78	CISC-RT-000810 (Manual)	☐	☐
1.79	CISC-RT-000820 (Manual)	☐	☐
1.80	CISC-RT-000830 (Manual)	☐	☐
1.81	CISC-RT-000840 (Manual)	☐	☐
1.82	CISC-RT-000850 (Manual)	☐	☐
1.83	CISC-RT-000860 (Manual)	☐	☐
1.84	CISC-RT-000870 (Manual)	☐	☐
1.85	CISC-RT-000880 (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.86	CISC-RT-000890 (Manual)	☐	☐
1.87	CISC-RT-000900 (Manual)	☐	☐
1.88	CISC-RT-000910 (Manual)	☐	☐
1.89	CISC-RT-000920 (Manual)	☐	☐
1.90	CISC-RT-000930 (Manual)	☐	☐
1.91	CISC-RT-000940 (Manual)	☐	☐
1.92	CISC-RT-000950 (Manual)	☐	☐

Internal Only - General

Appendix: Change History

Date	Version	Changes for this version
Aug 21, 2025	1.0.0	Initial CIS Release

Internal Only - General