

# Cisco Security Response: Internet Key Exchange Resource Exhaustion Attack

Document ID: 70810

<http://www.cisco.com/warp/public/707/cisco-sr-20060726-ike.shtml>

## Revision 2.3

**Last Updated** 2008 July 25 1900 UTC (GMT)

**For Public Release** 2006 July 26 1600 UTC (GMT)

---

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

---

[Cisco Response](#)  
[Additional Information](#)  
[Revision History](#)  
[Cisco Security Procedures](#)

---

## Cisco Response

This is a Cisco response to an advisory published by an unaffiliated third party, Roy Hills, of NTA Monitor Ltd posted as of July 26, 2006 at <http://www.nta-monitor.com/posts/2006/07/cisco-concentrator-dos.html>, and entitled: Cisco VPN Concentrator IKE resource exhaustion DoS.

This issue is being tracked by the following Cisco Bug IDs:

- [CSCse70811](#) ([registered](#) customers only) (Cisco IOS® software)
- [CSCse89808](#) ([registered](#) customers only) (Cisco VPN 3000 Concentrators)
- [CSCsb51032](#) ([registered](#) customers only) and [CSCsb50996](#) ([registered](#) customers only) (Cisco PIX firewalls running pre-7.x code)
- [CSCse92254](#) ([registered](#) customers only) (Cisco PIX firewalls and Cisco ASA appliances running 7.x code)
- [CSCse92527](#) ([registered](#) customers only) (Cisco Firewall Services Module [FWSM] for Cisco Catalyst 6500 switches and Cisco 7600 Series routers)
- [CSCse96516](#) ([registered](#) customers only) (Cisco SAN-OS on MDS devices)
- [CSCek52553](#) ([registered](#) customers only) (Cisco IOS XR software)

We thank Roy Hills from NTA Monitor Ltd for reporting this issue to Cisco. We greatly appreciate the opportunity to work with researchers on security vulnerabilities, and welcome the opportunity to review and assist in product reports.

## Additional Information

### Vulnerability Impact Overview

Cisco devices which implement the IKE version 1 protocol may be vulnerable to an attack that attempts to exploit limitations of the IKE version 1 protocol in order to deplete available resources to negotiate IKE SAs

(Security Associations) and block legitimate IPSec peers from establishing new IKE SAs or rekey existing IKE SAs. The vulnerability is inherent to the IKE version 1 protocol and is not specific to any vendor implementation.

This attack may cause IKE resource depletion and VPN termination devices that are under attack may not allow legitimate VPN connection requests or in some cases drop already established connections during rekey. IOS devices without [Call Admission Control for IKE](#) configured may exhibit high CPU utilization if attacked and may not perform as expected. IKE CAC is not enabled by default on those IOS releases that support it.

**Networks designed with multiple dedicated VPN termination points at different ingress locations should only be minimally impacted as long as VPN client devices are configured to connect to alternate VPN termination points if their primary VPN termination point is not available.**

Organizations should follow their standard risk mitigation process to determine the potential impact of this issue. A document that may be used to aid in risk triage is available at <http://www.cisco.com/web/about/security/intelligence/vulnerability-risk-triage.html>.

## Mitigation Overview

On IOS platforms, the [Call Admission Control for IKE](#) feature will aid in preventing an attacker from exploiting this issue to exhaust CPU resources and impact service for non-IKE traffic. For Site-to-Site IPSec tunnels, Interface access-lists may also provide some mitigation though the source IP addresses of IKE messages may be spoofed, which may limit the effectiveness of these access-lists.

On PIX, ASA, Firewall Service module Firewalls, and the VPN 3000 Series Concentrators there are built-in resource protection mechanisms to prevent an attacker from exploiting this issue to exhaust CPU resources and impact service for non-IKE traffic.

## Device-Specific Mitigation and Identification

### Internet Edge Router

The Internet Edge router does not terminate IPSec but is positioned to mitigate and detect attempted exploitation of this issue.

#### *Mitigation*

- Interface Access-lists

Internet Edge routers can be configured with interface access-lists to drop packets that can be used to exploit this issue. Since the source IP address of peer VPN devices must be known beforehand, interface access-lists are most effective in protecting site-to-site IPSec tunnels.

The following ACL is specifically designed to permit IKE (UDP/500) traffic only from known IPSec peer IP addresses and drop IKE from all other IP addresses. The ACL should be applied to Internet facing IPv4 interfaces and should include topology-specific filters. In this example, the protected VPN head-end's IP address is 192.0.2.1 and IP addresses 192.0.2.100 and 192.0.2.101 are known site-to-site IPSec peers. All other IP addresses are denied.

```
access-list 150 permit udp host 192.0.2.100 host 192.0.2.1 eq 500
access-list 150 permit udp host 192.0.2.101 host 192.0.2.1 eq 500
access-list 150 deny udp any any eq 500
access-list 150 permit ip any any
```

```
interface serial 2/0
 ip access-group 150 in
```

These ACL statements may be deployed on the IOS Internet Edge router as part of a transit access-list which will protect the router itself and devices deployed behind it. Further information about transit ACLs is available in the white paper "Transit Access Control Lists: Filtering at Your Edge", available at <http://www.cisco.com/warp/public/707/tacl.html>.

Please note that filtering traffic with an interface access list will elicit the transmission of ICMP unreachable messages back to the source of the filtered traffic. This could have the undersired side effect high CPU utilization since the device needs to generate these ICMP unreachable messages. In IOS, ICMP unreachable generation is limited to one packet per 500 ms. ICMP unreachable generation can be disabled using the interface configuration command **no ip unreachable**. ICMP unreachable rate-limiting can be changed from the default 1 per 500 ms using the global configuration command **ip icmp rate-limit unreachable <1-4294967295 millisecond>**.

- Anti-Spoofing

This issue can be exploited by spoofed packets. Anti-spoof protection in the form of interface access-lists or unicast Reverse Path Forwarding can provide limited mitigation if properly configured. This mitigation should not be relied upon to provide complete mitigation since spoofed packets may still enter the network from the interface expected by uRPF or allowed by anti-spoofing access-lists. Also care must be taken to ensure that the appropriate uRPF mode (loose or strict) is configured to ensure that legitimate packets are not dropped.

Additional information about unicast Reverse Path Forwarding is available at [http://www.cisco.com/en/US/docs/ios/12\\_2t/12\\_2t13/feature/guide/ft\\_urpf.html](http://www.cisco.com/en/US/docs/ios/12_2t/12_2t13/feature/guide/ft_urpf.html).

- Committed Access Rate

Committed Access Rate (CAR) can be used to police or rate-limit IKE (UDP/500) traffic destined to a protected VPN head-end device. Unlike interface access-lists where the source IP addresses of peer VPN devices must be known beforehand, CAR can be applied to all IKE traffic preventing CPU processing power and memory resources from being fully consumed by incoming IKE requests. Since CAR is applied to all IKE traffic destined for the protected VPN head-end, legitimate IKE messages including initiator packets and tunnel rekeys will be rate-limited as well as attempts to exploit this issue.

The following CAR policy is specifically designed to rate-limit all IKE (UDP/500) traffic destined to the IOS VPN head-end's IP address of 192.0.2.1. The CAR policy is applied to the Edge router's Internet facing IPv4 interface.

```
access-list 111 permit udp any host 192.0.2.1 eq isakmp
interface serial 2/0
    rate-limit input access-group 111 32000 6000 12000 conform-action transmit exceed-
```

The rate-limit policy is applied to traffic matching access-list 111. Traffic at rates below 32,000 bits per second (bps) will be forwarded normally but above 32,000 bps packets will be dropped. 32,000 bps is loosely equivalent to 35 IKE messages per second. Burst and extend burst values are calculated using recommended [burst values](#).

**Network conditions, traffic, and device load may lead to significantly different results in other environments, therefore this information is provided as an example and should serve as the basis for testing in your environment, not as recommended settings for production devices.**

For additional information about traffic policing and Committed Access Rate, refer to:

– Policing and Shaping Overview

[http://www.cisco.com/en/US/docs/ios/12\\_2/qos/configuration/guide/qcfcpolsh.html](http://www.cisco.com/en/US/docs/ios/12_2/qos/configuration/guide/qcfcpolsh.html)

– Configuring Committed Access Rate

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products\\_configuration\\_guide\\_chapter09186a0080](http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a0080)

– Cisco IOS Embedded Event Manager

[http://www.cisco.com/en/US/products/ps6815/products\\_ios\\_protocol\\_group\\_home.html](http://www.cisco.com/en/US/products/ps6815/products_ios_protocol_group_home.html)

This mitigation method uses EEM (Embedded Event Manager) on the Internet edge router to protect IOS VPN routers that do not have the CAC feature enabled. This solution requires either 12.4(7a), 12.2(18)SXF4 or later IOS software.

Follow these steps:

1. Enable NetFlow on the IOS screening device.

2. Configure two numeric extended access-lists which will be used by EEM to rate-limit IKE messages.
3. Use EEM on the screening device to periodically check NetFlow information.
4. EEM will create additions to the numeric access-list which will be used to rate-limit excessive IKE messages to IOS VPN routers. EEM will also create and send Syslog messages when it detects excessive IKE message flows or an excessive number of packets in an IKE flow.

**Note:** Customers who have implemented GDOI (Group Domain of Interpretation) for either Secure Multicast or Dynamic Multipoint VPN (DMVPN) can implement the previous mitigation techniques as needed. Any references to port 500/UDP or "isakmp" in an access-list would need to be replaced with port 848/UDP, the GDOI port. Additional information on GDOI can be found in the [Related Information](#) section of this Security Response.

### Identification

- Interface Access-lists

Once the interface access-list is deployed, the command **show access-list 150** can be used to identify the number of packets being dropped. Dropped packets should be investigated to determine if they are attempts to exploit the issue.

Example output for **show access-list 150**:

```
Edge-Router#show access-list 150
Extended IP access list 150
 10 permit udp host 192.0.2.100 host 192.0.2.1 eq 500 (100 matches)
 20 permit udp host 192.0.2.101 host 192.0.2.1 eq 500 (141 matches)
 30 deny udp any any eq 500 (700 matches)
 40 permit ip any any (26154 matches)
```

In the above example 700 UDP/500 packets from untrusted IP addresses have been dropped by the access-list configured on interface Serial 2/0.

- Committed Access Rate

Once CAR is deployed, the commands **show access-list 111** and **show interface serial 2/0 rate-limit** can be used to identify packets sent to the VPN head-end device, conforming and non-conforming bits per second rates, and packets passed and dropped. High non-conforming traffic rates and high packet drop rates should be investigated to determine if they are attempts to exploit the issue.

```
edge-router#show access-list 111
Extended IP access list 111
 10 permit udp any host 192.0.2.1 eq isakmp (99769 matches)

edge-router#show interface serial 2/0 rate-limit
interface serial 2/0
  Input
  matches: access-group 111
  params: 32000 bps, 6000 limit, 12000 extended limit
  conformed 9459 packets, 1191834 bytes; action: transmit
  exceeded 281539 packets, 35473914 bytes; action: drop
  last packet: 0ms ago, current burst: 11954 bytes
  last cleared 00:06:21 ago, conformed 24000 bps, exceeded 743000 bps
```

In the above example the conforming traffic rate is 24,000 bits per second, non-conforming traffic rate is 743,000 bits per second. 9,459 packets have been transmitted and 281,539 packets have been dropped. High drop rates should be investigated to ensure configured rate-limit values are not dropping legitimate IKE messages or to determine if an attempt to exploit the issue is underway.

- NetFlow

NetFlow can be configured on the Internet Edge Router to determine if a VPN head-end device is

being subject to an IKE attack as shown in the example below:

```

edge-router#show ip cache flow
IP packet size distribution (17718080 total packets):
  1-32  64  96  128  160  192  224  256  288  320  352  384  416  448  480
    .927 .000 .000 .072 .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, 4456704 bytes
  1 active, 65535 inactive, 51 added
  8021 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 336520 bytes
  0 active, 16384 inactive, 0 added, 0 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
  last clearing of statistics never

Protocol      Total      Flows      Packets Bytes  Packets Active(Sec) Idle(Sec)
-----      Flows      /Sec      /Flow  /Pkt  /Sec      /Flow      /Flow
TCP-WWW        4         0.0         1      60     0.0        0.0       15.4
TCP-other       8         0.0         3      48     0.0         1.1       15.5
UDP-other     19       0.0       67338 112    6.5      140.6    15.3
ICMP           19         0.0      865037    28     84.0       236.8     13.8
Total:         50         0.0     354303    34     90.5       143.6     14.8

SrcIf      SrcIPAddress      DstIf      DstIPAddress      Pr SrcP DstP  Pkts
Fa1/0     192.0.2.200     Fa0/0     192.0.2.1       11 01F4 01F4 3003
Fa1/0      185.175.246.62   Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      0.230.148.122    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      85.209.13.76     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      80.128.151.37    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      248.109.89.75    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      157.9.168.72     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      31.176.6.0       Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      114.232.221.71   Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      175.7.125.23     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      170.134.34.58    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      18.3.184.14      Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      76.84.113.115    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      17.134.203.42    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      126.201.170.38   Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      229.63.43.57     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      193.147.96.69    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      204.5.89.41      Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      66.170.76.74     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      79.186.243.47    Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      81.150.19.62     Fa0/0      192.0.2.1         11 01F4 01F4    1
Fa1/0      81.79.175.95     Fa0/0      192.0.2.1         11 01F4 01F4    1
----- Output Truncated -----

```

In the example there is one active flow on UDP/500 (hex 01F4) from 192.0.2.200 destined to the IPSec router at IP address 192.0.2.1. The high packet count should serve as an indicator that this flow should be further investigated. Also the protocol statistics shown a high number of packets per flow for **UDP-other**. At the other extreme, a very high number of single packet flows on UDP/500 (hex 01F4) from multiple source IP addresses may be indicative of an attempt to exploit this issue or some other network problem that is interfering with IKE communications.

**Networks designed with multiple dedicated VPN termination points at different ingress locations should only be minimally impacted as long as VPN client devices are configured to connect to alternate VPN termination points if their primary VPN termination point is not available.**

- Call Admission Control for IKE

This feature is available beginning in Cisco IOS software versions 12.3(8)T and 12.2(33)SRA. The job of CAC is to protect the router from severe resource depletion and prevent crashes of the router. As the IKE protocol does not supply authenticated identity information about the peers until late in the IKE exchange, it is not possible for CAC to determine if a connection is from a trusted source until the authenticated identity is received. It is also not possible for CAC to determine if the request is a new session or a rekey event as the determination of this is not possible until the appropriate information is exchanged between peers after the session has been authenticated.

Call Admission Control for IKE (CAC) limits the number of simultaneous IKE security associations (SA) that a router can establish. This feature can limit IKE SA requests in two ways:

- Configure an absolute IKE SA limit. The router drops new IKE SA requests when the configured value is reached.
- Configure a system resource limit. The router drops new IKE SA requests when a percentage of system resources is being used. This value is based on the total consumed resources on the router, not just the resources used by IKE requests.

CAC is enabled in global configuration mode by entering one of the following commands:

- **crypto call admission limit ike sa number** This specifies the maximum number of IKE SAs that the router can establish before IKE begins rejecting new SA requests.
- **call admission limit percent** This instructs IKE to stop accepting new SA requests when the specified percentage of system resources is being used. At high IKE SA request level scanning, legitimate IPsec peers may not be able to connect to the router being scanned.

Additional information about the Call Admission Control for IKE feature is available at

[http://www.cisco.com/en/US/docs/ios/12\\_3t/12\\_3t8/feature/guide/gtcallik.html](http://www.cisco.com/en/US/docs/ios/12_3t/12_3t8/feature/guide/gtcallik.html).

In 12.4(6)T a further enhancement to CAC was added that allows the user to configure a limit on the number of in–negotiation (embryonic) IKE connections. This type of IKE connection represents either an aggressive mode IKE SA or a main mode IKE SA prior to its authentication and actual establishment.

An incoming IKE request is compared against the current count of in–negotiation sessions and the packet is dropped if the in–negotiation limit has been reached. As with other CAC settings, there is no way to determine if the incoming IKE request is a brand new connection or the rekey of an existing connection. This is because the only way to identify the event is to wait for the IKE SA to be authenticated and the presence or absence of the initial contact notify is determined.

The following command will allow the configured number of embryonic IKE SAs to start negotiation without contributing to the maximum number of IKE SAs allowed. This option is only available today in 12.4(6)T or later T train releases.

```
crypto call admission limit ike in-negotiation-sa
<10-99999> maximum in-negotiation IKE SA limit
```

- Interface Access–lists

IPsec head–end routers that do not support Call Admission Control for IKE can be configured with interface access–lists to drop packets that can be used to exploit this issue. Since the source IP address of peer VPN devices must be known beforehand, Interface access–lists are most effective in protecting site–to–site IPsec tunnels.

The following ACL is specifically designed to permit IKE (UDP/500) traffic only from known IPsec peer IP addresses and drop IKE from all other IP addresses. The ACL should be applied to Internet facing IPv4 interfaces and should include topology–specific filters. In this example, this VPN head–end router's IP address is 192.0.2.1 and IP addresses 192.0.2.100 and 192.0.2.101 are known site–to–site IPsec peers. All other IP addresses are denied.

```

access-list 150 permit udp host 192.0.2.100 host 192.0.2.1 eq 500
access-list 150 permit udp host 192.0.2.101 host 192.0.2.1 eq 500
access-list 150 deny udp any any eq 500
access-list 150 permit ip any any

interface serial 2/0
 ip access-group 150 in

```

These ACL statements may be deployed on the IOS VPN router as part of a transit access-list which will protect the router where the ACL is configured. Further information about transit ACLs is available in the white paper "Transit Access Control Lists: Filtering at Your Edge", available at <http://www.cisco.com/warp/public/707/tacl.html>.

- Control Plane Policing

IPSec head-end routers that do not support Call Admission Control for IKE can be configured with Control Plane Policing (CoPP) to drop or rate-limit packets that can be used to exploit this issue. To drop IKE packets, the source IP address of peer VPN devices must be known beforehand, and in this mode CoPP is most effective in protecting site-to-site IPSec tunnels. When CoPP is used to rate-limit IKE (UDP/500) traffic, it is not necessary to know the source IP addresses of peer VPN devices.

The following CoPP policy is specifically designed to permit IKE (UDP/500) traffic only from known IPSec peer IP addresses and drop IKE from all other IP addresses. In this example IP addresses 192.0.2.100 and 192.0.2.101 are known site-to-site IPSec peers. All other IP addresses are denied.

```

access-list 145 deny udp host 192.0.2.100 any eq 500
access-list 145 deny udp host 192.0.2.101 any eq 500
access-list 145 permit udp any any eq 500
access-list 145 deny ip any any

class-map allow-known-ike-class
 match access-group 145

policy-map allow-known-ike-policy
 class allow-known-ike-class
 drop

control-plane
 service-policy input allow-known-ike-policy

```

Please note that in the 12.0S, 12.2S, and 12.2SX Cisco IOS trains the policy-map syntax is different:

```

policy-map allow-known-ike-policy
 class allow-known-ike-class
 police 32000 1500 1500 conform-action drop exceed-action drop

```

The following CoPP policy is specifically designed to rate-limit IKE (UDP/500) traffic destined for this VPN head-end router:

```

access-list 146 permit udp any any eq 500
access-list 146 deny ip any any

class-map rate-limit-ike-class
 match access-group 146

policy-map rate-limit-ike-policy
 class rate-limit-ike-class
 police 32000 6000 12000 conform-action transmit exceed-action drop

control-plane
 service-policy input rate-limit-ike-policy

```

The police (rate-limit) policy is applied to traffic matching access-list 146. Traffic at rates below 32,000 bits per second (bps) will be forwarded normally but above 32,000 bps packets will be dropped. 32,000 bps is loosely equivalent to 35 IKE messages per second. Burst and extend burst values are calculated using recommended [burst values](#).

**Network conditions, traffic, and device load may lead to significantly different results in other environments, therefore this rate-limiting information is provided as an example and should serve as the basis for testing in your environment, not as recommended settings for production devices.**

Unlike interface access-lists, the CoPP policy will only be applied to traffic destined for the router itself. CoPP is available in Cisco IOS release trains 12.0S, 12.2SX, 12.2S, 12.3T, 12.4, and 12.4T. In the above CoPP example the access-list entries that match exploit packets with the **permit** action result in those packets being rate-limited by the policy-map **police** function or discarded by the policy-map **drop** function, while packets that match the **deny** actions are not affected by the policy-map.

Additional information on the configuration and use of the CoPP feature can be found at [http://www.cisco.com/en/US/docs/ios/12\\_3t/12\\_3t4/feature/guide/gtrtlmt.html](http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtrtlmt.html).

- Anti-Spoofing

This vulnerability can be exploited by spoofed packets. Anti-spoof protection in the form of interface access-lists or unicast Reverse Path Forwarding can provide limited mitigation if properly configured. This mitigation should not be relied upon to provide 100% mitigation since spoofed packets may still enter the network from the interface expected by uRPF or allowed by anti-spoofing access-lists. Also care must be taken to ensure that the appropriate uRPF mode (loose or strict) is configured to ensure that legitimate packets are not dropped.

Additional information about unicast Reverse Path Forwarding is available at [http://www.cisco.com/en/US/docs/ios/12\\_2t/12\\_2t13/feature/guide/ft\\_urpf.html](http://www.cisco.com/en/US/docs/ios/12_2t/12_2t13/feature/guide/ft_urpf.html).

**Note:** Customers who have implemented GDOI (Group Domain of Interpretation) for either Secure Multicast or Dynamic Multipoint VPN (DMVPN) can implement the previous mitigation techniques as needed. Any references to port 500/UDP or "isakmp" in an access-list would need to be replaced with port 848/UDP, the GDOI port. Additional information on GDOI can be found in the [Related Information](#) section of this Security Response.

### Identification

- High CPU

The first indication of an attempt to exploit this issue against an IOS device may be high CPU as shown below:

```
vpn-router#show process cpu
CPU utilization for five seconds: 99%/0%; one minute: 95%; five minutes: 51%
PID Runtime(ms)   Invoked    uSecs    5Sec    1Min    5Min  TTY Process
-----
----- Output Truncated -----
```

- Call Admission Control for IKE

After the router drops a packet due to exceeding a Call Admission Control for IKE limit, it will generate a warning level (severity 4) Syslog message. The Syslog message is %CRYPTO-4-IKE\_DENY\_SA\_REQ which has these formats:

– IKE Security Association Limit

**%CRYPTO-4-IKE\_DENY\_SA\_REQ : IKE denied an INCOMING SA request from [IP\_address] to [IP\_address] due to IKE SA LIMIT REACHED**

– System Resource Limit

**%CRYPTO-4-IKE\_DENY\_SA\_REQ : IKE denied an INCOMING SA request from [IP\_address] to [IP\_address] due to SYSTEM RESOURCES LOW**

Additional information on those syslog messages can be found at

[http://www.cisco.com/en/US/docs/ios/12\\_3t/system/messages/smg2tmsd.html#wp715560](http://www.cisco.com/en/US/docs/ios/12_3t/system/messages/smg2tmsd.html#wp715560).



In the above example 700 UDP/500 packets from untrusted IP addresses have been dropped by the access-list configured on interface Serial 2/0.

- Control Plane Policing

Once the CoPP configuration is deployed, the commands **show access-list 145** and **show policy-map control-plane all** can be used to identify packets dropped by the CoPP policy. Additionally SNMP queries using the Cisco QoS MIB **CISCO-CLASS-BASED-QOS-MIB** can be used to track packets being dropped by CoPP. Packets being dropped by CoPP should be investigated to determine if they are attempts to exploit the issue.

Example output for **show access-list 145**:

```
vpn-router#show access-list 145
Extended IP access list 145
 10 deny udp host 192.0.2.100 any eq isakmp (1000 matches)
 20 deny udp host 192.0.2.101 any eq isakmp (523 matches)
 30 permit udp any any eq isakmp (62931 matches)
 40 deny ip any any (15239 matches)
```

In the above example 62,931 IKE (UDP/500) packets from unknown IP addresses have been dropped by the access-list associated with CoPP.

Example output for **show policy-map control-plane all**:

```
vpn-router#show policy-map control-plane all
Control Plane

Service-policy input: allow-known-ike-policy

Class-map: allow-known-ike-class (match-all)
 62931 packets, 5694444 bytes
 5 minute offered rate 143000 bps, drop rate 143000 bps
Match: access-group 145
drop

Class-map: class-default (match-any)
 24 packets, 2497 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: any
```

In the above example class-map allow-known-ike-class has dropped 62,931 packets.

Example output for an SNMP query using the Cisco QoS MIB **CISCO-CLASS-BASED-QOS-MIB**.

```
Linux# tinsynmpget 10.89.236.153 ha5d7ogu355 1.3.6.1.4.1.9.9.166.1.15.1.1.13.1043.1045
1.3.6.1.4.1.9.9.166.1.15.1.1.13.1043.1045 = Counter32 62931
```

The output of this SNMP query matches the exceeded packets that were dropped shown in the CLI output above for class-map:

**allow-known-ike-class**

**1.3.6.1.4.1.9.9.166.1.15.1.1.13** is object **cbQosCMDropByte**.

**1043.1045** is the table index number for a specific control-plane service policy class-map. The table index can be determined by SNMP walking OID: 1.3.6.1.4.1.9.9.166.1. In this case, **1043.1045** indicates class map allow-known-ike-class.

- NetFlow

NetFlow can be used to determine if a router is being subject to an IKE attack as shown in the example below:

```
vpn-router#show ip cache flow
IP packet size distribution (17718080 total packets):
 1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
 .927 .000 .000 .072 .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, 4456704 bytes  
1 active, 65535 inactive, 51 added  
8021 aged polls, 0 flow alloc failures  
Active flows timeout in 30 minutes  
Inactive flows timeout in 15 seconds  
IP Sub Flow Cache, 336520 bytes  
0 active, 16384 inactive, 0 added, 0 added to flow  
0 alloc failures, 0 force free  
1 chunk, 1 chunk added  
last clearing of statistics never

Protocol	Total	Flows	Packets	Bytes	Packets	Active(Sec)	Idle(Sec)
-----	Flows	/Sec	/Flow	/Pkt	/Sec	/Flow	/Flow
TCP-WWW	4	0.0	1	60	0.0	0.0	15.4
TCP-other	8	0.0	3	48	0.0	1.1	15.5
<b>UDP-other</b>	<b>19</b>	<b>0.0</b>	<b>67338</b>	<b>112</b>	<b>6.5</b>	<b>140.6</b>	<b>15.3</b>
ICMP	19	0.0	865037	28	84.0	236.8	13.8
Total:	50	0.0	354303	34	90.5	143.6	14.8

SrcIf	SrcIPAddress	DstIf	DstIPAddress	Pr	SrcP	DstP	Pkts
<b>Fa1/0</b>	<b>192.0.2.200</b>	<b>Local</b>	<b>192.0.2.1</b>	<b>11</b>	<b>01F4</b>	<b>01F4</b>	<b>3003</b>
Fa1/0	185.175.246.62	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	0.230.148.122	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	85.209.13.76	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	80.128.151.37	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	248.109.89.75	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	157.9.168.72	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	31.176.6.0	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	114.232.221.71	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	175.7.125.23	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	170.134.34.58	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	18.3.184.14	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	76.84.113.115	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	17.134.203.42	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	126.201.170.38	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	229.63.43.57	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	193.147.96.69	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	204.5.89.41	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	66.170.76.74	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	79.186.243.47	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	81.150.19.62	Local	192.0.2.1	11	01F4	01F4	1
Fa1/0	81.79.175.95	Local	192.0.2.1	11	01F4	01F4	1
-----	Output Truncated	-----					

In the example, there is one active flow on UDP/500 (hex 01F4) from 192.0.2.200 destined to this IPsec router's 192.0.2.1 IP address. The high packet count should serve as an indicator that this flow should be further investigated. Also the protocol statistics shown a high number of packets per flow for **UDP-other**. At the other extreme, a very high number of single packet flows on UDP/500 (hex 01F4) from multiple source IP addresses may be indicative of an attempt to exploit this issue or some other network problem that is interfering with IKE communications.

## IOS-XR Routers

### Mitigation

GSR (c12000) and CRS-1 routers running IOS-XR software support software-based IPsec for locally sourced and terminated traffic only (used mostly for routing protocols). IKE and IPsec support was added on Release R2.0 for CRS-1 and on Release R3.2 for GSR. The IOS-XR software is designed to limit the impact of such an attack on system resources consumed by users already connected. IOS-XR has a hard-coded limit of 500 IKE SAs (which includes both established and embryonic IKE connections). If new IKE initiator

packets are received and the available IKE negotiation slots are full, the new request will be discarded. This prevents CPU processing power and memory resources from being fully consumed by incoming IKE requests. The finite IKE negotiation slot design prevents new tunnel requests from degrading existing tunnels and other functions performed by the router. The design goals of IOS-XR software are to allow the box to survive an attack without jeopardizing the tunnels already established.

Legitimate IKE initiator packets from valid users will have to compete for slots with an attacking stream of packets. An incoming attack stream of IKE initiator requests does not render the GSR/CRS-1 running IOS-XR incapable of connecting a valid user, it simply reduces the likelihood that an IKE negotiation slot will be available when the user request arrives.

While under this type of attack, the GSR/CRS-1 running IOS-XR software:

- Will not crash due to memory exhaustion
- Will not consume all available CPU cycles processing IKE requests
- Will not drop existing tunnels

Peer requests for IKE Phase-1 re-keys have to contend for IKE negotiation slots with attacking packets and will be discarded if no slots are available. However, if a tunnel re-key has not already been initiated by the peer, the GSR or CRS-1 running IOS-XR software will initiate one itself as the negotiated lifetime approaches. Re-keys initiated by the GSR/CRS-1 are not subject to the concurrent IKE negotiation limits since the IKE initiator request in this case is outbound. Given this, no re-keys of valid peers should be blocked by this type of attack.

- *Interface Access-lists*

Since in most cases the IP address of the IPsec peer is known, interface access-lists can be configured to drop packets that can be used to exploit this issue.

The following ACL is specifically designed to permit IKE (UDP/500) traffic only from known IPsec peer IP addresses and drop IKE from all other IP addresses. The ACL should be applied to Internet facing IPv4 interfaces and should include topology-specific filters. In this example, this GSR/CRS-1 router's IP address is 192.0.2.1 and IP addresses 192.0.2.100 and 192.0.2.101 are known IPsec peers. All other IP addresses are denied.

```
ipv4 access-list IKEfilter
  10 remark Only Allow IKE From Known Peers
  20 permit udp host 192.0.2.100 host 192.0.2.1 eq isakmp
  30 permit udp host 192.0.2.101 host 192.0.2.1 eq isakmp
  40 deny udp any any eq isakmp
  50 permit ipv4 any any
```

The following example shows how to apply the above IKEfilter ACL on packets inbound on the Packet-over-SONET (POS) interface 0/2/0/2:

```
RP/0/RP0/CPU0:router(config)#interface POS 0/2/0/2
RP/0/RP0/CPU0:router(config-if)#ipv4 access-group IKEfilter ingress <hardware-count>
```

**Note:** In order to be able to view ACL hit counts that occur in hardware, the keyword "hardware-count" must be added to the above command.

For additional information on configuring access-lists on Cisco IOS XR Software, please reference the following:

## Identification

- Interface Access-lists

Once the interface access-list is deployed, the command **show access-list ipv4 IKEfilter** can be used to identify the number of packets being dropped. Dropped packets should be investigated to determine if they are attempts to exploit the issue.

Example output for **show access-list ipv4 IKEfilter**:

```
RP/0/RP0/CPU0:CRS-A_IOX#show access-lists IKEfilter
ipv4 access-list IKEfilter
 10 remark Only Allow IKE From Known Peers
 20 permit udp host 192.0.2.100 host 192.0.2.1 eq isakmp
 30 permit udp host 192.0.2.101 host 192.0.2.1 eq isakmp
 40 deny udp any any eq isakmp (39 matches)
 50 permit ipv4 any any (753 matches)
RP/0/RP0/CPU0:CRS-A_IOX#
```

In the above example 39 UDP/500 packets from untrusted IP addresses have been dropped by the access-list IKEfilter .

## Cisco ASA 7.X Firewalls for VPN

### Mitigation

When valid IKE initiator packets are presented to a PIX or ASA running 7.X from an attacker, they will compete with IKE initiator packets from legitimate users attempting to connect. PIX and ASA running 7.X have a finite number of slots to use for incoming IKE negotiation as shown in the table below.

PIX/ASA 7.X Platform	Embryonic IKE SA Limit
PIX 515	400
PIX 525	600
PIX 535	1000
ASA 5510	50
ASA 5520	250
ASA 5540	1000
ASA 5550	1000

ASA software version 7.X is designed to limit the impact of such an attack on system resources consumed by users already connected. If new IKE initiator packets are received and the available IKE negotiation slots are full, the new request will be discarded. This prevents CPU processing power and memory resources from being fully consumed by incoming IKE requests. The finite IKE negotiation slot design prevents new tunnel requests from degrading existing tunnels and other functions performed by the firewall. The design goals of ASA software version 7.X are to allow the box to survive an attack without jeopardizing the tunnels already established.

Legitimate IKE initiator packets from valid users will have to compete for slots with an attacking stream of packets. An incoming attack stream of IKE initiator requests does not render the PIX or ASA running 7.X incapable of connecting a valid user, it simply reduces the likelihood that an IKE negotiation slot will be available when the user request arrives.

While under this type of attack, the PIX or ASA running 7.X:

- Will not crash due to memory exhaustion
- Will not consume all available CPU cycles processing IKE requests
- Will not drop existing tunnels

Peer requests for IKE Phase-1 re-keys have to contend for IKE negotiation slots with attacking packets and will be discarded if no slots are available. However, if a tunnel re-key has not already been initiated by the peer, the PIX or ASA running 7.X will initiate one itself as the negotiated lifetime approaches. Re-keys initiated by the ASA are not subject to the concurrent IKE negotiation limits since the IKE initiator request in this case is outbound. Given this, no re-keys of valid peers should be blocked by this type of attack.

SSL VPN (WebVPN) with the SVC Client is a functionally equivalent alternative to IPSec for Remote Access users. SSL has been shown to not be vulnerable to the same attack. For additional information on configuring SSL VPN refer to the following link:

[http://www.cisco.com/en/US/products/ps6120/products\\_configuration\\_guide\\_chapter09186a0080334071.html](http://www.cisco.com/en/US/products/ps6120/products_configuration_guide_chapter09186a0080334071.html)

### *Identification*

More IKE connection attempts than there are Embryonic IKE SA connection slots on the PIX or ASA will result in the generation of an informational level (severity 6) Syslog message. The Syslog message is **%PIX | ASA-6-713905** which has this format:

**%PIX-6-713905: IP = 192.0.2.145, Duplicate first packet detected. Ignoring packet.**

OR

**%ASA-6-713905: IP = 192.0.2.145, Duplicate first packet detected. Ignoring packet.**

Additional information on those syslog messages can be found at

[http://www.cisco.com/en/US/products/sw/secursw/ps2120/products\\_system\\_message\\_guide\\_chapter09186a008055fd2](http://www.cisco.com/en/US/products/sw/secursw/ps2120/products_system_message_guide_chapter09186a008055fd2)

The **IP =** shown is the source IP address of the IKE connection attempt. Multiple or continuous **%PIX | ASA-6-713905** messages should be investigated to determine if this issue is being exploited.

*show crypto isakmp stats* – There is a statistics value that is incremented every time an IKE negotiation fails because the peer (as Initiator) does not respond. The value is **Responder Fails** under the global IKE/isakmp stats. An abnormally high number of **Responder Fails** may indicate that the device has been attacked.

Here is an example output:

```
ASA(config)#show crypto isakmp stats
```

```
Global IKE Statistics
Active Tunnels: 1
Previous Tunnels: 13
In Octets: 31161
In Packets: 91
In Drop Packets: 0
In Notifys: 0
In P2 Exchanges: 13
In P2 Exchange Invalids: 0
In P2 Exchange Rejects: 0
In P2 Sa Delete Requests: 0
Out Octets: 16648
Out Packets: 120
```

```

Out Drop Packets: 0
Out Notifys: 62
Out P2 Exchanges: 0
Out P2 Exchange Invalids: 0
Out P2 Exchange Rejects: 0
Out P2 Sa Delete Requests: 12
Initiator Tunnels: 0
Initiator Fails: 0
Responder Fails: 8405
System Capacity Fails: 0
Auth Fails: 0
Decrypt Fails: 0
Hash Valid Fails: 0
No Sa Fails: 0

```

In the above example 8405 IKE negotiations failed because the responder did not reply.

## Cisco PIX 6.X Firewalls for VPN

### *Mitigation*

When valid IKE initiator packets are presented to a PIX from an attacker, they will compete with IKE initiator packets from legitimate users attempting to connect. PIX's have a finite number of slots to use for incoming IKE negotiation as shown in the table below.

PIX 6.X Platform	Embryonic IKE SA Limit
All PIX Firewall platforms running 6.X software	500

The PIX is designed to limit the impact of such an attack on system resources consumed by users already connected. If new IKE initiator packets are received and the available IKE negotiation slots are full, the new request will be discarded. This prevents CPU processing power and memory resources from being fully consumed by incoming IKE requests. The finite IKE negotiation slot design prevents new tunnel requests from degrading existing tunnels and other functions performed by the firewall. The design goals of the PIX are to allow the box to survive an attack without jeopardizing the tunnels already established.

Legitimate IKE initiator packets from valid users will have to compete for slots with an attacking stream of packets. An incoming attack stream of IKE initiator requests does not render the PIX incapable of connecting a valid administrator, it simply reduces the likelihood that an IKE negotiation slot will be available when the user request arrives.

While under this type of attack, the PIX:

- Will not crash due to memory exhaustion
- Will not consume all available CPU cycles processing IKE requests
- Will not drop existing tunnels

Peer requests for IKE Phase-1 re-keys have to contend for IKE negotiation slots with attacking packets and will be discarded if no slots are available. However, if a tunnel re-key has not already been initiated by the peer, the PIX will initiate one itself as the negotiated lifetime approaches. Re-keys initiated by the PIX are not subject to the concurrent IKE negotiation limits since the IKE initiator request in this case is outbound. Given this, no re-keys of valid peers should be blocked by this type of attack.

### *Identification*

More IKE connection attempts than there are Embryonic IKE SA connection slots on the PIX will result in the generation of an error level (severity 3) Syslog message. The Syslog message is **%PIX-3-ISAKMP: Exceeded embryonic limit.**

## Cisco 3000 Series VPN Concentrators

### Mitigation

When valid IKE initiator packets are presented to a VPN3000 Concentrator from an attacker, they will compete with IKE initiator packets from legitimate end-users attempting to connect. VPN3000 Concentrators have a finite number of slots to use for incoming IKE negotiation as shown in the table below.

VPN3000 Platform	Embryonic IKE SA Limit
VPN 3005	15
VPN 3015	15
VPN 3020	40
VPN 3030	40
VPN3060	40
VPN 3080	40

The VPN3000 system is designed to limit the impact of such an attack on system resources consumed by users already connected. If new IKE initiator packets are received and the available IKE negotiation slots are full, the new request will be discarded. This prevents CPU processing power and memory resources from being fully consumed by incoming IKE requests. The finite IKE negotiation slot design prevents new tunnel requests from degrading existing tunnels. The design goals of the VPN3000 are to allow the box to survive an attack without jeopardizing the tunnels already established and other functions performed by the device.

Legitimate IKE initiator packets from valid end users will have to compete for slots with an attacking stream of packets. An incoming attack stream of IKE initiator requests does not render the VPN3000 incapable of connecting a valid user, it simply reduces the likelihood that an IKE negotiation slot will be available when the user request arrives.

While under this type of attack, the VPN3000:

- Will not crash due to memory exhaustion
- Will not consume all available CPU cycles processing IKE requests
- Will not drop existing tunnels

Peer requests for IKE Phase-1 re-keys have to contend for IKE negotiation slots with attacking packets and will be discarded if no slots are available. However, if a tunnel re-key has not already been initiated by the peer, the VPN3000 will initiate one itself as the negotiated lifetime approaches. Re-keys initiated by the VPN3000 are not subject to the concurrent IKE negotiation limits since the IKE initiator request in this case is outbound. Given this, no re-keys of valid peers should be blocked by this type of attack.

### Identification

If an IKE initiator request is received by a VPN3000 Concentrator but there is no available slot to process it, an event will be generated to indicate this condition. The event generated is defined as:

- **Event ID:** IKE/191
- **Severity:** 4

- **Event Text:** Maximum concurrent IKE negotiations exceeded.
- **Explanation:** This condition indicates a high connection rate. The concentrator will limit the maximum number of concurrent IKE negotiations to prevent CPU lock-up during crypto key generation and user authentications.

An example of this event is:

**28304 07/10/2006 08:08:33.500 SEV=4 IKE/191 RPT=278 192.0.2.111 Maximum concurrent IKE negotiations exceeded.**

Note that the Repeat Count (RPT) and source IP Address are indicated in the event header.

The occurrence of this event does not always indicate a malicious attack. This can occur when valid user connection requests are received close together. This condition can also be induced if there are a number of hardware clients all attempting to re-establish connection if there was a network outage or after a maintenance reboot of the VPN3000 Concentrator. The software clients and hardware clients will retry their requests and ultimately connect as negotiation slots open up.

## Firewall Services Module

### *Mitigation*

When valid IKE initiator packets are presented to a Firewall Service Module (FWSM) from an attacker, they will compete with IKE initiator packets from legitimate administrators attempting to connect. FWSMs have a finite number of slots to use for incoming IKE negotiation as shown in the table below.

FWSM Software Version	Embryonic IKE SA Limit
FWSM 2.X Single Mode	5
FWSM 2.X Multi Mode	10
FWSM 3.X All Modes	50

The FWSM is designed to limit the impact of such an attack on system resources consumed by administrators already connected. If new IKE initiator packets are received and the available IKE negotiation slots are full, the new request will be discarded. This prevents CPU processing power and memory resources from being fully consumed by incoming IKE requests. The finite IKE negotiation slot design prevents new tunnel requests from degrading existing tunnels. The design goals of the FWSM are to allow the box to survive an attack without jeopardizing the tunnels already established.

Legitimate IKE initiator packets from valid administrators will have to compete for slots with an attacking stream of packets. An incoming attack stream of IKE initiator requests does not render the FWSM incapable of connecting a valid administrator, it simply reduces the likelihood that an IKE negotiation slot will be available when the administrator request arrives.

While under this type of attack, the FWSM:

- Will not crash due to memory exhaustion
- Will not consume all available CPU cycles processing IKE requests
- Will not drop existing tunnels

Peer requests for IKE Phase-1 re-keys have to contend for IKE negotiation slots with attacking packets and will be discarded if no slots are available. However, if a tunnel re-key has not already been initiated by the peer, the FWSM will initiate one itself as the negotiated lifetime approaches. Re-keys initiated by the FWSM

are not subject to the concurrent IKE negotiation limits since the IKE initiator request in this case is outbound. Given this, no re-keys of valid peers should be blocked by this type of attack.

### Identification

- FWSM 2.X

More than 5 IKE connections or connection attempts to the FWSM in single mode or to an individual context in multi mode, will result in the generation of a notification level (severity 5) Syslog message. The Syslog message is **%FWSM-5-321001**, which has this format:

**%FWSM-5-321001: Resource 'ipsec' limit of 5 reached for context 'single\_vf'**

Additional information on those syslog messages can be found at

[http://www.cisco.com/en/US/products/hw/switches/ps708/products\\_system\\_message\\_guide\\_chapter09186a00](http://www.cisco.com/en/US/products/hw/switches/ps708/products_system_message_guide_chapter09186a00)

The Resource should be 'ipsec' and context should be 'single\_vf' in single mode or a specific context in multimode. Multiple or continuous **%FWSM-5-321001** messages should be investigated to determine if this issue is being exploited.

- FWSM 3.X

More than 5 IKE connections or connection attempts to the FWSM in single mode or to an individual context in multi mode, will result in the generation of an informational level (severity 6) Syslog message. The Syslog message is **%FWSM-6-713905**, which has this format:

**%FWSM-6-713905: IP = 192.0.2.150, Duplicate first packet detected. Ignoring packet.**

Additional information on those syslog messages can be found at

[http://www.cisco.com/en/US/products/hw/switches/ps708/products\\_system\\_message\\_guide\\_chapter09186a00](http://www.cisco.com/en/US/products/hw/switches/ps708/products_system_message_guide_chapter09186a00)

The **IP =** shown is the source IP address of the IKE connection attempt. Multiple or continuous **%FWSM-6-713905** messages should be investigated to determine if this issue is being exploited.

## MDS

### Mitigation

- MDS IP Interface ACLs

MDS Gigabit Ethernet interfaces are not generic switch/router interfaces but are meant to terminate specific IP storage protocols (FC-IP/iSCSI). IKE and IPsec are only supported on specific crypto-enabled Gigabit Ethernet interfaces like those on the DS-X9302-14K9, 2-port 1-Gigabit Ethernet IPS, 14-port 1/2-Gbps Fibre Channel module. Since the source IP address of iSCSI clients and FCIP peers must be known beforehand, Interface access-lists are most effective in protecting site-to-site IPsec tunnels for storage.

For non-IPsec enabled internet facing Gigabit Ethernet IPS interfaces and management interface, the following ACL specifically designed to drop IKE from all IP addresses should be applied.

```
ip access-list ikeblock deny udp any any eq port 500
ip access-list ikeblock permit ip any any

interface mgmt 0
 ip access-group ikeblock in
```

For IPsec enabled interfaces, the following ACL is specifically designed to permit IKE (UDP/500) traffic only from known IPsec peer IP addresses and to drop IKE from all other IP addresses. The ACL should be applied to Internet facing GigabitEthernet IPS interfaces on the MPS-14/2 module and should include topology-specific filters. In this example, the protected IPS head-end's IP address is 192.0.2.1 and IP addresses 192.0.2.100 and 192.0.2.101 are known site-to-site IPsec peers. All other IP addresses are denied.

```
ip access-list ikeacl permit udp 192.0.2.100 0.0.0.0 192.0.2.1 0.0.0.0 eq port 500
ip access-list ikeacl permit udp 192.0.2.101 0.0.0.0 192.0.2.1 0.0.0.0 eq port 500
ip access-list ikeacl deny udp any any eq port 500
ip access-list ikeacl permit ip any any
```

```
interface GigabitEthernet1/1
 ip access-group ikeacl in
```

For additional information on configuring IKE/IPSec and IP ACLs on the MDS 9000 series products, please reference the following:

– Cisco MDS 9000 Family Configuration Guide, Release 2.x: Configuring IPSec Network Security

[http://www.cisco.com/en/US/docs/storage/san\\_switches/mds9000/sw/rel\\_2\\_x/san-os/configuration/guide/ipe.html](http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/rel_2_x/san-os/configuration/guide/ipe.html)

– Cisco MDS 9000 Family Configuration Guide, Release 2.x: Configuring IP Access Control Lists

[http://www.cisco.com/en/US/docs/storage/san\\_switches/mds9000/sw/rel\\_2\\_x/san-os/configuration/guide/ipac.html](http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/rel_2_x/san-os/configuration/guide/ipac.html)

## Identification

- Interface Access-lists

Once the interface access-list is deployed, the command **show ip access-list ikeacl** can be used to identify the number of packets being dropped. Dropped packets should be investigated to determine if they are attempts to exploit the issue.

Example output for **show ip access-list ikeacl**:

```
mds9600#show ip access-list ikeacl
ip access-list ikeacl permit udp 192.0.2.100 0.0.0.0 192.0.2.1 0.0.0.0 eq port 500 (0 matches)
ip access-list ikeacl permit udp 192.0.2.101 0.0.0.0 192.0.2.1 0.0.0.0 eq port 500 (0 matches)
ip access-list ikeacl deny udp any any eq port 500 (700 matches)
ip access-list ikeacl permit ip any any (26154 matches)
```

In the above example, 700 UDP/500 (IKE) packets from untrusted IP addresses have been dropped by the access-list ( ikeacl ) configured on interface gigabitEthernet 1/1.

- Syslog

A large number of syslog messages of the format **2006 Aug 8 07:01:06 192.0.1.10**

**%IKE-3-PHASE1\_NEGOTIATION\_FAILED: IKEv1: Phase 1 negotiation failed for peer 192.0.2.100** is an indication of half-open (embryonic) IKE connections which fail to establish a tunnel. These messages should be investigated to determine if this issue is being exploited.

## Revision History

Revision 2.3	2008-July-28	Added bug id CSCsb50996 to the list of bugs affecting Cisco PIX firewalls running pre-7.x code.
Revision 2.2	2006-September-01	Added information about GDOI to IOS mitigation sections.
Revision 2.1	2006-August-17	Added Cisco Bug ID for IOS-XR devices. Edited IOS-XR section to fix typos.
Revision 2.0	2006-August-08	Added detection and mitigation steps for affected devices.
Revision 1.3	2006-August-08	Added Cisco Bug ID for SAN-OS on MDS devices.
Revision 1.2	2006-August-01	Added Cisco Bug IDs for FWSM, ASA, and PIX devices

		running v7.x code.
Revision 1.1	2006-07-26	Updated Cisco Bug ID for VPN 3000 concentrators.
Revision 1.0	2006-07-26	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](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 IOS Secure Multicast – Questions and Answers](#)
- [Cisco IOS Secure Multicast – Whitepaper](#)
- [Implementing Group Domain of Interpretation in a Dynamic Multipoint VPN – Whitepaper](#)
- [Cisco Security Response: Internet Key Exchange Resource Exhaustion Attack](#)
- [NTA Monitor's Cisco VPN Concentrator IKE resource exhaustion Advisory](#)
- [Cisco MySDN: Internet Key Exchange Protocol Version 1 Denial of Service Vulnerability](#)
- [RFC 2408 – Internet Security Association and Key Management Protocol \(ISAKMP\)](#)
- [RFC 2409 – The Internet Key Exchange \(IKE\)](#)

---

[Contacts & Feedback](#) | [Help](#) | [Site Map](#) © 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Sep 01, 2006

Document ID: 70810

---