

Cisco Applied Mitigation Bulletin: Identifying and Mitigating Exploitation of the Cisco IOS Software Tunnels Vulnerability

<http://www.cisco.com/warp/public/707/cisco-amb-20090923-tunnels.shtml>

Revision 1.1

Last Updated 2009 September 29 1500 UTC (GMT)

For Public Release 2009 September 23 1600 UTC (GMT)

Please provide your [feedback](#) on this document.

Contents

[Cisco Response](#)

[Device-Specific Mitigation and Identification](#)

[Additional Information](#)

[Revision History](#)

[Cisco Security Procedures](#)

[Related Information](#)

Cisco Response

This Applied Mitigation Bulletin is a companion document to the PSIRT Security Advisory *Cisco IOS Software Tunnels Vulnerability* and provides identification and mitigation techniques that administrators can deploy on Cisco network devices.

Vulnerability Characteristics

Cisco devices that are running certain versions of Cisco IOS Software are vulnerable to a Denial of Service (DoS) attack if configured for IP (GRE or IPinIP) tunnels and Cisco Express Forwarding (CEF). This vulnerability can be exploited remotely without authentication and without end-user interaction. Successful exploitation of this vulnerability may result

in a denial of service (DoS) condition. The attack vector for exploitation is through GRE packets using IP protocol 47, IPinIP packets using IP protocol 4, and IPv6 over IPv4 packets using IP protocol 41. An attacker could exploit this vulnerability using spoofed packets. This vulnerability has been assigned CVE identifiers CVE-2009-2872 and CVE-2009-2873.

Information about vulnerable, unaffected, and fixed software is available in the PSIRT Security Advisory, which is available at the following link: <http://www.cisco.com/warp/public/707/cisco-sa-20090923-tunnels.shtml>.

Mitigation Technique Overview

Cisco devices provide several countermeasures for this vulnerability. Administrators are advised to consider these protection methods to be general security best practices for infrastructure devices and the traffic that transits the network. This section of the document provides an overview of these techniques.

Cisco IOS Software can provide effective means of exploit prevention using the following methods:

- Infrastructure access control lists (iACLs)
- Unicast Reverse Path Forwarding (Unicast RPF)
- IP source guard (IPSG)

These protection mechanisms filter and drop, as well as verify the source IP address of, packets that are attempting to exploit this vulnerability.

The proper deployment and configuration of Unicast RPF provides an effective means of protection against attacks that use packets with spoofed source IP addresses. Unicast RPF should be deployed as close to all traffic sources as possible.

The proper deployment and configuration of IPSG provides an effective means of protection against spoofing attacks at the access layer.

Effective means of exploit prevention can also be provided by the Cisco ASA 5500 Series Adaptive Security Appliance, the Cisco PIX 500 Series Security Appliance, and the Firewall Services Module (FWSM) for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers using the following:

- Transit access control lists (tACLs)
- Unicast RPF

These protection mechanisms filter and drop, as well as verify the source IP address of, packets that are attempting to exploit this vulnerability.

Effective use of Cisco Intrusion Prevention System (IPS) event actions provides visibility into and protection against attacks that attempt to exploit this vulnerability.

Cisco IOS NetFlow records can provide visibility into network-based exploitation attempts.

Cisco IOS Software, Cisco ASA appliances, Cisco PIX security appliances, and FWSM firewalls can provide visibility through syslog messages and the counter values displayed in the output from **show** commands.

Risk Management

Organizations are advised to follow their standard risk evaluation and mitigation processes to determine the potential impact of this vulnerability. Triage refers to sorting projects and prioritizing efforts that are most likely to be successful. Cisco has provided documents that can help organizations develop a risk-based triage capability for their information security teams. [Risk Triage for Security Vulnerability Announcements](#) and [Risk Triage and Prototyping](#) can help organizations develop repeatable security evaluation and response processes.

Device-Specific Mitigation and Identification



Caution: The effectiveness of any mitigation technique depends on specific customer situations such as product mix, network topology, traffic behavior, and organizational mission. As with any configuration change, evaluate the impact of this configuration prior to applying the change.

Specific information about mitigation and identification is available for these devices:

- [Cisco IOS Routers and Switches](#)
- [Cisco IOS NetFlow](#)
- [Cisco ASA, PIX, and FWSM Firewalls](#)
- [Cisco Intrusion Prevention System](#)

Cisco IOS Routers and Switches

Mitigation: Infrastructure Access Control Lists

To protect infrastructure devices and minimize the risk, impact, and effectiveness of direct infrastructure attacks, administrators are advised to deploy infrastructure access control lists (iACLs) to perform policy enforcement of traffic sent to infrastructure equipment. Administrators can construct an iACL by explicitly permitting only authorized traffic sent to infrastructure devices in accordance with existing security policies and configurations. For the maximum protection of infrastructure devices, deployed iACLs should be applied in the ingress direction on all interfaces to which an IP address has been configured. An iACL workaround cannot provide complete protection against this vulnerability when the attack originates from a trusted source address.

The iACL policies below deny unauthorized IPv4 and IPv6 GRE packets on IP protocol 47, IPinIP packets on IP protocol 4, and IPv6 over IPv4 packets on IP protocol 41 that are sent to affected devices. In the following examples, 192.168.60.0/24 (IPv4) and 2001:DB8:1:128::/64 (IPv6) represent the IP address space that is used by the affected devices. Care should be taken to allow required traffic for routing and administrative access prior to denying all unauthorized traffic. Whenever possible, infrastructure address space should be distinct from the address space used for user and services segments. Using this addressing methodology will assist with the construction and deployment of iACLs.

Additional information about iACLs is in [Protecting Your Core: Infrastructure Protection Access Control Lists](#).

```
ip access-list extended Infrastructure-ACL-Policy

!
!-- Include explicit permit statements for trusted
!-- sources that require access on the vulnerable protocols
```

```
!  
permit gre host 192.168.100.1 192.168.60.0 0.0.0.255  
permit ipinip host 192.168.100.1 192.168.60.0 0.0.0.255  
permit 41 host 192.168.100.1 192.168.60.0 0.0.0.255
```

```
!  
!-- The following vulnerability-specific access control entries  
!-- (ACEs) can aid in identification of attacks  
!
```

```
deny gre any 192.168.60.0 0.0.0.255  
deny ipinip any 192.168.60.0 0.0.0.255  
deny 41 any 192.168.60.0 0.0.0.255
```

```
!  
!-- Explicit deny ACE for traffic sent to addresses configured within  
!-- the infrastructure address space  
!
```

```
deny ip any 192.168.60.0 0.0.0.255
```

```
!  
!-- Permit/deny all other Layer 3 and Layer 4 traffic in accordance  
!-- with existing security policies and configurations  
!  
!-- Apply iACL to interfaces in the ingress direction  
!
```

```
interface GigabitEthernet0/0  
ip access-group Infrastructure-ACL-Policy in
```

Note: This iACL will deny all vulnerable IPv6 packets to the affected devices.

```
ipv6 access-list IPv6-Infrastructure-ACL-Policy
```

```
!  
!-- The following vulnerability-specific access control entries  
!-- (ACEs) can aid in identification of attacks to global and  
!-- link local addresses  
!
```

```
deny 47 any 2001:DB8:1:128::/64  
deny 4 any 2001:DB8:1:128::/64
```

```
!  
!-- Permit other required traffic to the infrastructure address  
!-- range and allow IPv6 Neighbor Discovery packets, which  
!-- include Neighbor Solicitation packets and Neighbor  
!-- Advertisement packets
```

```

!
permit icmp any any nd-ns
permit icmp any any nd-na

!
!-- Explicit deny for all other IP traffic to the global
!-- infrastructure address range
!

deny ipv6 any 2001:db8:1:128::/64

!
!-- Permit/deny all other Layer 3 and Layer 4 traffic
!-- in accordance with existing security policies and
configurations
!
!-- Apply iACL to interfaces in the ingress direction
!

interface GigabitEthernet0/0
  ipv6 traffic-filter IPv6-Infrastructure-ACL-Policy in

```

Note that filtering with an interface access list will elicit the transmission of ICMP unreachable messages back to the source of the filtered traffic. Generating these messages could have the undesired effect of increasing CPU utilization on the device. In Cisco IOS Software, ICMP unreachable generation is limited to one packet every 500 milliseconds by default. ICMP unreachable message generation can be disabled using the interface configuration command **no ip unreachable** (IPv4) and **no ipv6 unreachable** (IPv6). ICMP unreachable rate limiting can be changed from the default using the global configuration command **ip icmp rate-limit unreachable** *interval-in-ms* (IPv4) and **ipv6 icmp error-interval** *interval-in-ms* [**bucket-size**] (IPv6).

Mitigation: Spoofing Protection

Unicast Reverse Path Forwarding

The vulnerability that is described in this document can be exploited by spoofed IP packets. Administrators can deploy and configure Unicast Reverse Path Forwarding (Unicast RPF) as a protection mechanism against spoofing.

Unicast RPF is configured at the interface level and can detect and drop packets that lack a verifiable source IP address. Administrators should not rely on Unicast RPF to provide complete spoofing protection because spoofed packets may enter the network through a Unicast RPF-enabled interface if an appropriate return route to the source IP address exists. Administrators are advised to take care to ensure that the appropriate Unicast RPF mode (loose or strict) is configured during the deployment of this feature because it can drop legitimate traffic that is transiting the network. In an enterprise environment, Unicast RPF might be enabled at the Internet edge and the internal access layer on the user-supporting Layer 3 interfaces.

Additional information is in the [Unicast Reverse Path Forwarding Loose Mode Feature Guide](#).

For additional information about the configuration and use of Unicast RPF, reference the [Understanding Unicast Reverse](#)

[Path Forwarding](#) Applied Intelligence white paper.

IP Source Guard

IP source guard (IPSG) is a security feature that restricts IP traffic on nonrouted, Layer 2 interfaces by filtering packets based on the DHCP snooping binding database and manually configured IP source bindings. Administrators can use IPSG to prevent attacks from an attacker who attempts to spoof packets by forging the source IP address and/or the MAC address. When properly deployed and configured, IPSG coupled with strict mode Unicast RPF provides the most effective means of spoofing protection for the vulnerability that is described in this document.

Additional information about the deployment and configuration of IPSG is in [Configuring DHCP Features and IP Source Guard](#).

Identification: Infrastructure Access Control Lists

After the administrator applies the iACL to an interface, the **show ip access-lists** command will identify the number of GRE packets on IP protocol 47, IPinIP packets on IP protocol 4, and IPv6 over IPv4 packets on IP protocol 41 that have been filtered on interfaces on which the iACL is applied. Administrators should investigate filtered packets to determine whether they are attempts to exploit this vulnerability. Example output for **show ip access-lists Infrastructure-ACL-Policy** follows:

```
router# show ip access-lists Infrastructure-ACL-Policy
Extended IP access list Infrastructure-ACL-Policy
 10 permit gre host 192.168.100.1 192.168.60.0 0.0.0.255 (27 matches)
 20 permit ipinip host 192.168.100.1 192.168.60.0 0.0.0.255 (13
matches)
 30 permit 41 host 192.168.100.1 192.168.60.0 0.0.0.255 (11 matches)
 40 deny gre any 192.168.60.0 0.0.0.255 (18 matches)
 50 deny ipinip any 192.168.60.0 0.0.0.255 (11 matches)
 60 deny 41 any 192.168.60.0 0.0.0.255 (6 matches)
 70 deny ip any 192.168.60.0 0.0.0.255 (19 matches)
router#
```

In the preceding example, access list Infrastructure-ACL-Policy has dropped **18 GRE** packets on **IP protocol 47** for access control list entry (ACE) line 40, **11 IPinIP** packets on **IP protocol 4** for ACE line 50, and **6 IPv6 over IPv4** packets on **IP protocol 41** for ACE line 60.

Example output for **show ipv6 access-list IPv6-Infrastructure-ACL-Policy** follows.

```
router#show ipv6 access-list IPv6-Infrastructure-ACL-Policy
IPv6 access list IPv6-Infrastructure-ACL-Policy
 deny 47 any 2001:DB8:1:128::/64 (12 matches) sequence 10
 deny 4 any 2001:DB8:1:128::/64 (7 matches) sequence 20
 permit icmp any any nd-ns sequence 30
 permit icmp any any nd-na sequence 40
 deny ipv6 any 2001:DB8:1:128::/64 sequence 50
router#
```

In the preceding example, access list IPv6-Infrastructure-ACL-Policy has dropped **12 GRE** packets on **IP protocol 47** for

access control list entry (ACE) line 10 and 7 IPinIP packets on IP protocol 4 for access control list entry (ACE) line 20.

For additional information about investigating incidents using ACE counters and syslog events, reference the [Identifying Incidents Using Firewall and IOS Router Syslog Events](#) Applied Intelligence white paper.

Administrators can use Embedded Event Manager to provide instrumentation when specific conditions are met, such as ACE counter hits. The Applied Intelligence white paper [Embedded Event Manager in a Security Context](#) provides additional details about how to use this feature.

Identification: Access List Logging

The **log** and **log-input** access control list (ACL) option will cause packets that match specific ACEs to be logged. The **log-input** option enables logging of the ingress interface in addition to the packet source and destination IP addresses and ports.



Caution: Access control list logging can be very CPU intensive and must be used with extreme caution. Factors that drive the CPU impact of ACL logging are log generation, log transmission, and process switching to forward packets that match log-enabled ACEs.

For Cisco IOS Software, the **ip access-list logging interval** *interval-in-ms* command can limit the effects of process switching induced by ACL logging. The **logging rate-limit** *rate-per-second* [**except** *loglevel*] command limits the impact of log generation and transmission.

The CPU impact from ACL logging can be addressed in hardware on the Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers with Supervisor Engine 720 or Supervisor Engine 32 using optimized ACL logging.

For additional information about the configuration and use of ACL logging, reference the [Understanding Access Control List Logging](#) Applied Intelligence white paper.

Identification: Spoofing Protection Using Unicast Reverse Path Forwarding

With Unicast RPF properly deployed and configured throughout the network infrastructure, administrators can use the **show cef interface** *type slot/port internal*, **show ip interface**, **show cef drop**, and **show ip traffic** commands to identify the number of packets that Unicast RPF has dropped.

Note: The **show command** | **begin** *regex* and **show command** | **include** *regex* command modifiers are used in the following examples to minimize the amount of output that administrators will need to parse to view the desired information. Additional information about command modifiers is in the [showcommand](#) sections of the Cisco IOS Configuration Fundamentals Command Reference.

```
router# show cef interface GigabitEthernet 0/0 internal | include drop
--          CLI Output Truncated          --
  ip verify: via=rx (allow default), acl=0, drop=18, sdrop=0
router#
```

Note: **show cef interface** *type slot/port internal* is a hidden command that must be fully entered at the command-line interface. Command completion is not available for it.

```

router# show ip interface GigabitEthernet 0/0 | begin verify
--      CLI Output Truncated      --
IP verify source reachable-via RX, allow default, allow self-ping
11 verification drops
0 suppressed verification drops
router#

```

```

router# show cef drop
CEF Drop Statistics
Slot  Encap_fail  Unresolved  Unsupported      No_route      No_adj
ChkSum_Err
RP           27           0           0           18
0           0
router#

```

```

router# show ip traffic

```

```

IP statistics:
Rcvd: 68051015 total, 2397325 local destination
      43999 format errors, 0 checksum errors, 33 bad hop count
      2 unknown protocol, 929 not a gateway
      21 security failures, 190123 bad options, 542768 with options
Opts: 352227 end, 452 nop, 36 basic security, 1 loose source route
      45 timestamp, 59 extended security, 41 record route
      53 stream ID, 3 strict source route, 40 alert, 45 cipso, 0 ump
      361634 other
Frag: 0 reassembled, 10008 timeouts, 56866 couldn't reassemble
      0 fragmented, 0 fragments, 0 couldn't fragment
Bcast: 64666 received, 0 sent
Mcast: 1589885 received, 2405454 sent
Sent: 3001564 generated, 65359134 forwarded
Drop: 4256 encapsulation failed, 0 unresolved, 0 no adjacency
      18 no route, 18 unicast RPF, 0 forced drop
      0 options denied
Drop: 0 packets with source IP address zero
Drop: 0 packets with internal loop back IP address
--      CLI Output Truncated      --
router#

```

In the preceding **show cef drop** and **show ip traffic** examples, Unicast RPF has dropped **18 IP packets** received globally on all interfaces with Unicast RPF configured because of the inability to verify the source address of the IP packets within the Forwarding Information Base of Cisco Express Forwarding.

Cisco IOS NetFlow

Identification: Traffic Flow Identification Using NetFlow Records

Administrators can configure Cisco IOS NetFlow on Cisco IOS routers and switches to aid in the identification of traffic flows that may be attempts to exploit the vulnerability. Administrators are advised to investigate flows to determine whether they are attempts to exploit the vulnerability or whether they are legitimate traffic flows.

```
router# show ip cache flow
IP packet size distribution (82889388 total packets):
  1-32  64  96  128  160  192  224  256  288  320  352  384  416
448  480
  .000 .359 .239 .084 .073 .017 .048 .019 .001 .000 .007 .002 .000 .000 .001

  512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
  .001 .000 .034 .023 .082 .000 .000 .000 .000 .000 .000
```

IP Flow Switching Cache, 4456704 bytes
 71 active, 65465 inactive, 9197254 added
 327737798 aget polls, 0 flow alloc failures
 Active flows timeout in 2 minutes
 Inactive flows timeout in 60 seconds

IP Sub Flow Cache, 533256 bytes
 71 active, 16313 inactive, 9197254 added, 9197254 added to flow
 0 alloc failures, 2161 force free
 1 chunk, 6 chunks added
 last clearing of statistics never

Protocol	Total	Flows	Packets	Bytes	Packets	Active(Sec)
Idle(Sec)	Flows	/Sec	/Flow	/Pkt	/Sec	/Flow
Flow						
TCP-Telnet	10125	0.0	8	43	0.0	2.2
24.5						
TCP-FTP	6612	0.0	2	51	0.0	0.8
47.8						
TCP-FTPD	2973	0.0	72	736	0.0	0.3
47.0						
TCP-WWW	149876	0.0	39	738	1.6	23.7
17.5						
TCP-SMTP	3696	0.0	1	80	0.0	0.2
48.1						
TCP-X	2853	0.0	1	43	0.0	0.0
50.6						
TCP-BGP	2761	0.0	1	43	0.0	0.0
50.1						
TCP-NNTP	2771	0.0	1	43	0.0	0.0
49.8						
TCP-Frag	2	0.0	105	1436	0.0	17.0
60.2						
TCP-other	6613046	1.8	6	246	11.7	1.0
28.7						
UDP-DNS	137965	0.0	4	67	0.1	21.7
50.3						
UDP-NTP	204064	0.0	1	76	0.0	5.9

58.2						
UDP-TFTP	4	0.0	1	28	0.0	0.0
60.3						
UDP-other	877188	0.2	32	125	7.9	19.5
52.2						
ICMP	797314	0.2	1	78	0.4	3.2
60.0						
IGMP	124136	0.0	2	39	0.0	58.9
42.2						
IPINIP	4	0.0	16	137	0.0	84.8
26.5						
IPv6INIP	5	0.0	12	143	0.0	69.7
31.8						
GRE	373	0.0	1	20	0.0	0.3
60.3						
IP-other	261415	0.0	10	95	0.7	92.3
16.4						
Total:	9197183	2.5	9	230	22.8	7.2
34.3						

SrcIf	SrcIPAddress	DstIf	DstIPAddress	Pr	SrcP
DstP Pkts					
Gi0/0	0.0.0.0	Null	255.255.255.255	11	0044
0043 8					
Gi0/0	192.168.208.127	Gi0/0	172.18.104.132	29	BC2C
1A29 7					
Gi0/0	192.168.208.80	Null	224.0.0.22	02	0000
0094 1					
Gi0/0	192.168.208.79	Null	224.0.0.22	02	0000
0094 2					
Gi0/0	192.168.208.127	Null	224.0.0.22	02	0000
0094 1					
Gi0/0	192.168.208.76	Gi0/1	192.168.150.70	29	A0C8
00A1 52					
Gi0/0	192.168.208.76	Gi0/1	192.168.150.65	29	A0C8
00A1 44					
Gi0/0	192.168.10.110	Gi0/1	192.168.60.10	2F	0000
0000 11					
Gi0/0	192.168.10.116	Gi0/1	192.168.60.10	2F	0000
0000 7					
Gi0/0	192.168.10.86	Gi0/1	192.168.60.15	04	0000
0000 2					
Gi0/0	192.168.10.1	Gi0/1	192.168.60.100	04	0000
0000 3					
Gi0/0	192.168.10.14	Gi0/1	192.168.60.100	2F	0000
0000 10					
Gi0/0	192.168.208.76	Gi0/0	192.168.128.22	11	A0C8
00A1 4					
Gi0/0	192.168.208.76	Local	192.168.128.20	11	A0C8
00A1 1					

```

Gi0/0      192.168.10.77   Gi0/1      192.168.60.8    2F 0000
0000      56
Gi0/0      192.168.208.76   Gi0/1      192.168.128.3   11 A0C8
00A1      1
Gi0/0      192.168.10.31    Gi0/1      192.168.60.8    04 0000
0000      1
Gi0/0      192.168.10.24    Gi0/1      192.168.60.10   2F 0000
0000      5
Gi0/0      192.168.10.36    Gi0/1      192.168.60.10   2F 0000
0000      1
Gi0/0      192.168.10.32    Gi0/1      192.168.60.8    04 0000
0000      1
router#

```

In the preceding example, there are multiple flows for **GRE on IP protocol 47 (Pr = hex value 2F)**, **IPinIP on IP protocol 4 (Pr = hex value 04)**, and **IPv6 over IPv4 on IP protocol 41 (Pr = hex value 29)**.

The packets in these flows may be spoofed and may indicate an attempt to exploit this vulnerability. Administrators are advised to compare these flows to baseline utilization for GRE traffic sent on IP protocol 47, IPinIP traffic sent on IP protocol 4, and IPv6 over IPv4 traffic sent on IP protocol 41 and also investigate the flows to determine whether they are sourced from untrusted hosts or networks.

To view only the traffic flows for GRE packets on IP protocol 47 (hex value 2F), IPinIP packets on IP protocol 4 (hex value 04), and IPv6 over IPv4 packets on IP protocol 41 (hex value 29) the command **show ip cache flow | include SrcIf|_(2F|04|29)_** will display the related NetFlow records as shown here:

GRE, IPinIP, and IPv6 over IPv4 Flows

```

router#show ip cache flow | include SrcIf|_(2F|04|29)_

SrcIf      SrcIPAddress      DstIf      DstIPAddress      Pr SrcP
DstP      Pkts
Gi0/0      192.168.10.18     Gi0/1      192.168.60.100   2F 0000
0000      4
Gi0/0      192.168.10.1     Gi0/1      192.168.60.100   04 0000
0000      8
Gi0/0      192.168.10.110   Gi0/1      192.168.60.1     2F 0000
0000      21
Gi0/0      192.168.10.73    Gi0/1      192.168.60.4     04 0000
0000      33
Gi0/0      192.168.10.70    Gi0/1      192.168.60.1     04 0000
0000      13
Gi0/0      192.168.10.59    Gi0/1      192.168.60.1     2F 0000
0000      9
Gi0/0      192.168.10.1     Gi0/1      192.168.60.10    04 0000
0000      8
Gi0/0      192.168.208.127  Gi0/0      172.18.104.132   29 BC2C
1A29      7
Gi0/0      192.168.208.76   Gi0/1      192.168.150.70   29 A0C8
00A1      52

```

```
Gi0/0          192.168.208.76  Gi0/1          192.168.150.65  29 A0C8
00A1          44
router#
```

Identification: Traffic Flow Identification Using IPv6 NetFlow Records

Administrators can configure Cisco IOS IPv6 NetFlow on Cisco IOS routers and switches to aid in the identification of traffic flows that may be attempts to exploit this vulnerability. Administrators are advised to investigate flows to determine whether they are attempts to exploit this vulnerability or whether they are legitimate traffic flows. The following example shows a sample IPv6 NetFlow configuration.

Note that NetFlow has been configured to export records to a workstation with the IPv4 address of 192.168.132.91 using UDP port 10000:

```
ipv6 flow-export destination 192.168.132.91 10000
ipv6 flow-capture packet-length
ipv6 flow-capture ttl
ipv6 flow-capture vlan-id
ipv6 flow-capture icmp
ipv6 flow-capture ip-id
ipv6 flow-capture mac-addresses
```

!

```
interface GigabitEthernet0/0
 ip address 192.168.208.20 255.255.255.0
 ipv6 address 2001:DB8:1:208::20/64
 ipv6 flow ingress
```

Administrators are advised to investigate flows to determine whether they are attempts to exploit this vulnerability or whether they are legitimate traffic flows.

```
router#show ipv6 flow cache
IP packet size distribution (50078919 total packets):
  1-32  64   96  128  160  192  224  256  288  320  352  384  416
448   480
  .000 .990 .001 .008 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
.000

  512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
  .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 475168 bytes
 8 active, 4088 inactive, 6160 added
1092984 age polls, 0 flow alloc failures
Active flows timeout in 30 minutes
Inactive flows timeout in 15 seconds

IP Sub Flow Cache, 33928 bytes
 16 active, 1008 inactive, 12320 added, 6160 added to flow
 0 alloc failures, 0 force free
```

1 chunk, 1 chunk added

SrcAddress	InpIf	DstAddress	OutIf	Prot	SrcPrt	DstPrt
Packets						
2001:DB...06::201	Gi0/0	2001:DB...28::20	Local	0x04	0x0000	0x0000
1464K						
2001:DB...6A:5BA6	Gi0/0	2001:DB...28::21	Gi0/1	0x3A	0x0000	0x8000
1191						
2001:DB...6A:5BA6	Gi0/0	2001:DB...134::3	Gi0/1	0x3A	0x0000	0x8000
1191						
2001:DB...6A:5BA6	Gi0/0	2001:DB...128::4	Gi0/1	0x3A	0x0000	0x8000
1192						
2001:DB...6A:5BA6	Gi0/0	2001:DB...128::2	Gi0/1	0x2F	0x0000	0x0000
1597						
2001:DB...6A:5BA6	Gi0/0	2001:DB...128::3	Gi0/1	0x3A	0x0000	0x8000
1192						
2001:DB...6A:5BA6	Gi0/0	2001:DB...146::3	Gi0/1	0x3A	0x0000	0x8000
1192						
2001:DB...6A:5BA6	Gi0/0	2001:DB...144::4	Gi0/1	0x3A	0x0000	0x8000
119						

To permit display of the full 128-bit IPv6 address, use the **terminal width 132** exec mode command.

In the preceding example, there are multiple IPv6 flows for GRE on IP protocol 47 (**Prot hex value = 0x2F**) and for IPinIP on IP protocol 4 (**Prot hex value = 0x04**). The packets in these flows may be spoofed and may indicate an attempt to exploit this vulnerability. Administrators are advised to compare these flows to baseline utilization for GRE traffic sent on IP protocol 47 and IPinIP traffic sent on IP protocol 4 and also investigate the flows to determine whether they are sourced from untrusted hosts or networks.

As shown in the following example, to view only the traffic flows for IPv6 GRE packets on IP protocol 47 (**Prot hex value = 0x2F**) and for IPinIP on IP protocol 4 (**Prot hex value = 0x04**), use the **show ipv6 flow cache | include SrcAddress|(0x2F|0x04)** command to display the related NetFlow records:

```
router#show ipv6 flow cache | include SrcAddress|(0x2F|0x04)
-
SrcAddress      InpIf      DstAddress      OutIf      Prot  SrcPrt  DstPrt
Packets
2001:DB...06::201 Gi0/0      2001:DB...28::20 Local      0x2F  0x0000  0x0000
56K
2001:DB...06::201 Gi0/0      2001:DB...28::20 Local      0x04  0x0000  0x0000
74
2001:DB...06::201 Gi0/0      2001:DB...28::20 Local      0x2F  0x0000  0x0000
356
router#
```

Cisco ASA, PIX, and FWSM Firewalls

Mitigation: Transit Access Control Lists

To protect the network from traffic that enters the network at ingress access points, which may include Internet connection points, partner and supplier connection points, or VPN connection points, administrators are advised to deploy tACLs to perform policy enforcement. Administrators can construct a tACL by explicitly permitting only authorized traffic to enter the network at ingress access points or permitting authorized traffic to transit the network in accordance with existing security policies and configurations. A tACL workaround cannot provide complete protection against this vulnerability when the attack originates from a trusted source address.

The tACL policies below deny unauthorized IPv4 and IPv6 GRE packets on IP protocol 47, IPinIP packets on IP protocol 4, and IPv6 over IPv4 packets on IP protocol 41 that are sent to affected devices. In the following examples, 192.168.60.0/24 (IPv4) and 2001:DB8:1:128::/64 (IPv6) represent the IP address space that is used by the affected devices. Care should be taken to allow required traffic for routing and administrative access prior to denying all unauthorized traffic.

Additional information about tACLs is in [Transit Access Control Lists: Filtering at Your Edge](#).

```
!  
!-- Include explicit permit statements for trusted sources  
!-- that require access on the vulnerable protocols  
!  
  
access-list tACL-Policy extended permit gre host 192.168.100.1  
192.168.60.0 255.255.255.0  
access-list tACL-Policy extended permit ipinip host 192.168.100.1  
192.168.60.0 255.255.255.0  
access-list tACL-Policy extended permit 41 host 192.168.100.1  
192.168.60.0 255.255.255.0  
  
!  
!-- The following vulnerability-specific access control entries  
!-- (ACEs) can aid in identification of attacks  
!  
  
access-list tACL-Policy extended deny gre any 192.168.60.0 255.255.255.0  
access-list tACL-Policy extended deny ipinip any 192.168.60.0  
255.255.255.0  
access-list tACL-Policy extended deny 41 any 192.168.60.0 255.255.255.0  
  
!  
!-- Permit/deny all other Layer 3 and Layer 4 traffic in accordance  
!-- with existing security policies and configurations  
!  
!-- Explicit deny for all other IP traffic  
!  
  
access-list tACL-Policy extended deny ip any any  
  
!  
!-- Apply tACL to interface(s) in the ingress direction  
!
```

```
access-group tACL-Policy in interface outside
```

Note: This tACL will deny all vulnerable IPv6 packets to the affected devices.

```
!-- The following vulnerability-specific access control entries
!-- (ACEs) can aid in identification of attacks
!

ipv6 access-list IPv6-tACL-Policy deny gre any 2001:db8:1:128::/64
ipv6 access-list IPv6-tACL-Policy deny ipinip any 2001:db8:1:128::/64

!
!-- Permit/deny all other Layer 3 and Layer 4 traffic in accordance
!-- with existing security policies and configurations
!
!-- Explicit deny for all other IP traffic
!

ipv6 access-list IPv6-Transit-ACL-Policy deny ip any any

!
!-- Apply tACL to interfaces in the ingress direction
!

access-group IPv6-Transit-ACL-Policy in interface outside
```

Mitigation: Spoofing Protection Using Unicast Reverse Path Forwarding

The vulnerability that is described in this document can be exploited by spoofed IP packets. Administrators can deploy and configure Unicast RPF as a protection mechanism against spoofing.

Unicast RPF is configured at the interface level and can detect and drop packets that lack a verifiable source IP address. Administrators should not rely on Unicast RPF to provide complete spoofing protection because spoofed packets may enter the network through a Unicast RPF-enabled interface if an appropriate return route to the source IP address exists. In an enterprise environment, Unicast RPF might be enabled at the Internet edge and at the internal access layer on the user-supporting Layer 3 interfaces.

For additional information about the configuration and use of Unicast RPF, reference the Cisco Security Appliance Command Reference for [ipverify reverse-path](#) and the [Understanding Unicast Reverse Path Forwarding](#) Applied Intelligence white paper.

Identification: Transit Access Control Lists

After the tACLs have been applied to an interface, administrators can use the **show access-list** command (IPv4) and **show ipv6 access-list** command (IPv6) to identify the number of GRE packets on IP protocol 47, IPinIP packets on IP protocol 4, and IPv6 over IPv4 packets on IP protocol 41 that have been filtered. Administrators are advised to investigate filtered packets to determine whether they are attempts to exploit this vulnerability. Example outputs below display the respective show access-list outputs:

```

firewall# show access-list tACL-Policy
access-list tACL-Policy; 7 elements
access-list tACL-Policy line 1 extended permit gre host 192.168.100.1
192.168.60.0 255.255.255.0 (hitcnt=56)
access-list tACL-Policy line 2 extended permit ipinip host
192.168.100.1 192.168.60.0 255.255.255.0 (hitcnt=9)
access-list tACL-Policy line 3 extended permit 41 host 192.168.100.1
192.168.60.0 255.255.255.0 (hitcnt=11)
access-list tACL-Policy line 4 extended deny gre any 192.168.60.0
255.255.255.0 (hitcnt=27)
access-list tACL-Policy line 5 extended deny ipinip any 192.168.60.0
255.255.255.0 (hitcnt=13)
access-list tACL-Policy line 6 extended deny 41 any 192.168.60.0
255.255.255.0 (hitcnt=7)
access-list tACL-Policy line 7 extended deny ip any any (hitcnt=123)
firewall#

```

In the preceding example, access list *tACL-Policy* has dropped **27 GRE** packets on **IP protocol 47**, **13 IPinIP** packets on **IP protocol 4**, and **7 IPv6 over IPv4** packets on **IP protocol 41** received from an untrusted host or network. In addition, syslog message *106023* can provide valuable information, which includes the source and destination IP address, the source and destination port numbers, and the IP protocol for the denied packet.

```

firewall#show ipv6 access-list IPv6-tACL-Policy
ipv6 access-list IPv6-TACL-Policy; 3 elements
ipv6 access-list IPv6-tACL-Policy line 1 deny gre any 2001:
db8:1:128::/64 (hitcnt=15)
ipv6 access-list IPv6-tACL-Policy line 2 deny ipinip any 2001:
db8:1:128::/64 (hitcnt=12)
ipv6 access-list IPv6-tACL-Policy line 3 deny ip any any (hitcnt=9)
firewall#

```

In the preceding example, access list *IPv6-tACL-Policy* has dropped **15 GRE** packets on **IP protocol 47** and **12 IPinIP** packets on **IP protocol 4** received from an untrusted host or network. In addition, syslog message *106023* can provide valuable information, which includes the source and destination IP address, the source and destination port numbers, and the IP protocol for the denied packet.

Identification: Firewall Access List Syslog Messages

Firewall syslog message *106023* will be generated for packets denied by an ACE that does not have the **log** keyword present. Additional information about this syslog message is in [Cisco Security Appliance System Log Message - 106023](#).

Information about configuring syslog for the Cisco ASA 5500 Series Adaptive Security Appliance or the Cisco PIX 500 Series Security Appliance is in [Monitoring the Security Appliance - Configuring and Managing Logs](#). Information about configuring syslog on the FWSM for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers is in [Monitoring the Firewall Services Module](#).

In the following example, the **show logging | grep regex** command extracts syslog messages from the logging buffer on the firewall. These messages provide additional information about denied packets that could indicate potential attempts to

exploit the vulnerability that is described in this document. It is possible to use different regular expressions with the **grep** keyword to search for specific data in the logged messages.

Additional information about regular expression syntax is in [Creating a Regular Expression](#).

```
firewall# show logging | grep 106023
      Sep 08 2009 08:58:33: %ASA-4-106023: Deny protocol 47
src outside:192.168.60.1
      dst inside:192.168.60.164 by access-group
"tACL-Policy"
      Sep 08 2009 08:58:33: %ASA-4-106023: Deny protocol 4
src outside:192.168.1.1
      dst inside:192.168.60.54 by access-group "tACL-
Policy"
      Sep 08 2009 08:58:33: %ASA-4-106023: Deny protocol 47
src outside:192.168.1.1
      dst inside:192.168.60.17 by access-group "tACL-
Policy"
      Sep 08 2009 08:58:33: %ASA-4-106023: Deny protocol 04
src outside:192.168.60.233
      dst inside:192.168.60.91 by access-group "tACL-
Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 4
src outside:192.168.10.53
      dst inside:192.168.60.10 by access-group "tACL-
Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 4
src outside:192.168.10.1
      dst inside:192.168.60.10 by access-group "tACL-
Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 4
src outside:192.168.10.1
      dst inside:192.168.60.10 by access-group "tACL-
Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 41
src outside:192.168.11.1
      dst inside:192.168.60.70 by access-group "tACL-
Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 47
src outside:2001:db8:5:122::
      dst inside:2001:db8:1:128:: by access-group
"IPv6-tACL-Policy"
      Sep 08 2009 09:09:31: %ASA-4-106023: Deny protocol 4
src outside:2001:db8:5:122::
      dst inside:2001:db8:1:128:: by access-group
"IPv6-tACL-Policy"

firewall#
```

In the preceding example, the messages logged for the tACL *tACL-Policy* some potentially spoofed **GRE** packets for **IP**

protocol 47, IPinIP packets for **IP protocol 4**, and **IPv6 over IPv4** packets for **IP protocol 41** sent to the address block assigned to the infrastructure devices.

Additional information about syslog messages for ASA and PIX security appliances is in [Cisco Security Appliance System Log Messages](#). Additional information about syslog messages for the FWSM is in [Catalyst 6500 Series Switch and Cisco 7600 Series Router Firewall Services Module Logging System Log Messages](#).

For additional information about investigating incidents using syslog events, reference the [Identifying Incidents Using Firewall and IOS Router Syslog Events](#) Applied Intelligence white paper.

Identification: Spoofing Protection Using Unicast Reverse Path Forwarding

Firewall syslog message *106021* will be generated for packets denied by Unicast RPF. Additional information about this syslog message is in [Cisco Security Appliance System Log Message - 106021](#).

Information about configuring syslog for the Cisco ASA 5500 Series Adaptive Security Appliance or the Cisco PIX 500 Series Security Appliance is in [Monitoring the Security Appliance - Configuring and Managing Logs](#). Information about configuring syslog on the FWSM for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers is in [Monitoring the Firewall Services Module](#).

In the following example, the **show logging | grep regex** command extracts syslog messages from the logging buffer on the firewall. These messages provide additional information about denied packets that could indicate potential attempts to exploit these the vulnerability that is described in this document. It is possible to use different regular expressions with the **grep** keyword to search for specific data in the logged messages.

Additional information about regular expression syntax is in [Creating a Regular Expression](#).

```
firewall# show logging | grep 106021
Sep 08 2009 11:27:15: %ASA-1-106021: Deny IP reverse path check
from
                192.168.60.213 to 192.168.60.10 on interface
outside
Sep 08 2009 11:27:15: %ASA-1-106021: Deny IP reverse path check
from
                192.168.60.208 to 192.168.60.10 on interface
outside
Sep 08 2009 11:27:15: %ASA-1-106021: Deny IP reverse path check
from
                192.168.60.1 to 192.168.60.10 on interface
outside
Sep 08 2009 11:27:15: %ASA-1-106021: Deny IP reverse path check
from
                192.168.60.161 to 192.168.60.10 on interface outside
Sep 08 2009 11:27:45: %ASA-1-106021: Deny GRE reverse path
check from
                192.168.60.1 to 192.168.60.10 on interface
outside
Sep 08 2009 11:27:45: %ASA-1-106021: Deny GRE reverse path
```

```

check from
                192.168.60.1 to 192.168.60.10 on interface
outside
    Sep 08 2009 11:27:45: %ASA-1-106021: Deny GRE reverse path
check from
                192.168.60.1 to 192.168.60.10 on interface
outside
    Sep 08 2009 11:27:45: %ASA-1-106021: Deny GRE reverse path
check from
                192.168.60.1 to 192.168.60.10 on interface
outside
    Sep 08 2009 11:27:45: %ASA-1-106021: Deny GRE reverse path
check from
                192.168.60.41 to 192.168.60.10 on interface
outside
firewall#

```

The **show asp drop** command can also identify the number of packets that the Unicast RPF feature has dropped, as shown in the following example:

```

firewall# show asp drop frame rpf-violated
    Reverse-path verify failed (rpf-violated)                33
firewall#

```

In the preceding example, Unicast RPF has dropped **33 IP packets** received on interfaces with Unicast RPF configured. Absence of output indicates that the Unicast RPF feature on the firewall has not dropped packets.

For additional information about debugging accelerated security path dropped packets or connections, reference the Cisco Security Appliance Command Reference for [show asp drop](#).

Cisco Intrusion Prevention System

Mitigation: Cisco IPS Signature Event Actions

Administrators can use Cisco Intrusion Prevention System (IPS) appliances and services modules to provide threat detection and help prevent attempts to exploit this vulnerability. For sensors running Cisco IPS version 6.x or 5.x that are up to date with signature updates, this vulnerability can be detected. Due to the nature of this vulnerability the Signature Update Number, Signature Name, and/or Signature ID cannot be provided at this time. The signature is enabled by default and is configured with a default event action of **deny-packet-inline**.

Administrators can configure Cisco IPS sensors to perform an event action when an attack is detected. The configured event action performs preventive or deterrent controls to help protect against an attack that is attempting to exploit the vulnerability that is described in this document.

Exploits that use spoofed IP addresses may cause a configured event action to inadvertently deny traffic from trusted sources.

Cisco IPS sensors are most effective when deployed in inline protection mode combined with the use of an event action. Automatic Threat Prevention for Cisco IPS 6.x sensors that are deployed in inline protection mode provides threat

prevention against an attack that is attempting to exploit the vulnerability that is described in this document. Threat prevention is achieved through a default override that performs an event action for triggered signatures with a *riskRatingValue* greater than 90.

Cisco IPS 5.x sensors that are deployed in inline protection mode require an event action configured on a per-signature basis. Alternatively, administrators can configure an override that can perform an event action for any signatures that are triggered and are calculated as a high-risk threat. Using an event action on sensors deployed in inline protection mode provides the most effective exploit prevention.

For additional information about the risk rating and threat rating calculation, reference [Risk Rating and Threat Rating: Simplify IPS Policy Management](#).

Additional Information

THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS AND DOES NOT IMPLY ANY KIND OF GUARANTEE OR WARRANTY, INCLUDING THE WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR USE. YOUR USE OF THE INFORMATION ON THE DOCUMENT OR MATERIALS LINKED FROM THE DOCUMENT IS AT YOUR OWN RISK. CISCO RESERVES THE RIGHT TO CHANGE OR UPDATE THIS DOCUMENT AT ANY TIME.

Revision History

Revision 1.1	2009-September-29	Updated to include IPv6 over IPv4 traffic sent on IP protocol 41 as a vector of exploitation.
Revision 1.0	2009-September-23	Initial public release

Cisco Security Procedures

Complete information on reporting security vulnerabilities in Cisco products, obtaining assistance with security incidents, and registering to receive security information from Cisco, is available on Cisco's worldwide website at http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html. This includes instructions for press inquiries regarding Cisco security notices. All Cisco security advisories are available at <http://www.cisco.com/go/psirt>.

Related Information

- [Cisco Applied Mitigation Bulletins](#)
- [Cisco Guide to Harden Cisco IOS Devices](#)
- [Security Intelligence Operations](#)
- [Cisco IOS NetFlow - Home Page on Cisco.com](#)
- [Cisco IOS NetFlow White Papers](#)
- [NetFlow Performance Analysis](#)
- [Cisco Network Foundation Protection White Papers](#)
- [Cisco Network Foundation Protection Presentations](#)
- [A Security-Oriented Approach to IP Addressing](#)

- [Cisco Firewall Products - Home Page on Cisco.com](#)
 - [Cisco 6.x Intrusion Prevention System](#)
 - [Cisco IPS 6.x Signature Downloads](#)
 - [Cisco IPS Signature Search Page](#)
-

Help us help you.

Please rate this document.

- Excellent
- Good
- Average
- Fair
- Poor

This document solved my problem.

- Yes
- No
- Just browsing

Suggestions for improvement:

(256 character limit)

[Home](#)

[How to Buy](#)

[Login](#)

[Profile](#)

[Feedback](#)

[Site Map](#)

[Help](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 - 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)