# BIG-IP® System: DoS Protection and Protocol Firewall Implementations

Version 13.1

# Table of Contents

# Detecting and Protecting Against DoS, DDoS, and Protocol Attacks

## About detecting and protecting against DoS, DDoS, and protocol attacks

Attackers can target the BIG-IP® system in a number of ways. The BIG-IP system addresses several possible DoS, DDoS, SIP, and DNS attack routes. These DoS attack prevention methods are available when theBIG-IP® Advanced Firewall Manager™ is licensed and provisioned.

### DoS and DDoS attacks

Denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks attempt to render a machine or network resource unavailable to users. DoS attacks involve the efforts of one or more sources to disrupt the services of one or more hosts connected to the Internet.

With Advanced Firewall Manager, you can configure the system to automatically track traffic and CPU usage patterns over time, and adapt automatically to possible DoS attacks across a range of DoS vectors. You can initiate DoS detection for the whole system, and in profiles that are associated with specific virtual servers. Configure responses to system-level DoS attack vectors in the DoS Device Configuration.

Automatic threshold configuration is available for a range of non-error packet types on the AFM system. Use automatic thresholds to adapt responses to DoS attack vectors based on the traffic history on the system.

With AFM, you can also configure manual responses to DoS vectors. For non-error packets, you can specify absolute packet-per-second limits for attack detection (reporting and logging), percentage increase thresholds for detection, and absolute rate limits on a wide variety of packets that attackers can leverage as attack vectors.

You can also enable Bad Actor detection on a per-vector basis to identify IP addresses that engage in attacks where one IP address is targeting many destinations; the system can automatically blacklist Bad Actor IP addresses with specific thresholds and time limits. In addition, you can use Attacked Destination Detection to determine IP addresses that are being attacked from many sources (many to one attacks). The attacked addresses are added to a list and packets are rate limited to that attacked address.

### DNS and SIP flood (or DoS) attacks

Denial-of-service (DoS) or flood attacks attempt to overwhelm a system by sending thousands of requests that are either malformed or simply attempt to overwhelm a system using a particular DNS query type or protocol extension, or a particular SIP request type. The BIG-IP system allows you to track such attacks, using the DoS Protection profile.

### DoS Sweep and Flood attacks

A sweep attack is a network scanning technique that sweeps your network by sending packets, and using the packet responses to determine responsive hosts. Sweep and flood attack prevention allows you to configure system thresholds for packets that conform to typical sweep or flood attack patterns. This configuration is set in the DoS Device Configuration.

### Malformed DNS packets

Malformed DNS packets can be used to consume processing power on the BIG-IP system, ultimately causing slowdowns like a DNS flood. The BIG-IP system drops malformed DNS packets, and allows you to configure how you track such attacks. This configuration is set in the DoS Protection profile.

### Malformed SIP packets

Malformed SIP request packets can be used to consume processing power on the BIG-IP system, ultimately causing slowdowns like a SIP flood. The BIG-IP system drops malformed SIP packets, and allows you to configure how you track such attacks. This configuration is set in the DoS Protection profile.

### Protocol exploits

Attackers can send DNS requests using unusual DNS query types or OpCodes. The BIG-IP system can be configured to allow or deny certain DNS query types, and to deny specific DNS OpCodes. When you configure the system to deny such protocol exploits, the system tracks these events as attacks. This configuration is set in the DNS Security profile.

## About profiles for DoS and protocol service attacks

On the BIG-IP® system, you can use different types of profiles to detect and protect against system DoS attacks, to rate limit possible attacks, and to automatically blacklist IP addresses when identified as Bad Actors. You can configure settings for specific protocol attacks for DNS and SIP, and other network attacks.

### DoS Protection profile

With the DoS Protection profile you can configure settings for DoS protection that you can apply to a virtual server, to protect a specific application or server. You can configure the DoS profile to provide specific attack prevention at a more granular level than the Device DoS profile. In a DoS Profile, you can:

- Configure automatic thresholds for each profile, and for specific DoS vectors, to allow the system to adjust the configuration for DoS attack detection automatically over time.
- Define a source IP address whitelist, to allow legitimate addresses to pass through the DoS protection checks.
- Define settings for DNS protocol error detection, which allows you to configure a percentage rate increase over time and a packets-per-second threshold to trigger logging, as well as a hard rate limit on DNS protocol error packets.
- Define packet-per-second detection-limit, percentage rate increases, and packet-per-second rate limiting for DNS record types.
- Define settings for SIP protocol error detection, which allows you to configure a percentage rate increase over time and a packets-per-second threshold to trigger logging, as well as a hard rate limit on SIP protocol error packets.
- Define specific packet-per-second rate increases, percentage rate increases, and packet-per-second rate limiting for SIP request methods.
- Configure identification, rate limiting, and automatic blacklisting of Bad Actors for supported attack vectors, according to various detection criteria.
- Offload blacklisting of Bad Actor IP addresses to edge routers using BGP.
- Configure identificaton, rate limiting, and classification of attacked destinations.

### DNS Protocol Security Profile

The DNS Security Profile is a separate profile that you specify in a DNS service profile, to provide security features. The DNS Security Profile allows you to configure the BIG-IP system to exclude (drop) or include (allow) packets of specific DNS query record types. You can also configure the profile to exclude (drop) the DNS QUERY header OpCode.

### HTTP Protocol Security Profile

The HTTP Security Profile allows you to configure the AFM system to perform HTTP protocol checks, HTTP request checks, and to present a blocking page if a check fails. You can attach an HTTP Security profile to a virtual server.

***Important:*** *You can attach an HTTP security profile only to a virtual server that is already configured with an HTTP profile.*

### SSH Proxy Protocol Security Profile

The SSH Proxy Security Profile allows you to configure the AFM system to allow or block SSH proxy commands, based on criteria including user name,

# Detecting and Preventing System DoS and DDoS Attacks

## About configuring the BIG-IP system to detect and prevent DoS and DDoS attacks

DoS and DDoS attack detection and prevention is enabled by the BIG-IP® Advanced Firewall Manager™ (AFM™) Device DoS Configuration for system-wide DoS protection, and by DoS Profiles for virtual servers. DoS detection features allow you to detect possible attacks on the system and on particular applications, and to rate limit possible attack vectors. AFM also enables further attack mitigation, including automatic identification and blacklisting of attacking IP addresses, and automatic configuration of DoS attack vector thresholds based on system analysis. DoS detection and prevention features are enabled with an Advanced Firewall Manager license, which also includes protocol DoS detection support that can be configured on a per-virtual-server basis.

- At the virtual server level, detect malicious or malformed DNS and SIP protocol errors, and report anomalies by percentage increase, or by absolute packets per second.
- At the virtual server level, rate limit malicious or malformed DNS and SIP protocol error packets.
- At the virtual server level and system-wide, manually configure detection of potential DoS vector attacks by rate increase or absolute packets per second, and rate limit or leak limit such packets.
- System-wide, automatically detect potential attacks across a wide range of DoS attack vectors, and rate limit or leak limit such packets,
- At the virtual server level, detect repeat attackers for SIP, DNS, and other attack vectors and automatically blacklist their IP addresses, with configurable thresholds and blacklist duration.
- System-wide, detect repeat attackers for a wide range of attack vectors and automatically blacklist their IP addresses, with configurable thresholds and blacklist duration.
- At the virtual server level and system-wide, advertise blacklisted IP addresses to BGP routers, per DoS vector and per IP intelligence category. With this option, once an IP address is identified for blacklisting, all further blacklisting of IP addresses is handled by upstream routers, until the blacklist entry is automatically removed.

### Task list

*Detecting and protecting against system-wide DoS and DDoS attacks*
*Automatically detecting and protecting against system-wide DoS and DDoS attacks*
*Configuring manual thresholds for DoS and DDoS vectors*

## Detecting and protecting against system-wide DoS and DDoS attacks

The BIG-IP® system handles DoS and DDoS attacks with preconfigured responses. With DoS Protection Device Configuration, you can automatically or manually set detection thresholds and internal rate or leak limits for a range of DoS and DDoS attack vectors.

***Note:*** *Not all settings apply to all DoS vectors. For example, some vectors cannot use automatic thresholds, and some vectors cannot be automatically blacklisted.*

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Properties**.
   The DoS Protection Device Configuration screen opens.
2. From the **Log Publisher** list, select the destination to which the BIG-IP system sends DoS and DDoS logs.

You can review, create, and update log publishers in **System** > **Logs** > **Configuration** > **Log Publishers**.

3. Configure the **Threshold Sensitivity**.

   Select **Low**, **Medium**, or **High**. A lower setting means the automatic threshold algorithm is less sensitive to changes in traffic and CPU usage, but will also trigger fewer false positives.
4. From the **Eviction Policy** list, select the eviction policy to apply globally.

   ***Note:** The global context requires an eviction policy. If you do not apply a custom eviction policy, the system default policy,* `default-eviction-policy` *is applied and selected in this field.*
5. For **Relearn**, click **Start Relearning** to start relearning auto thresholds.

   Auto thresholds are calculated from the system start. If you have made changes to the system since then, and want the system to adjust automatic DoS thresholds because of these changes, use this option.
6. To specify a system-wide DoS address list containing addresses that do not need to be checked for DoS attacks, type the name of the list in the **Whitelist Address List** field.

   ***Note:** Available address lists appear on the right side of the screen, in the Shared Objects pane. You can view, edit, and add address lists there.*
7. To apply a system-wide rich DoS whitelist, click **Add Whitelist**, and type the information to define the packets to allow.

   You can define up to eight rich whitelists.
8. At the top of the screen, from Device Configuration, choose Network Security, DNS Security, and SIP Security to configure relevant attack responses per vector.

   The screen displays all the available attack vectors for the given type.

   ***Note:** Network Security vectors are listed in categories to make the list more manageable. Click the* + *next to a category to expand it.*
9. To enable (or disable) auto thresholds for one or more attack types, select the check box next to the vector name or names, and from the **Set Threshold** button at the bottom of the screen, select **Fully-automatic**. Select **Manual** to disable auto thresholds and set properties manually.

   ***Note:** To work accurately, using fully-automatic thresholds requires some amount of historical data on the system gathered through observing normal traffic. Therefore, it is recommended that you not enforce auto thresholds directly after installation.*

   ***Tip:** You can select all vectors by clicking the check box at the top of the list. However, some vectors do not support automatic thresholds. Deselect these vectors before you select **Fully-automatic** to avoid an error.*
10. Similarly, to set the state for one or more attack types, select the check box next to the vector name or names, and from the **Set State** list at the bottom of the screen, select **Mitigate**, **Detect Only**, or **Disable**.
    The state you click is set for all selected vectors.
11. In the **Attack Type** column, click the name of any attack type to edit the settings.
    The attack settings appear on the right, in the **Properties** pane.
12. To enforce the DoS vector, make sure the **State** is set to **Mitigate** (watch, learn, alert, and mitigate) .

    Other options allow you to **Detect Only** (watch, learn, and alert) or **Learn Only** (collect stats, no mitigation),

***Caution:** For most DoS vectors, you want to enforce the vector, which is the default setting. Set a vector to **Disabled** (no stat collection, no mitigation) only when you find that enforcement of the vector is disrupting legitimate traffic. For example, if you test a legitimate packet with the packet tester and find a DoS vector is preventing packet transmission, you can adjust the thresholds or disable the vector to remedy the issue.*

13. Set the **Threshold Mode** for the vector.
    - If the attack allows automatic threshold configuration, you can select **Fully Automatic** or **Manual Detection/Auto Mitigation** to configure automatic or partially automatic thresholds.
    - To configure thresholds manually, click **Fully Manual**.
14. Adjust the other settings for the DoS vector for fully automatic, partially maual, or fully manual threshold configuration.
15. Click the **Update** button.
    The selected configuration is updated, and the changes appear on the Device Configuration screen.
16. Repeat the previous steps for any other attack types for which you want to change the configuration.

You have now configured the system to provide custom responses to possible DoS and DDoS attacks, and to allow such attacks to be identified in system logs and reports.

Next, you can configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system.

## Automatically detecting and protecting against system-wide DoS and DDoS attacks

The BIG-IP® system handles DoS and DDoS attacks with preconfigured responses. With the DoS Protection Device Configuration, you can automatically or manually set detection thresholds and internal rate or leak limits for a range of DoS and DDoS attack vectors. Use this task to configure automatic thresholds for the system, and for adjusting individual DoS vectors.

***Note:** Not all settings apply to all DoS vectors. For example, some vectors do not support automatic thresholds, and some vectors do not include bad actor detection or automatic blacklisting.*

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Properties**.
   The DoS Protection Device Configuration screen opens.
2. From the **Log Publisher** list, select the destination to which the BIG-IP system sends DoS and DDoS logs.

   You can review, create, and update log publishers in **System** > **Logs** > **Configuration** > **Log Publishers**.
3. Configure the **Threshold Sensitivity**.

   Select **Low**, **Medium**, or **High**. A lower setting means the automatic threshold algorithm is less sensitive to changes in traffic and CPU usage, but will also trigger fewer false positives.
4. From the **Eviction Policy** list, select the eviction policy to apply globally.

   ***Note:** The global context requires an eviction policy. If you do not apply a custom eviction policy, the system default policy,* `default-eviction-policy` *is applied and selected in this field.*
5. For **Relearn**, click **Start Relearning** to start relearning auto thresholds.

   Auto thresholds are calculated from the system start. If you have made changes to the system since then, and want the system to adjust automatic DoS thresholds because of these changes, use this option.
6. To specify a system-wide DoS address list containing addresses that do not need to be checked for DoS attacks, type the name of the list in the **Whitelist Address List** field.

***Note:** Available address lists appear on the right side of the screen, in the Shared Objects pane. You can view, edit, and add address lists there.*

7. To apply a system-wide rich DoS whitelist, click **Add Whitelist**, and type the information to define the packets to allow.

   You can define up to eight rich whitelists.
8. At the top of the screen, from Device Configuration, choose Network Security, DNS Security, and SIP Security to configure relevant attack responses per vector.

   The screen displays all the available attack vectors for the given type.

***Note:** Network Security vectors are listed in categories to make the list more manageable. Click the* + *next to a category to expand it.*

9. To enable auto thresholds for one or more attack types, select the check box next to the vector name or names, and from the **Set Threshold** button at the bottom of the screen, select **Fully-automatic**.

***Note:** To work accurately, using fully-automatic thresholds requires some amount of historical data on the system gathered through observing normal traffic. Therefore, it is recommended that you not enforce auto thresholds directly after installation.*

***Tip:** You can select all vectors by clicking the check box at the top of the list. However, some vectors do not support automatic thresholds. Deselect these vectors before you select **Fully-automatic** to avoid an error.*

10. In the **Attack Type** column, click the name of any attack type to edit the settings.
    The attack settings appear on the right, in the **Properties** pane.
11. To enforce the DoS vector, make sure the **State** is set to **Mitigate** (watch, learn, alert, and mitigate) .

    Other options allow you to **Detect Only** (watch, learn, and alert) or **Learn Only** (collect stats, no mitigation),

***Caution:** For most DoS vectors, you want to enforce the vector, which is the default setting. Set a vector to **Disabled** (no stat collection, no mitigation) only when you find that enforcement of the vector is disrupting legitimate traffic. For example, if you test a legitimate packet with the packet tester and find a DoS vector is preventing packet transmission, you can adjust the thresholds or disable the vector to remedy the issue.*

12. For **Threshold Mode**, select **Fully Automatic**.

***Note:** You cannot configure automatic thresholds for every DoS vector. In particular, for error packets you can manually specify only **Detection Threshold EPS**, **Detection Threshold Percent**, and the **Mitigation Threshold EPS**.*

***Note:** If automatic thresholds are available, you can configure automatic thresholds, partially manual, or manual thresholds for that DoS vector. When you select one configuration setting, the options for the other setting no longer appear.*

13. In the **Attack Floor EPS** field, specify the number of events per second of the vector type to allow at a minimum, before automatically calculated thresholds are determined.

    Because automatic thresholds take time to be reliably established, this setting defines the minimum packets allowed before automatic thresholds are calculated.
14. In the **Attack Ceiling EPS** field, specify the absolute maximum allowable for events of this type, before automatically calculated thresholds are determined.

    Because automatic thresholds take time to be reliably established, this setting rate limits packets to the events per second setting, when specified. To set no hard limit, set this to **Infinite**.

15. If the vector includes other settings, such as Bad Actor Detection and Attacked Destination Detection, configure them as needed. If using automatic blacklisting with Bad Actor Detection, be sure to assign a global IP intelligence policy to the device (**Security** > **Network Firewall** > **IP Intelligence** > **Policies**).
16. Click the **Update** button.
    The selected vector is updated, and the DoS Protection Device Configuration screen refreshes.
17. Repeat the previous steps for any other attack types for which you want to change the configuration.

Now you have configured the system to automatically detect and respond to possible DoS and DDoS attacks, and to identify such attacks in system logs and reports.

Configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system.

## Configuring manual thresholds for DoS and DDoS vectors

You manually configure thresholds for a DoS vector when you want to configure specific settings, or when the vector does not allow automatic threshold configuration.

***Note:*** *Not all settings apply to all DoS vectors. For example, some vectors cannot be automatically blacklisted.*

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Properties**.
   The DoS Protection: Device Configuration Properties screen opens.
2. At the top of the screen, from Device Configuration, choose Network Security, DNS Security, and SIP Security to configure relevant attack responses per vector.

   The screen displays all the available attack vectors for the given type.

   ***Note:*** *Network Security vectors are listed in categories to make the list more manageable. Click the* **+** *next to a category to expand it.*
3. In the **Attack Type** column, click the name of any attack type to edit the settings.
   The attack settings appear on the right, in the **Properties** pane.
4. For **Threshold Mode**, select **Fully Manual**.
5. From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
6. From the **Detection Threshold Percent** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in percentage of traffic) for the attack detection threshold. If packets of the specified types cross the percentage threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
7. From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.

8. To log traffic that the system identifies as a DoS attack according to the automatic thresholds, enable **Simulate Auto Threshold.**

   ***Note:*** *This setting applies only to vectors that can be configured for automatic thresholds. It allows you to see the results of automatic thresholds on the selected DoS vector without actually affecting traffic. When you enable this setting, the current system-computed thresholds for automatic thresholds are displayed for this vector. Automatic thresholds are not applied to packets unless the* ***Threshold Mode*** *is changed for the vector.*

9. To detect IP address sources from which possible attacks originate, enable **Bad Actor Detection**.

   ***Note:*** *Bad Actor Detection is not available for every vector.*

10. In the **Per Source IP Detection Threshold EPS** field, specify the number of events of this type per second from one IP address that identifies the IP source as a bad actor, for purposes of attack detection and logging.
11. In the **Per Source IP Mitigation Threshold EPS** field, specify the number of events of this type per second from one IP address, above which rate limiting or leak limiting occurs.
12. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

    ***Important:*** *For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy:* ***Security*** > ***Network Firewall*** > ***IP Intelligence*** > ***Policies****. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

13. Specify the **Sustained Attack Detection Time**, in seconds, after which an IP address is blacklisted.
14. To change the duration for which the address is blacklisted, specify the duration in seconds in the **Category Duration Time** field. The default duration for an automatically blacklisted item is 4 hours (`14400` seconds).

    After this time period, the IP address is removed from the blacklist.
15. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** > ***Options*** > ***External Redirection*** > ***Blacklist Publisher****.*

16. To set thresholds for attacked destinations, select **Attacked Destination Detection**.
    a) In the **Per Destination IP Detection Threshold EPS** field, specify the number of events per second that IP source as a bad actor, for purposes of attack detection and logging.
    b) In the **Per Destination IP Mitigation Threshold EPS** field, specify the number of events per second headed to one IP address, above which rate limiting occurs.
    c) To automatically blacklist bad actor IP addresses, select **Add Destination Address to Category**.

       For DoS protection, the blacklist category is set to **denial_of_service** automatically.
    d) Specify the **Sustained Attack Detection Time**, in seconds, after which an IP address is blacklisted.
    e) To set the duration the destination address remains blacklisted, specify the **Category Duration Time** in seconds. The default is `900` seconds.
    f) To allow destination IP blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.
17. Click the **Update** button.
    The selected configuration is updated, and the DoS Protection Device Configuration screen opens again.
18. Repeat the previous steps for any other attack types for which you want to manually configure thresholds.

Now you have configured the system to provide custom responses to possible DoS and DDoS attacks, and to allow such attacks to be identified in system logs and reports, rate-limited, and blacklisted when specified.

Configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system. Configure a Blacklist Publisher, if necessary, to advertise routes for blacklist entries.

## Device DoS attack types

You can specify particular auto or manual thresholds, rate increases, rate limits, enforcement, and other parameters for supported device DoS attack types, to more accurately detect, track, and rate limit attacks.

***Important:*** *All hardware-supported vectors are performed in hardware on vCMP® guests, provided that the vCMP guests have the same software version as the vCMP host.*

### Network Security vectors

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| Flood | Ethernet Broadcast Packet | ether-brdcst-pkt | Ethernet broadcast packet flood | Yes |
| Flood | Ethernet Multicast Packet | ether-multicst-pkt | Ethernet destination is not broadcast, but is multicast | Yes |
| Flood | ARP Flood | arp-flood | ARP packet flood | Yes |
| Flood | IP Fragment Flood | ip-frag-flood | Fragmented packet flood with IPv4 | Yes |
| Flood | IGMP Flood | igmp-flood | Flood with IGMP packets (IPv4 packets with IP protocol number 2) | Yes |
| Flood | Routing Header Type 0 | routing-header-type-0 | Routing header type zero is present in flood packets | Yes |
| Flood | IPv6 Fragment Flood | ipv6-frag-flood | Fragmented packet flood with IPv6 | No |
| Flood | IGMP Fragment Flood | igmp-frag-flood | Fragmented packet flood with IGMP protocol | Yes |
| Flood | TCP SYN Flood | tcp-syn-flood | TCP SYN flood | Yes |
| Flood | TCP SYN ACK Flood | tcp-synack-flood | TCP SYN/ACK flood | Yes |
| Flood | TCP RST Flood | tcp-rst-flood | TCP RST flood | Yes |
| Flood | TCP Window Size | tcp-window-size | The TCP window size in packets is above the maximum. To tune this value, in `tmsh`: `modify sys db dos.tcplowwindowsize value`, where `value` is <=`128`. | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| Flood | ICMPv4 Flood | icmpv4-flood | Flood with ICMP v4 packets | Yes |
| Flood | ICMPv6 Flood | icmpv6-flood | Flood with ICMP v6 packets | Yes |
| Flood | UDP Flood | udp-flood | UDP flood attack | Yes |
| Flood | TCP SYN Oversize | tcp-syn-oversize | Detects TCP data SYN packets larger than the maximum specified by the dos.maxsynsize parameter. To tune this value, in `tmsh`: `modify sys db dos.maxsynsize value`. The default size is `64` and the maximum allowable value is `9216`. | Yes |
| Flood | TCP Push Flood | tcp-push-flood | TCP push packet flood | Yes |
| Flood | TCP BADACK Flood | tcp-ack-flood | TCP ACK packet flood | No |
| Bad Header - L2 | Ethernet MAC Source Address == Destination Address | ether-mac-sa-eq-da | Ethernet MAC source address equals the destination address | Yes |
| Bad Header - IPv4 | Bad IP Version | bad-ver | The IPv4 address version in the IP header is not 4 | Yes |
| Bad Header - IPv4 | Header Length Too Short | hdr-len-too-short | IPv4 header length is less than 20 bytes | Yes |
| Bad Header - IPv4 | Header Length > L2 Length | hdr-len-gt-l2-len | No room in layer 2 packet for IP header (including options) for IPv4 address | Yes |
| Bad Header - IPv4 | L2 Length >> IP Length | l2-len-ggt-ip-len | Layer 2 packet length is much greater than the payload length in an IPv4 address header and the layer 2 length is greater than the minimum packet size | Yes |
| Bad Header - IPv4 | No L4 | no-l4 | No layer 4 payload for IPv4 address | Yes |
| Bad Header - IPv4 | Bad IP TTL Value | bad-ttl-val | Time-to-live equals zero for an IPv4 address | Yes |
| Bad Header - IPv4 | TTL <= <tunable> | ttl-leq-one | An IP packet with a destination that is not multicast and that has a TTL greater than 0 and less than or equal to a tunable value, which is 1 by default. To tune this value, in `tmsh`: `modify sys db dos.iplowttl value`, where `value` is `1-4`. | Yes |
| Bad Header - IPv4 | IP Error Checksum | ip-err-chksum | The header checksum is not correct | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| Bad Header - IPv4 | IP Option Frames | ip-opt-frames | IPv4 address packet with `option.db variable tm.acceptipsourceroute` must be enabled to receive IP options. | Yes |
| Bad Header - IPv4 | Bad Source | ip-bad-src | The IPv4 source IP = `255.255.255.255` or `0xe0000000U` | Yes |
| Bad Header - IPv4 | IP Option Illegal Length | bad-ip-opt | Option present with illegal length | No |
| Bad Header - IPv4 | Unknown Option Type | unk-ipopt-type | Unknown IP option type | No |
| Bad Header - IGMP | Bad IGMP Frame | bad-igmp-frame | IPv4 IGMP packets should have a header >= 8 bytes. Bits 7:0 should be either 0x11, 0x12, 0x16, 0x22 or 0x17, or else the header is bad. Bits 15:8 should be non-zero only if bits 7:0 are 0x11, or else the header is bad. | Yes |
| Fragmentation | IP Fragment Too Small | ip-short-frag | IPv4 short fragment error | Yes |
| Fragmentation | IPv6 Fragment Too Small | ipv6-short-frag | IPv6 short fragment error | Yes |
| Fragmentation | IPV6 Atomic Fragment | ipv6-atomic-frag | IPv6 Frag header present with M=0 and FragOffset =0 | Yes |
| Fragmentation | ICMP Fragment | icmp-frag | ICMP fragment flood | Yes |
| Fragmentation | IP Fragment Error | ip-other-frag | Other IPv4 fragment error | Yes |
| Fragmentation | IPV6 Fragment Error | ipv6-other-frag | Other IPv6 fragment error | Yes |
| Fragmentation | IP Fragment Overlap | ip-overlap-frag | IPv4 overlapping fragment error | No |
| Fragmentation | IPv6 Fragment Overlap | ipv6-overlap-frag | IPv6 overlapping fragment error | No |
| Bad Header - IPv6 | Bad IPV6 Version | bad-ipv6-ver | The IPv6 address version in the IP header is not 6 | Yes |
| Bad Header - IPv6 | IPV6 Length > L2 Length | ipv6-len-gt-l2-len | IPv6 address length is greater than the layer 2 length | Yes |
| Bad Header - IPv6 | Payload Length < L2 Length | payload-len-ls-l2-len | Specified IPv6 payload length is less than the L2 packet length | Yes |
| Bad Header - IPv6 | Too Many Extension Headers | too-many-ext-hdrs | For an IPv6 address, there are more than <tunable> extended headers (the default is `4`). To tune this value, in `tmsh`: `modify sys db dos.maxipv6exthdrs value`, where `value` is `0-15`. | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| Bad Header - IPv6 | IPv6 duplicate extension headers | dup-ext-hdr | An extension header should occur only once in an IPv6 packet, except for the Destination Options extension header | Yes |
| Bad Header - IPv6 | IPv6 extension header too large | ext-hdr-too-large | An extension header is too large. To tune this value, in `tmsh`: `modify sys db dos.maxipv6extsize value`, where `value` is `0-1024`. | Yes |
| Bad Header - IPv6 | No L4 (Extended Headers Go To Or Past End of Frame) | l4-ext-hdrs-go-end | Extended headers go to the end or past the end of the L4 frame | Yes |
| Bad Header - IPv6 | Bad IPV6 Hop Count | bad-ipv6-hop-cnt | Both the terminated (cnt=0) and forwarding packet (cnt=1) counts are bad | Yes |
| Bad Header - IPv6 | IPv6 hop count <= <tunable> | hop-cnt-leq-one | The IPv6 extended header hop count is less than or equal to <tunable>. To tune this value, in `tmsh`: `modify sys db dos.ipv6lowhopcnt value`, where `value` is `1-4`. | Yes |
| Bad Header - IPv6 | IPv6 Extended Header Frames | ipv6-ext-hdr-frames | IPv6 address contains extended header frames | Yes |
| Bad Header - IPv6 | IPv6 extended headers wrong order | bad-ext-hdr-order | Extension headers in the IPv6 header are in the wrong order | Yes |
| Bad Header - IPv6 | Bad IPv6 Addr | ipv6-bad-src | IPv6 source IP = `0xff00::` | Yes |
| Bad Header - IPv6 | IPv4 Mapped IPv6 | ipv4-mapped-ipv6 | IPv4 address is in the lowest 32 bits of an IPv6 address. | Yes |
| Bad Header - TCP | TCP Header Length Too Short (Length < 5) | tcp-hdr-len-too-short | The Data Offset value in the TCP header is less than five 32-bit words | Yes |
| Bad Header - TCP | TCP Header Length > L2 Length | tcp-hdr-len-gt-l2-len | | Yes |
| Bad Header - TCP | Unknown TCP Option Type | unk-tcp-opt-type | Unknown TCP option type | Yes |
| Bad Header - TCP | Option Present With Illegal Length | opt-present-with-illegal-len | Option present with illegal length | Yes |
| Bad Header - TCP | TCP Option Overruns TCP Header | tcp-opt-overruns-tcp-hdr | The TCP option bits overrun the TCP header | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| Bad Header - TCP | Bad TCP Checksum | bad-tcp-chksum | The TCP checksum does not match | Yes |
| Bad Header - TCP | Bad TCP Flags (All Flags Set) | bad-tcp-flags-all-set | Bad TCP flags (all flags set) | Yes |
| Bad Header - TCP | Bad TCP Flags (All Cleared) | bad-tcp-flags-all-clr | Bad TCP flags (all cleared and SEQ#=0) | Yes |
| Bad Header - TCP | SYN && FIN Set | syn-and-fin-set | Bad TCP flags (SYN and FIN set) | Yes |
| Bad Header - TCP | FIN Only Set | fin-only-set | Bad TCP flags (only FIN is set) | Yes |
| Bad Header - TCP | TCP Flags - Bad URG | tcp-bad-urg | Packet contains a bad URG flag, this is likely malicious | Yes |
| Bad Header - ICMP | Bad ICMP Checksum | bad-icmp-chksum | An ICMP frame checksum is bad. Reuse the TCP or UDP checksum bits in the packet | Yes |
| Bad Header - ICMP | Bad ICMP Frame | bad-icmp-frame | The ICMP frame is either the wrong size, or not of one of the valid IPv4 or IPv6 types. Valid IPv4 types:<br>• 0 Echo Reply<br>• 3 Destination Unreachable<br>• 4 Source Quench<br>• 5 Redirect<br>• 8 Echo<br>• 11 Time Exceeded<br>• 12 Parameter Problem<br>• 13 Timestamp<br>• 14 Timestamp Reply<br>• 15 Information Request<br>• 16 Information Reply<br>• 17 Address Mask Request<br>• 18 Address Mask Reply<br>Valid IPv6 types:<br>• 1 Destination Unreachable<br>• 2 Packet Too Big<br>• 3 Time Exceeded<br>• 4 Parameter Problem<br>• 128 Echo Request<br>• 129 Echo Reply<br>• 130 Membership Query<br>• 131 Membership Report<br>• 132 Membership Reduction | Yes |
| Bad Header - ICMP | ICMP Frame Too Large | icmp-frame-too-large | The ICMP frame exceeds the declared IP data length or the maximum datagram length. To tune this value, in `tmsh`: | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| | | | `modify sys db dos.maxicmpframesize value`, where `value` is <=`65515`. | |
| Bad Header - UDP | Bad UDP Header (UDP Length > IP Length or L2 Length) | bad-udp-hdr | UDP length is greater than IP length or layer 2 length | Yes |
| Bad Header - UDP | Bad UDP Checksum | bad-udp-chksum | The UDP checksum is not correct | Yes |
| Other | Host Unreachable | host-unreachable | Host unreachable error | Yes |
| Other | TIDCMP | tidcmp | ICMP source quench attack | Yes |
| Other | LAND Attack | land-attack | Source IP equals destination IP address | Yes |
| Other | IP Unknown protocol | ip-unk-prot | Unknown IP protocol | No |
| Other | TCP Half Open | tcp-half-open | The number of new or untrusted TCP connections that can be established. Overrides the Global SYN Check threshold in Configuration > Local Traffic > General. | No |
| Other | IP uncommon proto | ip-uncommon-proto | Sets thresholds for and tracks packets containing IP protocols considered to be uncommon. By default, all IP protocols other than TCP, UDP, ICMP, IPV6-ICMP, and SCTP are on the IP uncommon protocol list. | Yes |
| Bad Header - DNS | DNS Oversize | dns-oversize | Detects oversized DNS headers. To tune this value, in `tmsh`: `modify sys db dos.maxdnssize value`, where `value` is `256-8192`. | Yes |
| Single Endpoint | Single Endpoint Sweep | sweep | Sweep on a single endpoint. You can configure packet types to check for, and packets per second for both detection and rate limiting. | No |
| Single Endpoint | Single Endpoint Flood | flood | Flood to a single endpoint. You can configure packet types to check for, and packets per second for both detection and rate limiting. | No |
| Bad Header-SCTP | Bad SCTP Checksum | bad-sctp-checksum | Bad SCTP packet checksum | No |

### DNS Security vectors

The system tracks and rate limits all UDP DNS packets (excluding those whitelisted). TCP DNS packets are also tracked but only for the DNS requests that reach a virtual server that has a DNS profile associated with it.

For vectors where VLAN is <tunable>, you can tune this value in `tmsh`: `modify sys db dos.dnsvlan value`, where `value` is `0-4094`.

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| DNS | DNS A Query | dns-a-query | DNS Query, DNS Qtype is A_QRY, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS AAAA Query | dns-aaaa-query | DNS Query, DNS Qtype is AAAA, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS Any Query | dns-any-query | DNS Query, DNS Qtype is ANY_QRY, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS AXFR Query | dns-axfr-query | DNS Query, DNS Qtype is AXFR, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS CNAME Query | dns-cname-query | DNS Query, DNS Qtype is CNAME, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS IXFR Query | dns-ixfr-query | DNS Query, DNS Qtype is IXFR, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS Malformed | dns-malformed | Malformed DNS packet | Yes |
| DNS | DNS MX Query | dns-mx-query | DNS Query, DNS Qtype is MX, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS NS Query | dns-ns-query | DNS Query, DNS Qtype is NS, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS OTHER Query | dns-other-query | DNS Query, DNS Qtype is OTHER, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS PTR Query | dns-ptr-query | DNS Query, DNS Qtype is PTR, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS Question Items != 1 | dns-qdcount-limit | DNS Query, DNS Qtype is ANY_QRY, the DNS query has more than one question. | Yes |
| DNS | DNS Response Flood | dns-response-flood | UDP DNS Port=53, packet and DNS header flags bit 15 is 1 (response), VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS SOA Query | dns-soa-query | DNS Query, DNS Qtype is SOA_QRY, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |
| DNS | DNS SRV Query | dns-srv-query | DNS Query, DNS Qtype is SRV, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| DNS | DNS TXT Query | dns-txt-query | DNS Query, DNS Qtype is TXT, VLAN is <tunable> in tmsh using `dos.dnsvlan`. | Yes |

### SIP Security vectors

| DoS category | Attack name | Dos vector name | Information | Hardware accelerated |
|---|---|---|---|---|
| SIP | SIP ACK Method | sip-ack-method | SIP ACK packets | Yes |
| SIP | SIP BYE Method | sip-bye-method | SIP BYE packets | Yes |
| SIP | SIP CANCEL Method | sip-cancel-method | SIP CANCEL packets | Yes |
| SIP | SIP INVITE Method | sip-invite-method | SIP INVITE packets | Yes |
| SIP | SIP Malformed | sip-malformed | Malformed SIP packets | Yes |
| SIP | SIP MESSAGE Method | sip-message-method | SIP MESSAGE packets | Yes |
| SIP | SIP NOTIFY Method | sip-notify-method | SIP NOTIFY packets | Yes |
| SIP | SIP OPTIONS Method | sip-options-method | SIP OPTIONS packets | Yes |
| SIP | SIP OTHER Method | sip-other-method | Other SIP method packets | Yes |
| SIP | SIP PRACK Method | sip-prack-method | SIP PRACK packets | Yes |
| SIP | SIP PUBLISH Method | sip-publish-method | SIP PUBLISH packets | Yes |
| SIP | SIP REGISTER Method | sip-register-method | SIP REGISTER packets | Yes |
| SIP | SIP SUBSCRIBE Method | sip-subscribe-method | SIP SUBSCRIBE packets | Yes |
| SIP | SIP URI Limit | sip-uri-limit | Packets that exceed the SIP URI limit | Yes |

# Preventing Global DoS Sweep and Flood Attacks

## About DoS sweep and flood attack prevention

A *sweep attack* is a network scanning technique that typically sweeps your network by sending packets, and using the packet responses to determine live hosts. Typical attacks use ICMP to accomplish this.

The sweep vector tracks packets by source address. Packets from a specific source that meet the defined single endpoint sweep criteria, and exceed the rate limit, are dropped. You can also configure the sweep vector to automatically blacklist an IP address from which the sweep attack originates.

***Important:*** *The sweep mechanism protects against a flood attack from a single source, whether that attack is to a single destination host, or multiple hosts.*

A *flood attack* is a an attack technique that floods your network with packets of a certain type, in an attempt to overwhelm the system. A typical attack might flood the system with SYN packets without then sending corresponding ACK responses. UDP flood attacks flood your network with a large amount of UDP packets, requiring the system to verify applications and send responses.

The flood vector tracks packets per destination address. Packets to a specific destination that meet the defined Single Endpoint Flood criteria, and exceed the rate limit, are dropped.

The BIG-IP® system can detect such attacks with a configurable detection threshold, and can rate limit packets from a source when the detection threshold is reached.

You can configure DoS sweep and flood prevention to detect and prevent floods and sweeps of ICMP, UDP, TCP SYN without ACK, or any IP packets that originate from a single source address, according to the threshold setting. Both IPv4 and IPv6 are supported. The sweep vector acts first, so a packet flood from a single source address to a single destination address is handled by the sweep vector.

You can configure DoS sweep and flood prevention through DoS Protection >Device Configuration > Network Security.

#### Task list

*Detecting and protecting against single endpoint DoS flood attacks*
*Detecting and protecting against DoS sweep attacks*
*Detecting and protecting against UDP flood attacks*
*Allowing addresses to bypass global DoS checks*

### Detecting and protecting against single endpoint DoS flood attacks

With the DoS Protection Device Configuration screen settings, you can set detection thresholds and rate limits for DoS flood attacks.

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Network Security**.
   The Network Security screen opens to Device Configuration.
2. In the **Category** column, expand the **Single-Endpoint** category.
3. Click **Single Endpoint Flood**.
   The **Single Endpoint Flood** Properties pane opens on the right side of the screen.
4. On the Properties pane, for **State**, select **Mitigate**.
5. From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.

- Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
- Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.

6. From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.
7. Select the **Add Destination Address to Category** check box to enable automatic blacklisting.
8. From the **Category Name** list, select a black list category to apply to automatically blacklisted addresses.
9. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
10. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).
11. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** > ***Options*** > ***External Redirection*** > ***Blacklist Publisher****.*
12. In the **Packet Type** area, select the packet types you want to detect for this attack type in the **Available** list, and click **<<** to move them to the **Selected** list.
13. Click the **Update** button.
    The flood attack configuration is updated on the Device Protection screen.

Now you have configured the system to provide protection against DoS flood attacks, and to allow such attacks to be identified in system logs and reports.

Configure sweep attack prevention, and configure any other DoS responses, in the DoS device configuration. Configure whitelist entries for addresses that you specifically want to bypass all DoS checks. Configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system.

## Detecting and protecting against DoS sweep attacks

With the DoS Protection Device Configuration screen settings, you can set detection thresholds and rate limits for DoS sweep attacks, and automatically blacklist IP addresses that you detect perpetrating such attacks.

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Network Security**.
   The Network Security screen opens to Device Configuration.
2. In the **Category** column, expand the **Single-Endpoint** category.
3. Click **Single Endpoint Sweep**.
   The Single Endpoint Sweep Properties pane opens on the right side of the screen.
4. On the Properties pane, for **State**, select **Mitigate**.
5. From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.

- Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
- Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.

6. From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.
7. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

   ***Important:*** *For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy:* ***Security*** *>* ***Network Firewall*** *>* ***IP Intelligence*** *>* ***Policies****. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

8. From the **Category Name** list, select a black list category to apply to automatically blacklisted addresses.
9. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
10. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).
11. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** *>* ***Options*** *>* ***External Redirection*** *>* ***Blacklist Publisher****.*

12. In the **Packet Type** area, select the packet types you want to detect for this attack type in the **Available** list, and click **<<** to move them to the **Selected** list.
13. Click the **Update** button.
    The sweep attack configuration is updated on the Device Protection screen.

Now you have configured the system to provide protection against DoS sweep attacks, to allow such attacks to be identified in system logs and reports, and to automatically add such attackers to a blacklist of your choice.

Configure flood attack prevention, and configure any other DoS responses, in the DoS device configuration. Configure whitelist entries for addresses that you specifically choose to bypass all DoS checks. Configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system.

## Detecting and protecting against UDP flood attacks

With the DoS Protection Device Configuration screen settings, you can set detection thresholds and rate limits for UDP flood attacks.

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Network Security**.
   The Network Security screen opens to Device Configuration.
2. In the **Category** column, expand the **Flood** category.
3. Click **UDP Flood**.
   The UDP Flood Properties pane opens on the right side of the screen.

4. On the Properties pane, for **State**, select **Mitigate**.
5. For **Threshold Mode**, select **Fully Manual**.
6. From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
7. From the **Detection Threshold Percent** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in percentage of traffic) for the attack detection threshold. If packets of the specified types cross the percentage threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
8. From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.
   - Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.
   - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.
9. Select **Simulate Auto Threshold** to log the results of the current automatic thresholds, when enforcing manual thresholds.
10. To detect IP address sources from which possible attacks originate, enable **Bad Actor Detection**.

    ***Note:*** *Bad Actor Detection is not available for every vector.*

11. In the **Per Source IP Detection Threshold EPS** field, specify the number of events of this type per second from one IP address that identifies the IP source as a bad actor, for purposes of attack detection and logging.
12. In the **Per Source IP Mitigation Threshold EPS** field, specify the number of events of this type per second from one IP address, above which rate limiting or leak limiting occurs.
13. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

    ***Important:*** *For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy:* ***Security*** > ***Network Firewall*** > ***IP Intelligence*** > ***Policies****. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

14. From the **Category Name** list, select a black list category to apply to automatically blacklisted addresses.
15. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
16. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).
17. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** > ***Options*** > ***External Redirection*** > ***Blacklist Publisher****.*

18. Select **Attacked Destination Detection** to configure automatic blacklisting for attacked destination IP addresses.
19. From the **Port List Type** list, select **Include All Ports** or **Exclude All Ports**.

    An *Include* list checks all the ports you specify in the Port List, using the specified threshold criteria, and ignores all others.

    An *Exclude* list excludes all the ports you specify in the Port List from checking, using the specified threshold criteria, and checks all others. To check all UDP ports, specify an empty exclude list.
20. In the **UDP Port List** area, type a port number to add to an exclude or include UDP port list.
21. In the **UDP Port List** area, select the mode for each port number you want to add to an exclude or include UDP port list.

    - **None** does not include or exclude the port.
    - **Source only** includes or excluded the port from source packets only.
    - **Destination only** includes or excludes the port for destination packets only.
    - **Both Source and Destination** includes or excludes the port in both source and destination packets.
22. Click the **Update** button.
    The UDP Flood attack configuration is updated on the DoS Device Configuration screen.

You have now configured the system to provide customized protection against UDP flood attacks, and to allow such attacks to be identified in system logs and reports.

Configure sweep and flood attack prevention, and configure any other DoS responses, in the DoS device configuration screens. Configure whitelist entries for addresses that you specifically choose to bypass all DoS checks. Configure SNMP traps, logging, and reporting for DoS attacks, to track threats to your system.

## Allowing addresses to bypass global DoS checks

You can specify whitelist addresses that the DoS Device Configuration do not subject to DoS checks. Whitelist entries are specified on a security address list, and can be configured directly on the Device DoS Configuration screen.

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Properties**.
   The DoS Protection Device Configuration screen opens.
2. In the **Whitelist Address List** field, begin typing the name of the address list to use as the whitelist, and select the address list when the name appears.
3. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the **+** next to **Address Lists**.
   The Address List Properties pane opens.
4. Type a **Name** for the address list.
5. Optionally, type a **Description** for the address list.
6. In the **Contents** field, type an address, and click **Add**. Repeat this step to add all items you want on the whitelist.

   You can type an IP address, a geographic location, or the name of another address list. Begin typing, and select the object when the name appears.
7. Click **Update** to update the address list.

   If this is a new address list, type and select the address list in the **Whitelist Address List** field.
8. Click **Commit Changes to System** to commit the whitelist to the device configuration.

You have now specified a whitelist to bypass DoS checks for specific addresses globally.

# Detecting and Preventing DNS DoS Attacks on a Virtual Server

## About preventing DNS DoS attacks on a virtual server

DNS DoS protection is a type of protocol security. DNS DoS attack detection and prevention serves several functions:

- To detect and report on DNS packets based on behavior characteristics of the sender, or characteristics of the packets, without enforcing any rate limits.
- To detect, report on, and rate limit DNS packets based on behavior characteristics that signify specific known attack vectors.
- To identify Bad Actor IP addresses from which attacks appear to originate, by detecting packets per second from a source, and to apply rate limits to such IP addresses.
- To blacklist Bad Actor IP addresses, with configurable detection times, blacklist durations, and blacklist categories, and allow such IP addresses to be advertised to edge routers to offload blacklisting.

You can use the DNS DoS Protection profile to configure the percentage increase over the system baseline, which indicates that a possible attack is in process on a particular DNS query type, or an increase in anomalous packets. You can also rate limit packets of known vectors. You can configure settings manually, and for many vectors you can allow AFM to manage thresholds automatically.

You can specify an address list as a whitelist, that the DoS checks allow. Whitelisted addresses are passed by the DoS profile, without being subject to the checks in the DoS profile.

Per-virtual server DoS protection requires that your virtual server includes a DoS profile that includes DNS security.

### Task list

*Detecting and protecting against DNS DoS attacks with a DoS profile*

*Creating a custom DNS profile to firewall DNS traffic*

*Assigning a DNS profile to a virtual server*

*Associating a DoS profile with a virtual server*

*Allowing addresses to bypass DoS profile checks*

*Creating a logging profile to log DNS attacks*

*Logging DoS events on a virtual server*

## Detecting and protecting against DNS DoS attacks with a DoS profile

You can configure DNS attack settings in a DoS profile that already exists, or create a new one.

The BIG-IP® system handles DNS attacks that use malformed packets, protocol errors, and malicious attack vectors. Protocol error attack detection settings detect malformed and malicious packets, or packets that are employed to flood the system with several different types of responses, by detecting packets per second and detecting percentage increase in packets over time. You can configure settings to identify and rate limit possible DNS attacks with a DoS profile.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.

2. Click **Create**.
   The New DoS Profile screen opens.
3. In the **Name** field, type the name for the profile.
4. Click **Finished**.
   The DoS Protection: DoS Profiles screen opens.
5. Click the name of the DoS profile you want to modify.
6. Select the **Threshold Sensitivity**.

   Select **Low**, **Medium**, or **High**. A lower setting means the automatic threshold algorithm is less sensitive to changes in traffic and CPU usage, but will also trigger fewer false positives.
7. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
8. To configure DNS security settings, click **Protocol Security**, and choose **DNS Security**.
9. To configure enforcement and settings for a DNS vector, in the **Attack Type** column, click the vector name.
   The vector properties pane opens on the right.
10. From the **State** list, choose the appropriate enforcement option.

    - Select **Mitigate** to enforce the configured DoS vector by examining packets, logging the results of the vector, learning patterns, alerting to trouble, and mitigating the attack (watch, learn, alert, and mitigate).
    - Select **Detect Only** to configure the vector, log the results of the vector without applying rate limits or other actions, and alerting to trouble (watch, learn, and alert).
    - Select **Learn Only** to configure the vector, log the results of the vector, without applying rate limits or other actions (watch and learn).
    - Select **Disabled** to disable logging and enforcement of the DoS vector (no stat collection, no mitigation).
11. For **Threshold Mode**, select whether to have the system determine thresholds for the vector (**Fully Automatic**), have partially automatic settings (**Manual Detection / Auto Mitigation**), or, you can control the settings (**Fully Manual**).

    The settings differ depending on the option you select. Here, we describe the settings for automatic threshold configuration. If you want to set thresholds manually, select one of the manual options and refer to online Help for details on the settings.
12. To allow the DoS vector thresholds to be automatically adjusted, for **Threshold Mode**, select **Fully Automatic**.
    a) In the **Attack Floor EPS** field, type the number of events per second of the vector type to allow at a minimum, before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting defines the minimum packets allowed before automatic thresholds are calculated.
    b) In the **Attack Ceiling EPS** field, specify the absolute maximum allowable for packets of this type before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting rate limits packets to the events per second setting, when specified. To set no hard limit, set this to **Infinite**.
13. To detect IP address sources from which possible attacks originate, enable **Bad Actor Detection**.

    ***Note:*** *Bad Actor Detection is not available for every vector.*
14. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

    ***Important:*** *For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy:* ***Security*** *>* ***Network***

***Firewall*** *>* ***IP Intelligence*** *>* ***Policies****. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

15. From the **Category Name** list, select the blacklist category to which to add blacklist entries generated by **Bad Actor Detection**.
16. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
17. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).
18. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** *>* ***Options*** *>* ***External Redirection*** *>* ***Blacklist Publisher****.*

19. Click **Update** to save your changes.

You have now configured a DoS Protection profile to provide custom responses to malicious DNS protocol attacks, to allow such attacks to be identified in system logs and reports, and to allow rate limiting and other actions when such attacks are detected. DNS queries on particular record types you have configured in the DNS Query Attack Detection area are detected as attacks at your specified thresholds and rate increases, and rate limited as specified.

Associate a DNS profile with a virtual server to enable the virtual server to handle DNS traffic. Associate the DoS Protection profile with a virtual server to apply the settings in the profile to traffic on that virtual server.

## Creating a custom DNS profile to firewall DNS traffic

Ensure that you have a DNS security profile created before you configure this system DNS profile.

You can create a custom DNS profile to configure the BIG-IP® system firewall traffic through the system.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **DNS**.
   The DNS profile list screen opens.
2. Click **Create**.
   The New DNS Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. In the General Properties area, from the **Parent Profile** list, accept the default **dns** profile.
5. Select the **Custom** check box.
6. In the DNS Traffic area, from the **DNS Security** list, select **Enabled**.
7. In the DNS Traffic area, from the **DNS Security Profile Name** list, select the name of the DNS firewall profile.
8. Click **Finished**.

Assign the custom DNS profile to the virtual server that handles the DNS traffic that you want to firewall.

## Assigning a DNS profile to a virtual server

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.

3. From the **Configuration** list, select **Advanced**.
4. From the **DNS Profile** list, select the profile you want to assign to the virtual server.
5. Click **Update**.

The virtual server now handles DNS traffic.

## Associating a DoS profile with a virtual server

You must first create a DoS profile separately, to configure denial-of-service protection for applications, the DNS protocol, or the SIP protocol. For application-level DoS protection, the virtual server requires an HTTP profile (such as the default http).

You add denial-of-service protection to a virtual server to provide enhanced protection from DoS attacks, and track anomalous activity on the BIG-IP® system.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, from the Security menu, choose Policies.
4. To enable denial-of-service protection, from the **DoS Protection Profile** list, select **Enabled**, and then, from the **Profile** list, select the DoS profile to associate with the virtual server.
5. Click **Update** to save the changes.

DoS protection is now enabled, and the DoS Protection profile is associated with the virtual server.

## Allowing addresses to bypass DoS profile checks

You can specify whitelisted addresses that the DoS Profile does not subject to DoS checks. Whitelist entries are specified on a security address list, and can be configured directly on the DoS Profile screen.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click the name of the DoS profile you want to modify.
3. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
4. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the + next to **Address Lists**.
   The Address List Properties pane opens.
5. Type a **Name** for the address list.
6. In the **Contents** field, type an address, and click **Add**. Repeat this step to add all items you want on the whitelist.

   You can type an IP address, a geographic location, or the name of another address list. Begin typing, and select the object when the name appears.
7. Click **Update** to create the address list.

   If this is a new address list, type and select the address list name in the **Default Whitelist** field.
8. Click **Update** to update the DoS Profile.

You have now configured a whitelist of addresses to bypass DoS checks for a DoS profile.

## Creating a logging profile to log DNS attacks

Create a custom logging profile to log DNS DoS events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.

The Logging Profiles list screen opens.

2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a name for the logging profile.
4. Select the **Protocol Security** check box.
5. In the DNS Security area, from the **Publisher** list, select a destination to which the BIG-IP system sends DNS log entries.
6. Select the **Log Dropped Requests** check box, to enable the BIG-IP system to log dropped DNS requests.
7. Select the **Log Filtered Dropped Requests** check box, to enable the BIG-IP system to log DNS requests dropped due to DNS query/header-opcode filtering.

   ***Note:*** *The system does not log DNS requests that are dropped due to errors in the way the system processes DNS packets.*

8. Select the **Log Malformed Requests** check box to enable the BIG-IP system to log malformed DNS requests.
9. Select the **Log Rejected Requests** check box to enable the BIG-IP system to log rejected DNS requests.
10. Select the **Log Malicious Requests** check box to enable the BIG-IP system to log malicious DNS requests.
11. From the **Storage Format** list, select how the BIG-IP system formats the log.

| Option | Description |
|---|---|
| **None** | Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example:<br>`"management_ip_address","bigip_hostname","context_type",`<br>`"context_name","src_ip","dest_ip","src_port",`<br>`"dest_port","vlan","protocol","route_domain",`<br>`"acl_rule_name","action","drop_reason` |
| **Field-List** | Allows you to:<br>• Select, from a list, the fields to be included in the log.<br>• Specify the order the fields display in the log.<br>• Specify the delimiter that separates the content in the log. The default delimiter is the comma character. |
| **User-Defined** | Allows you to:<br>• Select, from a list, the fields to be included in the log.<br>• Cut and paste, in a string of text, the order the fields display in the log. |

12. In the Logging Profile Properties, select the **DoS Protection** check box.
    The DoS Protection tab opens.
13. In the DNS DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS DoS events.

    You can specify publishers for other DoS types in the same profile, for example, for SIP or Application DoS Protection.
14. Click **Finished**.

Assign this custom DoS Protection Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

# Detecting and preventing SIP DoS Attacks on a Virtual Server

## About detecting and preventing SIP DoS attacks on a virtual server

*Session Initiation Protocol (SIP)* is a signaling protocol that is typically used to control communication sessions, such as voice and video calls over IP.

SIP DoS attack detection and prevention serves several functions:

- To detect and report on SIP packets based on behavior characteristics of the sender or characteristics of the packets, without enforcing any rate limits.
- To detect, report on, and rate limit SIP packets based on behavior characteristics that signify specific known attack vectors.
- To identify Bad Actor IP addresses from which attacks appear to originate, by detecting packets per second from a source, and to apply rate limits to such IP addresses.
- To blacklist Bad Actor IP addresses, with configurable detection times, blacklist durations, and blacklist categories, and allow such IP addresses to be advertised to edge routers to offload blacklisting.

You can use a SIP DoS profile to specify the percentage increase over the system baseline, which indicates that a possible attack is in process on a particular SIP method, or an increase in anomalous packets. You can also rate limit packets of known vectors. For all SIP vectors except sip-malformed, the system can manage thresholds automatically or manually. You can manually set thresholds for malformed SIP packets.

You can specify an address list as a whitelist, that the DoS checks allow. Whitelisted addresses are not subject to the checks configured in the DoS profile.

To protect a virtual server from SIP DoS attacks, you need to associate the virtual server with a DoS profile that includes SIP security.

***Important:*** *You must also create a SIP profile, and attach it to the virtual server being protected from SIP DoS attacks.*

### Task list

*Detecting and preventing SIP DoS attacks with a DoS profile*

*Creating a SIP profile for SIP DoS protection*

*Assigning a SIP profile to a virtual server*

*Associating a DoS profile with a virtual server*

*Allowing addresses to bypass DoS profile checks*

*Creating a custom SIP DoS Protection Logging profile*

*Logging DoS events on a virtual server*

## Detecting and preventing SIP DoS attacks with a DoS profile

This task describes how to create a new DoS profile and configure SIP settings to identify SIP attacks at the same time. However, you can also add SIP attack detection settings to an existing DoS profile. The BIG-IP® system handles SIP attacks that include malformed packets, protocol errors, and malicious attack vectors. Protocol error attack detection recognizes malformed and malicious packets, or packets that are employed to flood the system with several different types of responses.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click **Create**.
   The New DoS Profile screen opens.
3. In the **Name** field, type the name for the profile.
4. Click **Finished**.
   The DoS Protection: DoS Profiles screen opens.
5. Click the name of the DoS profile you want to modify.
6. Select the **Threshold Sensitivity**.

   Select **Low**, **Medium**, or **High**. A lower setting means the automatic threshold algorithm is less sensitive to changes in traffic and CPU usage, but will also trigger fewer false positives.
7. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
8. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the **+** next to **Address Lists**.
   The Address List Properties pane opens.
9. To configure SIP security settings, on the menu bar, select **Protocol Security**, and choose **SIP Security**.
10. To change the threshold or rate increase for a particular SIP vector, in the **Attack Type** column, click the vector name.
    The vector properties pane opens on the right.
11. From the **State** list, choose the appropriate enforcement option.
    - Select **Mitigate** to enforce the configured DoS vector by examining packets, logging the results of the vector, learning patterns, alerting to trouble, and mitigating the attack (watch, learn, alert, and mitigate).
    - Select **Detect Only** to configure the vector, log the results of the vector without applying rate limits or other actions, and alerting to trouble (watch, learn, and alert).
    - Select **Learn Only** to configure the vector, log the results of the vector, without applying rate limits or other actions (watch and learn).
    - Select **Disabled** to disable logging and enforcement of the DoS vector (no stat collection, no mitigation).
12. For **Threshold Mode**, select whether to have the system determine thresholds for the vector (**Fully Automatic**), have partially automatic settings (**Manual Detection / Auto Mitigation**), or, you can control the settings (**Fully Manual**).

    The settings differ depending on the option you select. Here, we describe the settings for automatic threshold configuration. If you want to set thresholds manually, select one of the manual options and refer to online Help for details on the settings.
13. To allow the DoS vector thresholds to be automatically adjusted, for **Threshold Mode**, select **Fully Automatic**.
    a) In the **Attack Floor EPS** field, type the number of events per second of the vector type to allow at a minimum, before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting defines the minimum packets allowed before automatic thresholds are calculated.
    b) In the **Attack Ceiling EPS** field, specify the absolute maximum allowable for packets of this type before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting rate limits packets to the events per second setting, when specified. To set no hard limit, set this to **Infinite**.

14. To detect IP address sources from which possible attacks originate, enable **Bad Actor Detection**.

    ***Note:*** *Bad Actor Detection is not available for every vector.*

15. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

    ***Important:*** *For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy:* ***Security*** *>* ***Network Firewall*** *>* ***IP Intelligence*** *>* ***Policies****. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

16. From the **Category Name** list, select the blacklist category to which to add blacklist entries generated by **Bad Actor Detection**.
17. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
18. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).
19. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** *>* ***Options*** *>* ***External Redirection*** *>* ***Blacklist Publisher****.*

20. Click **Update** to save your changes.

You have now configured a DoS profile to provide custom responses to malformed SIP attacks, SIP flood attacks, and to allow such attacks to be identified in system logs and reports.

Now you need to associate the DoS profile with a virtual server to apply the settings in the profile to traffic on that virtual server. When a SIP attack on a specific query type is detected, you can be alerted with various system monitors.

## Creating a SIP profile for SIP DoS protection

You can create a SIP profile if configuring SIP DoS protection.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **SIP (legacy)**.
   The SIP profile list screen opens.
2. Click **Create**.
   The New SIP Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Next to **Settings**, select the **Custom** check box.
5. Select the **SIP Firewall** check box.

   When enabled, the SIP Security settings configured in the DoS Profile apply to the virtual servers that use this profile.
6. Next to **Log Settings**, select the **Custom** check box.
7. From the **Log Publisher** list, select a destination to which the BIG-IP system sends log entries.

   You can specify publishers for other DoS types in the same profile, for example, for DNS, Network, or Application DoS Protection.
8. In the Log Settings area, from the **Logging Profile** list, select a custom Logging profile.
9. Modify other settings, as required.
10. Click **Update**.

A SIP profile is now configured for SIP DoS firewall features.

Assign this SIP profile to a virtual server, along with a DoS profile that includes SIP security, to provide SIP protocol DoS protection on a virtual server.

## Assigning a SIP profile to a virtual server

You need to have created a SIP profile already.

To apply the settings in the SIP profile to traffic, you associate the SIP profile with a virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. From the **SIP Profile** list, select the name of the SIP profile that you previously created.
5. Click **Update**.

The virtual server now uses the SIP settings from the SIP profile.

## Associating a DoS profile with a virtual server

You must first create a DoS profile separately, to configure denial-of-service protection for applications, the DNS protocol, or the SIP protocol. For application-level DoS protection, the virtual server requires an HTTP profile (such as the default http).

You add denial-of-service protection to a virtual server to provide enhanced protection from DoS attacks, and track anomalous activity on the BIG-IP® system.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, from the Security menu, choose Policies.
4. To enable denial-of-service protection, from the **DoS Protection Profile** list, select **Enabled**, and then, from the **Profile** list, select the DoS profile to associate with the virtual server.
5. Click **Update** to save the changes.

DoS protection is now enabled, and the DoS Protection profile is associated with the virtual server.

## Allowing addresses to bypass DoS profile checks

You can specify whitelisted addresses that the DoS Profile does not subject to DoS checks. Whitelist entries are specified on a security address list, and can be configured directly on the DoS Profile screen.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click the name of the DoS profile you want to modify.
3. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
4. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the **+** next to **Address Lists**.
   The Address List Properties pane opens.
5. Type a **Name** for the address list.
6. In the **Contents** field, type an address, and click **Add**. Repeat this step to add all items you want on the whitelist.

You can type an IP address, a geographic location, or the name of another address list. Begin typing, and select the object when the name appears.

7. Click **Update** to create the address list.

   If this is a new address list, type and select the address list name in the **Default Whitelist** field.
8. Click **Update** to update the DoS Profile.

You have now configured a whitelist of addresses to bypass DoS checks for a DoS profile.

## Creating a custom SIP DoS Protection Logging profile

Create a custom Logging profile to log SIP DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
4. In the SIP DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log SIP DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for DNS or Application DoS Protection.
5. Click **Finished**.

Assign this custom SIP DoS Protection Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

# Detecting and Preventing Network DoS Attacks on a Virtual Server

## About detecting and preventing Network DoS attacks on a virtual server

Network DoS protection is a type of security that collects several DoS checks in a DoS profile. Attack detection and prevention serves several functions:

- To detect and report on packets based on behavior characteristics of the sender or characteristics of the packets, without enforcing any rate limits.
- To detect, report on, and rate limit packets based on behavior characteristics that signify specific known attack vectors.
- To identify Bad Actor IP addresses from which attacks appear to originate, by detecting packets per second from a source, and to apply rate limits to such IP addresses.
- To blacklist Bad Actor IP addresses, with configurable detection times, blacklist durations, and blacklist categories, and allow such IP addresses to be advertised to edge routers to offload blacklisting.

You can configure the Network DoS Protection profile to detect possible attack vectors by packet-per-second or percentage-increase-over-time thresholds, which can indicate that a possible attack is in process. Such attacks can be logged and reported through system logging facilities. You can also rate limit packets of known vectors. You can configure settings manually, and for many vectors you can allow AFM to manage thresholds automatically.

You can specify an address list as a whitelist that the DoS checks allow. Whitelisted addresses are passed by the DoS profile, without being subject to the checks in the DoS profile.

Per-virtual server DoS protection requires that your virtual server includes a DoS profile that includes network security.

#### Task list

*Detecting and protecting a virtual server against network DoS attacks with a DoS profile*

*Associating a DoS profile with a virtual server*

*Allowing addresses to bypass DoS profile checks*

*Creating a custom Network Firewall Logging profile*

*Logging DoS events on a virtual server*

## Detecting and protecting a virtual server against network DoS attacks with a DoS profile

The BIG-IP® system handles network attacks that use malformed packets and malicious attack vectors. Possible malicious packets and attacks are detected by logging when packets exceed a threshold of packets per second, and by detecting the rate increase percentage in packets of a certain type over time. You can configure settings to identify and rate limit possible network attacks with a DoS profile. For many vectors, you can also automatically blacklist IP addresses.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click **Create**.
   The New DoS Profile screen opens.

3. In the **Name** field, type the name for the profile.
4. Click **Finished**.
   The DoS Protection: DoS Profiles screen opens.
5. Click the name of the DoS profile you want to modify.
6. Select the **Threshold Sensitivity**.

   Select **Low**, **Medium**, or **High**. A lower setting means the automatic threshold algorithm is less sensitive to changes in traffic and CPU usage, but will also trigger fewer false positives.
7. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
8. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the **+** next to **Address Lists**.
   The Address List Properties pane opens.
9. In the **Contents** field, type an address, and click **Add**. Repeat this step to add all items you want on the whitelist.

   You can type an IP address, a geographic location, or the name of another address list. Begin typing, and select the object when the name appears.
10. Click **Update**.
11. To configure network security settings, click **Network Security**.
12. To change the threshold or rate increase for a particular network attack, in the **Attack Type** column, click the name of the attack.
    The DoS attack Properties pane appears on the right side of the screen.
13. From the **State** list, choose the appropriate enforcement option.

    - Select **Mitigate** to enforce the configured DoS vector by examining packets, logging the results of the vector, learning patterns, alerting to trouble, and mitigating the attack (watch, learn, alert, and mitigate).
    - Select **Detect Only** to configure the vector, log the results of the vector without applying rate limits or other actions, and alerting to trouble (watch, learn, and alert).
    - Select **Learn Only** to configure the vector, log the results of the vector, without applying rate limits or other actions (watch and learn).
    - Select **Disabled** to disable logging and enforcement of the DoS vector (no stat collection, no mitigation).
14. To allow the DoS vector thresholds to be automatically adjusted, for **Threshold Mode**, select **Fully Automatic**.
    a) In the **Attack Floor EPS** field, type the number of events per second of the vector type to allow at a minimum, before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting defines the minimum packets allowed before automatic thresholds are calculated.

    b) In the **Attack Ceiling EPS** field, specify the absolute maximum allowable for packets of this type before automatically calculated thresholds are determined.

       Because automatic thresholds take time to be reliably established, this setting rate limits packets to the events per second setting, when specified. To set no hard limit, set this to **Infinite**.
15. To configure DoS vector thresholds manually, for **Threshold Mode**, select **Fully Manual**.
    a) From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.

       Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.

       Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.

    b) From the **Detection Threshold Percent** list, select **Specify** or **Infinite**.

Use **Specify** to set a value (in percentage of traffic) for the attack detection threshold. Use **Infinite** to set no value for the threshold.

c) From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.

Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.

Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.

16. From the **Detection Threshold EPS** list, select **Specify** or **Infinite**.
    - Use **Specify** to set a value (in events per second) for the attack detection threshold. If packets of the specified types cross the threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
    - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
17. From the **Detection Threshold Percent** list, select **Specify** or **Infinite**.
    - Use **Specify** to set a value (in percentage of traffic) for the attack detection threshold. If packets of the specified types cross the percentage threshold, an attack is logged and reported. The system continues to check every second, and registers an attack as long as the threshold is exceeded.
    - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not logged or reported based on this threshold.
18. From the **Mitigation Threshold EPS** list, select **Specify** or **Infinite**.
    - Use **Specify** to set a value (in events per second), which cannot be exceeded. If the number of events of this type exceeds the threshold, excess events are dropped until the rate no longer exceeds the threshold.
    - Use **Infinite** to set no value for the threshold. This specifies that this type of attack is not rate-limited.
19. Select **Simulate Auto Threshold** to log the results of the current automatic thresholds, when enforcing manual thresholds.
20. To detect IP address sources from which possible attacks originate, enable **Bad Actor Detection**.

***Note:** Bad Actor Detection is not available for every vector.*

21. In the **Per Source IP Detection Threshold EPS** field, specify the number of events of this type per second from one IP address that identifies the IP source as a bad actor, for purposes of attack detection and logging.
22. In the **Per Source IP Mitigation Threshold EPS** field, specify the number of events of this type per second from one IP address, above which rate limiting or leak limiting occurs.
23. To automatically blacklist bad actor IP addresses, select **Add Source Address to Category**.

***Important:** For this to work, you need to assign an IP Intelligence policy to the appropriate context (device, virtual server, or route domain). For the device, assign a global policy: **Security** > **Network Firewall** > **IP Intelligence** > **Policies**. For the virtual server or route domain, assign the IP Intelligence policy on the Security tab.*

24. From the **Category Name** list, select a black list category to apply to automatically blacklisted addresses.
25. In the **Sustained Attack Detection Time** field, specify the duration in seconds after which the attacking endpoint is blacklisted. By default, the configuration adds an IP address to the blacklist after one minute (60 seconds).
26. In the **Category Duration Time** field, specify the length of time in seconds that the address will remain on the blacklist. The default is `14400` seconds (4 hours).

27. To allow IP source blacklist entries to be advertised to edge routers so they will null route their traffic, select **Allow External Advertisement**.

    ***Note:*** *To advertise to edge routers, you must configure a Blacklist Publisher and Publisher Profile at* ***Security*** *>* ***Options*** *>* ***External Redirection*** *>* ***Blacklist Publisher****.*

28. Click **Update** to save your changes. The DoS vector is updated on the Network Security screen.

You have now configured a DoS Protection profile to analyze network packet behavior for DoS attacks, to allow specific configured attacks to be identified in system logs and reports, and to allow rate limiting of such attacks. DNS queries on particular record types you have configured in the DNS Query Attack Detection area are detected as attacks at your specified thresholds and rate increases, and rate limited as specified.

Associate the DoS profile with a virtual server to enable network DoS protection.

## DoS profile attack types

You can specify specific threshold, rate increase, rate limit, and other parameters for supported network DoS attack types, to more accurately detect, track, and rate limit attacks.

***Attention:*** *All hardware-supported vectors are performed in hardware on vCMP® guests, provided that the vCMP guests have the same software version as the vCMP host.*

| DoS Category | Attack Name | Dos Vector Name | Information | Hardware accelerated |
|---|---|---|---|---|
| + | TTL <= <tunable> | ttl-leq-one | An IP packet with a destination that is not multicast and that has a TTL greater than 0 and less than or equal to a tunable value, which is 1 by default. To tune this value, in tmsh: `modify sys db dos.iplowttl value`, where `value` is `1-4`. | Yes |
| + | IP Option Frames | ip-opt-frames | IPv4 address packet with `option.db variable tm.acceptipsourceroute` must be enabled to receive IP options | Yes |
| + | IPv6 extension header too large | ext-hdr-too-large | An extension header is too large. To tune this value, in tmsh: `modify sys db dos.maxipv6extsize value`, where `value` is `0-1024`. | Yes |
| + | IPv6 hop count <= <tunable> | hop-cnt-leq-one | The IPv6 extended header hop count is less than or equal to <tunable>. To tune this value, in tmsh: `modify sys db dos.ipv6lowhopcnt value`, where `value` is `1-4`. | Yes |
| + | IPv6 Extended Header Frames | ipv6-ext-hdr-frames | IPv6 address contains extended header frames | Yes |
| + | Too Many Extended Headers | too-many-ext-hdrs | For an IPv6 address, there are more than <tunable> extended headers (the default is `4`). To tune this value, in tmsh: `modify sys db dos.maxipv6exthdrs value`, where `value` is `0-15`. | Yes |

| DoS Category | Attack Name | Dos Vector Name | Information | Hardware accelerated |
|---|---|---|---|---|
| + | Option Present With Illegal Length | opt-present-with-illegal-len | Option present with illegal length | Yes |
| + | TCP Bad URG | tcp-bad-urg | Packet contains a bad URG flag, this is likely malicious | Yes |
| + | TCP Option Overruns TCP Header | tcp-opt-overruns-tcp-hdr | The TCP option bits overrun the TCP header. | Yes |
| + | Unknown TCP Option Type | unk-tcp-opt-type | Unknown TCP option type | Yes |
| + | ICMPv4 Flood | icmpv4-flood | Flood with ICMP v4 packets | Yes |
| + | ICMPv6 Flood | icmpv6-flood | Flood with ICMP v6 packets | Yes |
| + | IP Fragment Flood | ip-frag-flood | Fragmented packet flood with IPv4 | Yes |
| + | IPv6 Fragment Flood | ipv6-frag-flood | Fragmented packet flood with IPv6 | No |
| + | TCP RST Flood | tcp-rst-flood | TCP RST flood | Yes |
| + | TCP SYN ACK Flood | tcp-synack-flood | TCP SYN/ACK flood | Yes |
| + | TCP SYN Flood | tcp-syn-flood | TCP SYN flood | Yes |
| + | TCP Window Size | tcp-window-size | The TCP window size in packets exceeds the maximum. To tune this value, in tmsh: `modify sys db dos.tcplowwindowsize value`, where `value` is <=`128`. | Yes |
| + | TCP SYN Oversize | tcp-syn-oversize | Detects TCP data SYN packets larger than the maximum specified by the dos.maxsynsize parameter. To tune this value, in tmsh: `modify sys db dos.maxsynsize value`. The default size is `64` and the maximum allowable value is `9216`. | Yes |
| + | UDP Flood | udp-flood | UDP flood attack | Yes |
| + | ICMP Fragment | icmp-frag | ICMP fragment flood | Yes |
| + | Sweep | sweep | Sweep on a single endpoint. You can configure packet types to check for, and packets per second for both detection and rate limiting. You can also | No |

| DoS Category | Attack Name | Dos Vector Name | Information | Hardware accelerated |
|---|---|---|---|---|
| | | | configure automatic blacklisting for IPs that initiate sweep attacks, using the IP intelligence mechanism. | |
| + | Host Unreachable | host-unreachable | Host unreachable error | Yes |
| + | TIDCMP | tidcmp | ICMP source quench attack | Yes |

## Associating a DoS profile with a virtual server

You must first create a DoS profile separately, to configure denial-of-service protection for applications, the DNS protocol, or the SIP protocol. For application-level DoS protection, the virtual server requires an HTTP profile (such as the default http).

You add denial-of-service protection to a virtual server to provide enhanced protection from DoS attacks, and track anomalous activity on the BIG-IP® system.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, from the Security menu, choose Policies.
4. To enable denial-of-service protection, from the **DoS Protection Profile** list, select **Enabled**, and then, from the **Profile** list, select the DoS profile to associate with the virtual server.
5. Click **Update** to save the changes.

DoS protection is now enabled, and the DoS Protection profile is associated with the virtual server.

## Allowing addresses to bypass DoS profile checks

You can specify whitelisted addresses that the DoS Profile does not subject to DoS checks. Whitelist entries are specified on a security address list, and can be configured directly on the DoS Profile screen.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click the name of the DoS profile you want to modify.
3. If you have created a whitelist on the system, in the **Default Whitelist** field, begin typing the name of the address list to use as the whitelist, and select the list when the name appears.
4. To define an address list to use as a whitelist, on the right side of the screen in the Shared Objects pane, click the **+** next to **Address Lists**.
   The Address List Properties pane opens.
5. Type a **Name** for the address list.
6. In the **Contents** field, type an address, and click **Add**. Repeat this step to add all items you want on the whitelist.

   You can type an IP address, a geographic location, or the name of another address list. Begin typing, and select the object when the name appears.
7. Click **Update** to create the address list.

   If this is a new address list, type and select the address list name in the **Default Whitelist** field.
8. Click **Update** to update the DoS Profile.

You have now configured a whitelist of addresses to bypass DoS checks for a DoS profile.

## Creating a custom Network Firewall Logging profile

Create a custom Logging profile to log messages about BIG-IP® system Network Firewall events.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select the **Network Firewall** check box.
5. In the Network Firewall area, from the **Publisher** list, select the publisher the BIG-IP system uses to log Network Firewall events.
6. Set an **Aggregate Rate Limit** to define a rate limit for all combined network firewall log messages per second.

   Beyond this rate limit, log messages are not logged.
7. For the **Log Rule Matches** setting, select how the BIG-IP system logs packets that match ACL rules. You can select any or all of the options.

   | Option | Description |
   |---|---|
   | **Option** | Enables or disables logging of packets that match ACL rules configured with: |
   | **Accept** | `action=Accept` |
   | **Drop** | `action=Drop` |
   | **Reject** | `action=Reject` |

   When an option is selected, you can configure a rate limit for log messages of that type.
8. Select the **Log IP Errors** check box, to enable logging of IP error packets.

   When this setting is enabled, you can configure a rate limit for log messages of this type.
9. Select the **Log TCP Errors** check box, to enable logging of TCP error packets.

   When this is enabled, you can configure a rate limit for log messages of this type.
10. Select the **Log TCP Events** check box, to enable logging of open and close of TCP sessions.

    When this is enabled, you can configure a rate limit for log messages of this type.
11. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for Network Firewall log events.
12. Enable the **Log Geolocation IP Address** setting to specify that when a geolocation event causes a network firewall action, the associated IP address is logged.
13. From the **Storage Format** list, select how the BIG-IP system formats the log.

    | Option | Description |
    |---|---|
    | **None** | Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example: `"management_ip_address","bigip_hostname","context_type","context_name","src_ip","dest_ip","src_port","dest_port","vlan","protocol","route_domain","acl_rule_name","action","drop_reason` |
    | **Field-List** | Allows you to:<br>• Select, from a list, the fields to be included in the log. |

| Option | Description |
|---|---|
| | • Specify the order the fields display in the log.<br>• Specify the delimiter that separates the content in the log. The default delimiter is the comma character. |
| **User-Defined** | Allows you to:<br>• Select, from a list, the fields to be included in the log.<br>• Cut and paste, in a string of text, the order the fields display in the log. |

14. In the IP Intelligence area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log source IP addresses, which are identified and configured for logging by an IP Intelligence policy.

    ***Note:** The IP Address Intelligence feature must be enabled and licensed.*

15. Set an **Aggregate Rate Limit** to define a rate limit for all combined IP Intelligence log messages per second.

    Beyond this rate limit, log messages are not logged.
16. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for IP Intelligence log events.
17. In the Traffic Statistics area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log traffic statistics.
18. For the **Log Timer Events** setting, enable **Active Flows** to log the number of active flows each second.
19. For the **Log Timer Events** setting, enable **Reaped Flows**to log the number of reaped flows, or connections that are not established because of system resource usage levels.
20. For the **Log Timer Events** setting, enable **Missed Flows** to log the number of packets that were dropped because of a flow table miss. A flow table miss occurs when a TCP non-SYN packet does not match an existing flow.
21. For the **Log Timer Events** setting, enable **SYN Cookie (Per Session Challenge)** to log the number of SYN cookie challenges generated each second.
22. For the **Log Timer Events** setting, enable **SYN Cookie (White-listed Clients)** to log the number of SYN cookie clients whitelisted each second.
23. Click **Finished**.

Assign this custom network firewall Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

# Detecting Dynamic DoS Attacks

## About detecting dynamic DoS attacks

A *dynamic DoS attack* is a DoS attack that doesn't fit predefined DoS vector criteria. Using dynamic DoS attack detection, such attacks can be detected and mitigated automatically by AFM. Dynamic DoS detection creates vector signatures for attacks based on changing traffic patterns over time. When an attack is detected, a vector signature is created and added to a list of dynamic vectors. All packets are then checked against the dynamic vector, and mitigated according to internal logic. When packet processing on the system falls back to normal levels, the signature no longer appears as an attack, and is removed from the dynamic signature list.

### Detection modes

The following modes are available for dynamic DoS detection.

**Disabled**
In this mode, no dynamic DoS detection occurs.

**Learn-Only**
In this mode, the system establishes a baseline for packet processing on AFM. Learn-Only mode detects traffic patterns, establishes a baseline, and detects anomalies, but does not generate logs or dynamic DoS vector signatures. Attacks are not mitigated in Learn-Only mode.

**Enabled**
In this mode, the system monitors traffic, compares traffic changes over time, and detects anomalies. Attacks are logged, dynamic DoS vector signatures are generated, packets are compared to the dynamic DoS vector signature, and attacks are mitigated. When an attack ceases, the dynamic DoS vector signature is removed from the list of signatures.

### Mitigation Sensitivity

Mitigation sensitivity establishes how aggressively the system mitigates dynamic DoS attacks. You can set this to **None**, **Low**, **Medium**, or **High**. By default, mitigation sensitivity is set to **None**. **Low** sensitivity is the least aggressive, at the risk of allowing more attack packets through. **High** drops packets more aggressively, even when attack confidence is lower.

### Redirection/Scrubbing

You can configure redirection and scrubbing to handle mitigation of dynamic DoS signatures with an IP Intelligence category. The following parameters can be configured for redirection and scrubbing.

**Scrubbing Category**
You can select an IP Intelligence category for IP addresses blocked by dynamic DoS signatures. By default, the IP intelligence category for scrubbed IP addresses is **attacked_ips**.

**Scrubbing Advertisement Time**
This is the duration for which a mitigated IP is advertised to the IP Intelligence mechanism for scrubbing. The default is `300` seconds.

### Start Relearning

The **Start Relearning** option clears historical data, thresholds and signatures for the dynamic DoS detection system. The Dynamic DoS signature baseline is re-established. Relearning occurs for a period of 20 minutes.

## Detecting global dynamic DoS attacks

You enable dynamic DoS signatures at the device level to dynamically detect and mitigate network DoS attacks.

1. On the Main tab, click **Security** > **DoS Protection** > **Device Configuration** > **Network Security**.
   The Network Security screen opens to Device Configuration.
2. Under Dynamic Signatures, from the **Enforcement** list, select **Enabled**.
3. From the **Mitigation Sensitivity** list, select the sensitivity level for dropping packets.
   - Select **None** to generate and log dynamic signatures, without dropping packets.
   - To drop packets, set the mitigation level from **Low** to **High**. A setting of **Low** is least aggressive, but will also trigger fewer false positives. A setting of **High** is most aggressive, and the system may drop more false positive packets.
4. To have dynamic signatures handled by an IP Intelligence category, from the **Redirection/Scrubbing** list, select **Enabled**.
5. If you are using Redirection/Scrubbing to handle dynamic signatures, from the **Scrubbing Category** list, select the IP Intelligence category with which scrubbed packets are to be categorized.
6. In the **Scrubbing Advertisement Time** field, specify the amount of time to advertise the scrubbed IP address to the IP Intelligence category.
7. Click **Update** to save the device configuration.

## Detecting dynamic DoS network attacks on a virtual server

You enable dynamic DoS signatures on a virtual server to detect dynamic DoS attacks at a more granular level than the global level.

1. On the Main tab, click **Security** > **DoS Protection** > **DoS Profiles**.
   The DoS Profiles list screen opens.
2. Click the name of an existing DoS profile (or create a new one).
   The DoS Profile Properties screen for that profile opens.
3. To configure network security settings, click **Network Security**.
4. Under Dynamic Signatures, from the **Enforcement** list, select **Enabled**.
5. From the **Mitigation Sensitivity** list, select the sensitivity level for dropping packets.
   - Select **None** to generate and log dynamic signatures, without dropping packets.
   - To drop packets, set the mitigation level from **Low** to **High**. A setting of **Low** is least aggressive, but will also trigger fewer false positives. A setting of **High** is most aggressive, and the system may drop more false positive packets.
6. To have dynamic signatures handled by an IP Intelligence category, from the **Redirection/Scrubbing** list, select **Enabled**.
7. In the **Scrubbing Advertisement Time** field, specify the amount of time to advertise the scrubbed IP address to the IP Intelligence category.
8. Click **Update** to save the DoS profile.

You have configured the DoS profile to detect dynamic DoS vectors and mitigate such attacks.

Next, you associate the DoS profile with a virtual server to enable network DoS protection.

# SNMP Trap Configuration

## Overview: SNMP trap configuration

SNMP *traps* are definitions of unsolicited notification messages that the BIG-IP® alert system and the SNMP agent send to the SNMP manager when certain events occur on the BIG-IP system. Configuring SNMP traps on a BIG-IP system means configuring how the BIG-IP system handles traps, as well as setting the destination to which the notifications are sent.

The BIG-IP system stores SNMP traps in two specific files:

**`/etc/alertd/alert.conf`**
Contains default SNMP traps.

> ***Important:*** *Do not add or remove traps from the `/etc/alertd/alert.conf` file.*

**`/config/user_alert.conf`**
Contains user-defined SNMP traps.

### Task summary

Perform these tasks to configure SNMP traps for certain events and set trap destinations.

*Enabling traps for specific events*
*Setting v1 and v2c trap destinations*
*Setting v3 trap destinations*
*Viewing pre-configured SNMP traps*
*Creating custom SNMP traps*

## Enabling traps for specific events

You can configure the SNMP agent on the BIG-IP® system to send, or refrain from sending, notifications to the traps destinations.

1. On the Main tab, click **System** > **SNMP** > **Traps** > **Configuration**.
2. To send traps when an administrator starts or stops the SNMP agent, verify that the **Enabled** check box for the **Agent Start/Stop** setting is selected.
3. To send notifications when authentication warnings occur, select the **Enabled** check box for the **Agent Authentication** setting.
4. To send notifications when certain warnings occur, verify that the **Enabled** check box for the **Device** setting is selected.
5. Click **Update**.

The BIG-IP system automatically updates the `alert.conf` file.

## Setting v1 and v2c trap destinations

You specify the IP address of the SNMP manager in order for the BIG-IP® system to send notifications.

1. On the Main tab, click **System** > **SNMP** > **Traps** > **Destination**.
2. Click **Create**.

3. For the **Version** setting, select either `v1` or `v2c`.
4. In the **Community** field, type the community name for the SNMP agent running on the BIG-IP system.
5. In the **Destination** field, type the IP address of the SNMP manager.
6. In the **Port** field, type the port number on the SNMP manager that is assigned to receive the traps.
7. For the **Network** setting, select a trap network.

   The BIG-IP system sends the SNMP trap out of the network you select.
8. Click **Finished**.

## Setting v3 trap destinations

You specify the destination SNMP manager to which the BIG-IP® system sends notifications.

1. On the Main tab, click **System** > **SNMP** > **Traps** > **Destination**.
2. Click **Create**.
3. For the **Version** setting, select `v3`.
4. In the **Destination** field, type the IP address of the SNMP manager.
5. In the **Port** field, type the port number on the SNMP manager that is assigned to receive the traps.
6. For the **Network** setting, select a trap network.

   The BIG-IP system sends the SNMP trap out of the network you select.
7. From the **Security Level** list, select the level of security at which you want SNMP messages processed.

| Option | Description |
|---|---|
| **Auth, No Privacy** | Process SNMP messages using authentication but without encryption. When you use this value, you must also provide values for the **Security Name**, **Authentication Protocol**, and **Authentication Password** settings. |
| **Auth and Privacy** | Process SNMP messages using authentication and encryption. When you use this value, you must also provide values for the **Security Name**, **Authentication Protocol**, **Authentication Password**, **Privacy Protocol**, and **Privacy Password** settings. |

8. In the **Security Name** field, type the user name the system uses to handle SNMP v3 traps.
9. In the **Engine ID** field, type an administratively unique identifier for an SNMP engine. (This setting is optional.) You can find the engine ID in the `/config/net-snmp/snmpd.conf` file on the BIG-IP system.

   Note that this ID is identified in the file as the value of the oldEngineID token.
10. From the **Authentication Protocol** list, select the algorithm the system uses to authenticate SNMP v3 traps.

    When you set this value, you must also enter a value in the **Authentication Password** field.
11. In the **Authentication Password** field, type the password the system uses to handle an SNMP v3 trap.

    When you set this value, you must also select a value from the **Authentication Protocol** list.

    ***Note:** The authentication password must be at least 8 characters long.*
12. If you selected **Auth and Privacy** from the **Security Level** list, from the **Privacy Protocol** list, select the algorithm the system uses to encrypt SNMP v3 traps.

    When you set this value, you must also enter a value in the **Privacy Password** field.

13. If you selected **Auth and Privacy** from the **Security Level** list, in the **Privacy Password** field, type the password the system uses to handle an encrypted SNMP v3 trap.

    When you set this value, you must also select a value from the **Privacy Protocol** list.

    ***Note:*** *The authentication password must be at least 8 characters long.*

14. Click **Finished**.

## Viewing pre-configured SNMP traps

Verify that your user account grants you access to the advanced shell.

Pre-configured traps are stored in the `/etc/alertd/alert.conf` file. View these SNMP traps to understand the data that the SNMP manager can use.

Use this command to view the SNMP traps that are pre-configured on the BIG-IP® system: `cat /etc/alertd/alert.conf`.

## Creating custom SNMP traps

Verify that your user account grants you access to tmsh.

Create custom SNMP traps that alert the SNMP manager to specific SNMP events that occur on the network when the pre-configured traps do not meet all of your needs.

1. Log in to the command line.
2. Create a backup copy of the file `/config/user_alert.conf`, by typing this command: `cp /config/user_alert.conf` *`backup_file_name`*
   For example, type: `cp /config/user_alert.conf /config/user_alert.conf.backup`
3. With a text editor, open the file `/config/user_alert.conf`.
4. Add a new SNMP trap.

   The required format is:

```
alert alert_name "matched message" {
    snmptrap OID=".1.3.6.1.4.1.3375.2.4.0.XXX"
   }
```

- *`alert_name`* represents a descriptive name. The *`alert_name`* or *`matched_message`* value cannot match the corresponding value in any of the SNMP traps defined in the `/etc/alertd/alert.conf` or `/config/user_alert.conf` file.
- *`matched_message`* represents the text that matches the Syslog message that triggers the custom trap. You can specify either a portion of the Syslog message text or use a regular expression. Do not include the Syslog prefix information, such as the date stamp and process ID, in the match string.
- The *`XXX`* portion of the OID value represents a number that is unique to this OID. Specify any OID that meets all of these criteria:
  - Is in standard OID format and within the range `.1.3.6.1.4.1.3375.2.4.0.300` through `.1.3.6.1.4.1.3375.2.4.0.999`.
  - Is in a numeric range that can be processed by your trap receiving tool.
  - Does not exist in the MIB file `/usr/share/snmp/mibs/F5-BIGIP-COMMON-MIB.txt`.
  - Is not used in another custom trap.

As an example, to create a custom SNMP trap that is triggered whenever the system logs switchboard failsafe status changes, add the following trap definition to `/config/user_alert.conf`.

```
alert SWITCHBOARD_FAILSAFE_STATUS "Switchboard Failsafe (.*)" {
        snmptrap OID=".1.3.6.1.4.1.3375.2.4.0.500"
    }
```

This trap definition causes the system to log the following message to the file `/var/log/ltm`, when switchboard failsafe is enabled: `Sep 23 11:51:40 bigip1.askf5.com lacpd[27753]: 01160016:6: Switchboard Failsafe enabled`.

5. Save the file.
6. Close the text editor.
7. Restart the `alertd` daemon by typing this command: `bigstart restart alertd`

   If the `alertd` daemon fails to start, examine the newly-added trap entry to ensure that the format is correct.

# Configuring High-Speed Remote Logging of DoS Events

## Overview: Configuring DoS Protection event logging

You can configure the BIG-IP® system to log information about BIG-IP system denial-of-service (DoS) events, and send the log messages to remote high-speed log servers.

***Important:*** *The BIG-IP Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure DoS Protection event logging. Additionally, for high-volume logging requirements, such as DoS, ensure that the BIG-IP system sends the event logs to a remote log server.*

This illustration shows the association of the configuration objects for remote high-speed logging of DoS Protection events.

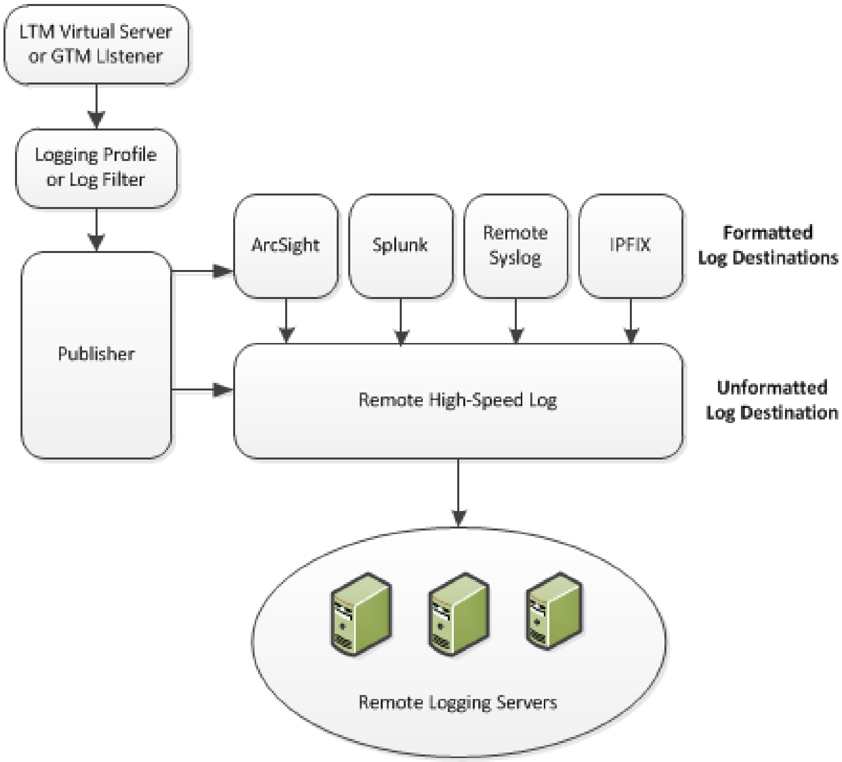
LTM Virtual Server or GTM Listener
Logging Profile or Log Filter
Publisher
ArcSight
Splunk
Remote Syslog
IPFIX
Formatted Log Destinations
Remote High-Speed Log
Unformatted Log Destination
Remote Logging Servers

**Figure 1: Association of remote high-speed logging configuration objects**

### Task summary

Perform these tasks to configure logging of DoS Protection events on the BIG-IP® system.

***Note:*** *Enabling logging impacts BIG-IP system performance.*

*Creating a pool of remote logging servers*
*Creating a remote high-speed log destination*
*Creating a formatted remote high-speed log destination*
*Creating a publisher*
*Creating a custom DoS Protection Logging profile*
*Logging DoS events on a virtual server*

*Disabling logging*

## About the configuration objects of DoS Protection event logging

When configuring remote high-speed logging of DoS Protection event logging, it is helpful to understand the objects you need to create and why, as described here:

| Object | Reason | Applies to |
|---|---|---|
| Pool of remote log servers | Create a pool of remote log servers to which the BIG-IP® system can send log messages. | Creating a pool of remote logging servers. |
| Destination (unformatted) | Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers. | Creating a remote high-speed log destination. |
| Destination (formatted) | If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination. | Creating a formatted remote high-speed log destination. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. | Creating a publisher. |
| DNS Logging profile | Create a custom DNS Logging profile to define the data you want the BIG-IP system to include in the DNS logs and associate a log publisher with the profile. | Creating a custom DoS Protecttion Logging profile. |
| LTM® virtual server | Associate a custom DNS profile with a virtual server to define how the BIG-IP system logs the DNS traffic that the virtual server processes. | Configuring an LTM virtual server for DoS Protection event logging. |

## Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
   - **DNS** > **Delivery** > **Load Balancing** > **Pools**
   - **Local Traffic** > **Pools**

   The Pool List screen opens.
2. Click **Create**.

The New Pool screen opens.

3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type a service number in the **Service Port** field, or select a service name from the list.

   ***Note:*** *Typical remote logging servers require port* `514`*.*

   c) Click **Add**.
5. Click **Finished**.

## Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

   ***Important:*** *If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the* ***Remote High-Speed Log*** *type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

   The BIG-IP system is configured to send an unformatted string of text to the log servers.
5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

## Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

   ***Important:*** *ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access*

*Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.*

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

   ***Important:*** *For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.*

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

## Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click **<<** to move the destination to the **Selected** list.

   ***Note:*** *If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

## Creating a custom DoS Protection Logging profile

Create a custom Logging profile to log DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
4. In the DNS DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for SIP or Application DoS Protection.
5. Click **Finished**.

Assign this custom DoS Protection Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

## Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

***Note:*** *You can disable and re-enable logging for a specific resource based on your network administration needs.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

# Implementation result

You now have an implementation in which the BIG-IP® system logs specific DoS Protection events and sends the logs to a specific location.

# Configuring High-Speed Remote Logging of DNS DoS Events

## Overview: Configuring DNS DoS Protection event logging

You can configure the BIG-IP® system to log information about BIG-IP system DNS denial-of-service (DoS) events, and send the log messages to remote high-speed log servers.

***Important:*** *The BIG-IP Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure DNS DoS Protection event logging. Additionally, for high volume logging requirements, such as DoS, ensure that the BIG-IP system sends the event logs to a remote log server.*

When configuring remote high-speed logging of DNS DoS Protection event logging, it is helpful to understand the objects you need to create and why, as described here:

| Object | Reason |
|---|---|
| Pool of remote log servers | Create a pool of remote log servers to which the BIG-IP system can send log messages. |
| Destination (unformatted) | Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers. |
| Destination (formatted) | If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. |
| Logging profile | Create a custom Logging profile to enable logging of user-specified data at a user-specified level, and associate a log publisher with the profile. |
| LTM® virtual server | Associate a custom Logging profile with a virtual server to define how the BIG-IP system logs security events on the traffic that the virtual server processes. |

This illustration shows the association of the configuration objects for remote high-speed logging of DoS Protection events.

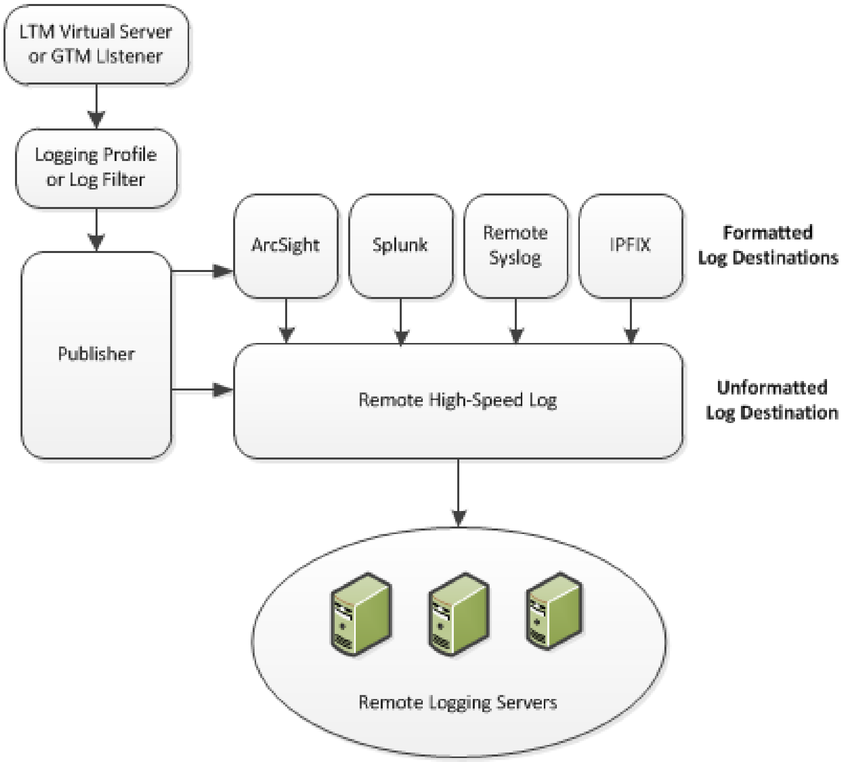
LTM Virtual Server or GTM Listener
Logging Profile or Log Filter
Publisher
ArcSight
Splunk
Remote Syslog
IPFIX
Formatted Log Destinations
Remote High-Speed Log
Unformatted Log Destination
Remote Logging Servers

**Figure 2: Association of remote high-speed logging configuration objects**

## Task summary

Perform these tasks to configure logging of DNS DoS Protection events on the BIG-IP® system.

***Note:*** *Enabling logging impacts BIG-IP system performance.*

*Creating a pool of remote logging servers*
*Creating a remote high-speed log destination*
*Creating a formatted remote high-speed log destination*
*Creating a publisher*
*Creating a custom DNS DoS Protection Logging profile*
*Logging DoS events on a virtual server*
*Disabling logging*

### Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
   - **DNS** > **Delivery** > **Load Balancing** > **Pools**
   - **Local Traffic** > **Pools**

   The Pool List screen opens.

2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type a service number in the **Service Port** field, or select a service name from the list.

   ***Note:*** *Typical remote logging servers require port* `514`*.*

   c) Click **Add**.
5. Click **Finished**.

## Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

   ***Important:*** *If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the* ***Remote High-Speed Log*** *type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

   The BIG-IP system is configured to send an unformatted string of text to the log servers.
5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

## Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

*Important: ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.*

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

   ***Important:*** *For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.*

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

## Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click **<<** to move the destination to the **Selected** list.

   ***Note:*** *If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

## Creating a custom DNS DoS Protection Logging profile

Create a custom Logging profile to log DNS DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a unique name for the logging profile.
4. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
5. In the DNS DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for SIP or Application DoS Protection.

6. Click **Finished**.

Assign this custom DNS DoS Protection Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

## Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

***Note:** You can disable and re-enable logging for a specific resource based on your network administration needs.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

# Implementation result

You now have an implementation in which the BIG-IP® system logs specific DoS Protection events and sends the logs to a specific location.

# About Logging DNS DoS Events to IPFIX Collectors

## Overview: Configuring IPFIX logging for DNS DoS

You can configure the BIG-IP® system to log information about DNS denial-of-service (DoS) events and send the log messages to remote IPFIX collectors.

IPFIX is a set of IETF standards. The BIG-IP system supports logging of DNS DoS events over the IPFIX protocol. IPFIX logs are raw, binary-encoded strings with their fields and field lengths defined by IPFIX templates. *IPFIX collectors* are external devices that can receive IPFIX templates and use them to interpret IPFIX logs.

The configuration process involves creating and connecting the following configuration objects:

| Object | Reason |
|---|---|
| Pool of IPFIX collectors | Create a pool of IPFIX collectors to which the BIG-IP system can send IPFIX log messages. |
| Destination | Create a log destination to format the logs in IPFIX templates, and forward the logs to the IPFIX collectors. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. |

### Task list

Perform these tasks to configure IPFIX logging of DNS DoS events on the BIG-IP system.

***Note:*** *Enabling IPFIX logging impacts BIG-IP system performance.*

*Assembling a pool of IPFIX collectors*
*Creating an IPFIX log destination*
*Creating a publisher*
*Creating a custom DNS DoS Protection Logging profile*

## Assembling a pool of IPFIX collectors

Before creating a pool of IPFIX collectors, gather the IP addresses of the collectors that you want to include in the pool. Ensure that the remote IPFIX collectors are configured to listen to and receive log messages from the BIG-IP® system.

These are the steps for creating a pool of IPFIX collectors. The BIG-IP system can send IPFIX log messages to this pool.

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each IPFIX collector that you want to include in the pool:

a) Type the collector's IP address in the **Address** field, or select a node address from the **Node List**.
b) Type a port number in the **Service Port** field.

   By default, IPFIX collectors listen on UDP or TCP port `4739` and Netflow V9 devices listen on port `2055`, though the port is configurable at each collector.
c) Click **Add**.

5. Click **Finished**.

## Creating an IPFIX log destination

A log destination of the **IPFIX** type specifies that log messages are sent to a pool of IPFIX collectors. Use these steps to create a log destination for IPFIX collectors.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **IPFIX**.
5. From the **Protocol** list, select **IPFIX** or **Netflow V9**, depending on the type of collectors you have in the pool.
6. From the **Pool Name** list, select an LTM® pool of IPFIX collectors.
7. From the **Transport Profile** list, select **TCP**, **UDP**, or any customized profile derived from TCP or UDP.
8. The **Template Retransmit Interval** is the time between transmissions of IPFIX templates to the pool of collectors. The BIG-IP system only retransmits its templates if the **Transport Profile** is a **UDP** profile.

   An *IPFIX template* defines the field types and byte lengths of the binary IPFIX log messages. The logging destination sends the template for a given log type (for example, NAT44 logs or customized logs from an iRule) before sending any of those logs, so that the IPFIX collector can read the logs of that type. The logging destination assigns a template ID to each template, and places the template ID into each log that uses that template.

   The log destination periodically retransmits all of its IPFIX templates over a UDP connection. The retransmissions are helpful for UDP connections, which are lossy.
9. The **Template Delete Delay** is the time that the BIG-IP device should pause between deleting an obsolete template and re-using its template ID. This feature is helpful for systems that can create custom IPFIX templates with iRules.
10. The **Server SSL Profile** applies Secure Socket Layer (SSL) or Transport Layer Security (TLS) to TCP connections. You can only choose an SSL profile if the **Transport Profile** is a **TCP** profile. Choose an SSL profile that is appropriate for the IPFIX collectors' SSL/TLS configuration.

    SSL or TLS requires extra processing and therefore slows the connection, so we only recommend this for sites where the connections to the IPFIX collectors have a potential security risk.
11. Click **Finished**.

## Creating a publisher

A publisher specifies where the BIG-IP® system sends log messages for IPFIX logs.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.

4. Use the **Log Destinations** setting to select an existing IPFIX destination (perhaps along with other destinations for your logs): click any destination name in the **Available** list, and click **<<** to move it to the **Selected** list.

   ***Important:*** *If you configure a log publisher to use multiple logging destinations, then, by default, all logging destinations must be available in order to log to each destination. Unless all logging destinations are available, no logging will occur. If you want to log to the available logging destinations when one or more destinations become unavailable, you must set the* `logpublisher.atomic` *db variable to* `false`*.*

5. Click **Finished**.

## Creating a custom DNS DoS Protection Logging profile

Create a custom Logging profile to log DNS DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a unique name for the logging profile.
4. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
5. In the DNS DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for SIP or Application DoS Protection.
6. Click **Finished**.

Assign this custom DNS DoS Protection Logging profile to a virtual server.

# Implementation result

Now you have an implementation in which the BIG-IP® system logs messages about DNS DoS events and sends the log messages to a pool of IPFIX collectors.

# Filtering DNS Packets

## About DNS protocol filtering

With a DNS security profile, you can filter DNS to allow or deny specific DNS query types, and to deny specific DNS OpCodes. The DNS security profile is attached to, and works with, a local traffic DNS profile to configure a range of DNS settings for a virtual server. Use DNS protocol filtering:

- To filter DNS query types or header OpCodes that are not necessary or relevant in your configuration, or that you do not want your DNS servers to handle.
- As a remediation tool to drop packets of a specific query type, if a DoS Protection Profile identifies anomalous DNS activity with that query type.

### Task list

*Filtering DNS traffic with a DNS security profile*
*Creating a custom DNS profile to firewall DNS traffic*
*Assigning a DNS profile to a virtual server*

## Filtering DNS traffic with a DNS security profile

The BIG-IP® system can allow or drop packets of specific DNS query types, or with specific opcodes, to prevent attacks or allow legitimate DNS traffic. You can use this to filter out header opcodes or query types that are not necessary on your system, or to respond to suspicious increases in packets of a certain type, as identified with the DNS security profile.

In this task, you create a DNS security profile and configure DNS security settings at the same time. However, you can also configure settings in a DNS security profile that already exists.

1. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **DNS**.
   The DNS Security Profiles list screen opens.
2. Click **Create**.
   The New Security Profile screen opens.
3. In the **Name** field, type the name for the profile.
4. From the **Query Type** list, select how to handle query types you add to the **Active** list.
   - Select **Inclusion** to allow packets with the DNS query types and header opcodes you add to the **Active** list, and drop all others.
   - Select **Exclusion** to deny packets with the DNS query types and header opcodes you add to the **Active** list, and allow all others.
5. In the **Query Type Filter** setting, move query types to filter for inclusion or exclusion from the **Available** list to the **Active** list.
6. In the **Header Opcode Exclusion** setting, move header types to filter for exclusion from the **Available** list to the **Active** list.

   ***Note:*** *Only the `query` opcode is available for header exclusion.*
7. Click **Finished** to save your changes.

Now you have configured the profile to include or exclude only specified DNS query types and header opcodes.

Specify this DNS security profile in a local traffic DNS profile attached to a virtual server.

## Creating a custom DNS profile to firewall DNS traffic

Ensure that you have a DNS security profile created before you configure this system DNS profile.

You can create a custom DNS profile to configure the BIG-IP® system firewall traffic through the system.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **DNS**.
   The DNS profile list screen opens.
2. Click **Create**.
   The New DNS Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. In the General Properties area, from the **Parent Profile** list, accept the default **dns** profile.
5. Select the **Custom** check box.
6. In the DNS Traffic area, from the **DNS Security** list, select **Enabled**.
7. In the DNS Traffic area, from the **DNS Security Profile Name** list, select the name of the DNS firewall profile.
8. Click **Finished**.

Assign the custom DNS profile to the virtual server that handles the DNS traffic that you want to firewall.

## Assigning a DNS profile to a virtual server

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. From the **DNS Profile** list, select the profile you want to assign to the virtual server.
5. Click **Update**.

The virtual server now handles DNS traffic.

# SSH Proxy Security

## Securing SSH traffic with the SSH Proxy

### Why use SSH proxy?

Network attacks are increasingly less visible, cloaked in SSL and SSH channels. The SSH Proxy feature provides a means to combat attacks in the SSH channel by providing visibility into SSH traffic and control over the commands that the users are executing in SSH channel. Administrators can control access on a per-user basis to SSH and the commands that a user can use in SSH.

### Challenges and problems that SSH proxy addresses

- Gives administrators visibility into user command activity in the SSH channel.
- Provides fine-grained control of SSH access commands on a per-user basis.
- Allows segmentation of access control for different users, allowing, for example, one user to download (but not upload) with SCP, while another user can upload and download with SCP. allowing SHELL access only to an administrator, and other examples.
- Control over SSH keep-alives that keep a session open indefinitely.

### Features of SSH Proxy

- Policy based SSH control capability
- Fine-grained control of SSH access on a per-user basis
- Visibility and control of SSH connection
- By controlling the SSH commands and session, datacenter admin can prevent advanced attacks from entering the datacenter.

### Current limits of SSH Proxy

- Supports SSH version 2.0 or above only
- SSH proxy is supported on a virtual server, not on a route domain or global context.
- SSH proxy auth key size is currently limited to 2K in this version.
- In this version, log profile configuration of SSH parameters is available only via tmsh.
- Elliptical Curve cypher (ECDHE) SSH keys are not supported for authentication in this version.

### Using SSH Proxy

You can use an SSH Proxy to secure SSH traffic on a virtual server, on a per-user basis. A working SSH proxy implementation requires

- An SSH proxy profile that defines actions for SSH channel commands
- A virtual server for the SSH server, configured for SSH traffic, and including the SSH proxy profile
- Authentication information for the SSH proxy

### Task summary

*Proxying SSH traffic with an SSH Proxy profile*
*Creating an SSH virtual server with SSH proxy security*
*Attaching an SSH proxy security profile to an existing virtual server*

## Proxying SSH traffic with an SSH Proxy profile

Configure an SSH proxy security profile to allow or deny SSH channel actions to specific users on a virtual server.

1. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **SSH Proxy**.
   The Protocol Security: Security Profiles: SSH Proxy screen opens.
2. Click **Create**.
   The New SSH Profile screen opens.
3. In the **Profile Name** field, type a unique name for the profile.
4. In the **Timeout** field, specify the timeout for an SSH session, in seconds.

   The timeout specifies how long the SSH connection will wait for a connection before returning an error. A setting of `0` means that the SSH connection attempt never times out.
5. To filter the list of SSH proxy permission rules, type the filter text in the **Filter Rules** field.

   ***Important:*** *The filter rules field is case sensitive.*
6. Edit an existing rule, or add a new rule.
   - To edit an existing rule, click the name of the rule. For example, click **Default Actions** to edit the default rule for a profile.
   - To add a new rule, click **Add New Rule**. A new line is added to the list of rules. Add a name to the rule to begin editing.
7. In the Users column, in the **add new user** field, type an SSH user name to which the rule applies, then click **Add**.

   ***Note:*** *You can not add users to the `Default Actions` rule.*
8. Configure the settings for each SSH channel action.
   - To allow the session to be set up for the SSH channel action, select **Allow**.
   - To deny an SSH channel action, and send a `command not accepted` message, select **Disallow**. Note that many SSH clients disconnect when this occurs.
   - To terminate an SSH connection by sending a reset message when a channel action is received, select **Terminate**.
9. To enable logging for an SSH action, select the **Log** check box.
10. When you finish editing
    - An existing rule, click **Done Editing**.
    - A new rule, click **Add Rule**.
11. When you are finished adding and editing rules, click **Commit Changes to System**.

The SSH proxy profile is saved to the system.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

### SSH channel actions

In an SSH proxy profile, you can configure whether to allow, disallow, or terminate SSH channel actions.

| Channel action | Description |
|---|---|
| Shell | Defines use of the `shell` command to open an SSH shell channel type. |
| Sub System | Defines the use of the `subsystem` command, to invoke remote commands that are defined on the server over the SSH tunnel. |

| Channel action | Description |
|---|---|
| SFTP Up | Defines the use of Secure File Transfer Protocol (`sftp`) to upload (`put`) files over the SSH tunnel. |
| SFTP Down | Defines the use of Secure File Transfer Protocol (`sftp`) to download (`get`) files over the SSH tunnel. |
| SCP Up | Defines the use of Secure Copy (`scp`) to copy files from a local directory to a remote directory over the SSH tunnel. |
| SCP Down | Defines the use of Secure Copy (`scp`) to copy files from a remote directory to a local directory over the SSH tunnel. |
| Rexec | Defines the use of `rexec` remote execution commands over the SSH tunnel. |
| Forward Local | Defines the use of the `-L` to do local port forwarding over the SSH tunnel. |
| Forward Remote | Defines the use of the `-R` to do remote port forwarding over the SSH tunnel. |
| Forward X11 | Defines the use of X11 forwarding over the SSH tunnel. |
| Agent | Defines the use of `ssh-agent` over the SSH tunnel. Agent forwarding specifies that the chain of SSH connections forwards key challenges back to the original agent, removing the need for passwords or private keys on intermediate machines. |

## Creating an SSH virtual server with SSH proxy security

Create an SSH virtual server to protect SSH connections with the SSH proxy.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

   ***Note:*** *The IP address you type must be available and not in the loopback network.*
5. In the **Service Port** field, type `22` or select **SSH** from the list.
6. From the **SSH Proxy Profile** list, select the SSH proxy profile to attach to the virtual server.
7. For the **Default Pool** setting, either select an existing pool from the list, or click the Create (+) button and create a new pool.

   The pool you create or select should contain your backend SSH server.
8. Click **Finished**.

The SSH virtual server appears in the Virtual Servers list.

## Attaching an SSH proxy security profile to an existing virtual server

You can add SSH proxy security to your SSH virtual server with SSH proxy profile.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.

2. In the **Name** column, click an SSH virtual server.
   The Properties screen for the virtual server opens.
3. From the **SSH Proxy Profile** list, select the SSH proxy profile to attach to the virtual server.
4. Click **Update** to save the changes.

You now have a virtual server configured so that the SSH proxy profile rules are applied to SSH traffic.

# Authenticating SSH proxy traffic

### What SSH authentication methods are supported?

SSH security supports public key authentication, password authentication, and keyboard-interactive authentication.

### Public key authentication

Public key authentication requires that both the SSH client and the SSH server must implement the security keys. With this method, each client must have a key pair generated using a supported encryption algorithm. When authentication occurs, the client sends a public key to the server. If the server finds the key in the list of allowed keys, the client encrypts data using the private key and sends the packet to the server with the public key.

### Password authentication

Password authentication is the simplest authentication method. The user specifies a username and password. This authentication method requires only one set of credentials for the user.

### Keyboard-interactive authentication

Keyboard-interactive authentication is a more complex form of password authentication, aimed specifically at the human operator as a client. During keyboard authentication prompts or questions are presented to the user. The user answers each prompt or question. The number and contents of the questions are virtually unlimited, so certain types of automated logins are also possible.

SSH client components support keyboard authentication via the `OnAuthenticationKeyboard` event. The client application should fill in the **Responses** parameter of the mentioned event with replies to questions contained in the **Prompts** parameter. Use `echo parameter` to specify whether the response is displayed on the screen, or masked. The number of responses must match the number of prompts or questions.

## Defining SSH proxy public key authentication

Before you configure public key authentication in the SSH proxy configuration, you must generate a public/private key pair. You can do this on the AFM system.

Configure tunnel keys for public key authentication to allow the SSH proxy to view tunnel trafffic.

1. On a system, type `ssh-keygen`.

   The system outputs:

   ```
   Generating public/private rsa key pair. Enter file in which to save the key
   (/root/.ssh/id_rsa):
   ```
2. Hit the **Enter** key to save the file.

   The system outputs:

   ```
   /root/.ssh/id_rsa already exists. Overwrite (y/n)?
   ```

3. Type `y` to save the file.

   The system prompts for a passphrase.

   `Enter passphrase (empty for no passphrase):`
4. Leave the passphrase and confirm passphrase fields blank, and hit `Enter`.

   The system outputs something like the following example. This output will be different on your system:

```
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
08:02:33:1a:8e:45:73:c0:eb:dc:fb:da:87:c5:2c:bf root@localhost.localdomain
The key's randomart image is:
+--[ RSA 2048]----+
|=o=..            |
|+*.o             |
|o....            |
|  .. . .         |
| o .  .oS        |
|  o . . +        |
|     . =         |
|    ... o        |
|    .oo.E.       |
+-----------------+
```

5. Copy the key from `id_rsa`.

   This is your private key, which you will add to the SSH proxy configuration.
6. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **SSH Proxy**. The Protocol Security: Security Profiles: SSH Proxy screen opens.
7. Click the name of the SSH proxy profile to edit. The SSH Profile screen opens.
8. Click the **Key Management** tab.
9. Click **Add New Auth Info**.
10. In the **Edit Auth Info Name** field, type a name for the authentication info settings.
    - To edit an existing rule, click the name of the rule. For example, click **Default Actions** to edit the default rule for a profile.
    - To add a new rule, click **Add New Rule**. A new line is added to the list of rules. Add a name to the rule to begin editing.
11. In the **Proxy Client Auth Private Key** field, paste the private key you have generated.

    You do not need to add the public key in the **Proxy Client Auth Public Key** field. This key is automatically generated.
12. In the **Proxy Server Auth Private Key** field, paste the private key of the client that will connect to the SSH proxy.
13. Click **Add**.
14. Click **Commit Changes to System**.
15. On the SSH client system, generate a private/public key pair with the command `ssh-keygen`.

    The system outputs:

```
Generating public/private rsa key pair.
Enter file in which to save the key (/home/user1/.ssh/id_rsa):
```

16. Click **Enter** or specify a different file location.
17. Type and confirm a passhphrase when prompted, or leave the fields blank to specify no passphrase.

The system outputs something like the following example. This output will be different on your system:

```
Your identification has been saved in /home/user1/.ssh/id_rsa.
Your public key has been saved in /home/user1/.ssh/id_rsa.pub.
The key fingerprint is:
25:26:7e:49:56:61:71:ca:23:ec:d1:49:6b:49:61:6b user1@Ubuntu-VM1
The key's randomart image is:
+--[ RSA 2048]----+
|          X+.    |
|       . O B     |
|      . O E      |
|     . * O .     |
|      . S        |
|       .         |
|                 |
|                 |
|                 |
+-----------------+
```

18. On the backend SSH server, modify the opensshd configuration file to look for public keys in multiple locations. In the opensshd file, uncomment the `AuthorizedKeysFile` line.
19. Specify a central authorized keys file by editing the AuthorizedKeysFile line as follows: `AuthorizedKeysFile %h/.ssh/authorized_keys /etc/ssh/authorized_keys`

    Note that you can specify your own path and filename for the authorized keys file on the SSH server.

    Restart the SSH daemon on the SSH server.
20. Copy the public key you created on the AFM system into the authorized keys file (for example `/etc/ssh/authorized_keys` with the following commands (the file location and name may differ, and the public key is an example only).

```
user1@Ubuntu-VM3:~$ cat /etc/ssh/authorized_keys
ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAQEAkCmU13s2/LVfm/eJ+HGesb8WeZ3A00iNX4S6ZDa7bOwb+f
jpr8rCwt4fWw8U7VwPaeqE35odBW7LhwQUXg5zL1KdxgguILVI2i/cDSkPKcaQKcUIvG+BrpYj
wky4T9tTKo2br+XQ92eWMh+xrVUwY4h2crpZxdng+YV+hUbqgJ+PHO4t0ozAYpgIul5C+2MTcN
zMuEYxbZqWdtNFtceAywu4CYZBwAZ3mCJbfW1wtFo6DG85tIo3LuaGXpA10jav1cC2szEo0OKT
0HUPJzYfSQiU/jHQv7Becwc9L8bOC6CxryTvx3Uq/Zf0ONQHhsyasIxg2wrVwzhbI1ctSyZgww==
root@localhost.localdomain
```

21. Copy the public key you created on the client system into the user authorized keys file (for example `/.ssh/authorized_keys` with the following commands (the file location and name may differ, and the public key is an example only).

```
user1@Ubuntu-VM3:~$ cat ~/.ssh/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDSMcf/wX3YZQAg+/RxbqXvXpIPVvnugCOYJm
uapYIze7Etc+192CB/zakmT3pKDyHHiVP1PwpP3jr99tY95llYg3p+A8nfv7+1UcwJYlS2EfYy
8qenb3Q4Mdtzrxr0AEjU/a4WXmGYd5h/ju5yRxQUt//q09PbxsEAf0qY05Tpax7R3rGl+15tf6
AI1a+poNGidfAAS1Pqc453qIXM1cp/PnOaKKzveQWBM2IIPenVxwdyX06Tn2OYBh4Rq4qUrt38
PyiYmKOYqQ/M4hD0R6/VLvF24i936uKfvBdkZcvePLGMpswQAteFzJA0JJjbWUIfvCYFCOLiFO
IATUGe9Nxl user1@Ubuntu-VM1
```

When the SSH server is added to a pool on a virtual server, and the SSH profile is attached to the virtual server, the client should now be able to make an SSH connection to the SSH server using the virtual server address.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

## Defining SSH proxy password or keyboard interactive authentication

Configure tunnel keys for password or keyboard interactive authentication to allow the SSH proxy to view tunnel trafffic.

1. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **SSH Proxy**.
   The Protocol Security: Security Profiles: SSH Proxy screen opens.
2. Click the name of the SSH proxy profile to edit.
   The SSH Profile screen opens.
3. Click the **Key Management** tab.
4. Click **Add New Auth Info**.
5. In the **Edit Auth Info Name** field, type a name for the authentication info settings.
   - To edit an existing rule, click the name of the rule. For example, click **Default Actions** to edit the default rule for a profile.
   - To add a new rule, click **Add New Rule**. A new line is added to the list of rules. Add a name to the rule to begin editing.
6. In the **Real Server Auth Public Key** field paste the public key from your backend server.

   The real server auth key must not be commented out in your sshd configuration. To make sure, on your backend SSH server, locate the file `etc/ssh/sshd_config`, and make sure the line `HostKey /etc/ssh/ssh_host_rsa_key` is not commented out.
7. In the **Proxy Server Auth Private Key** field, add a private key.

   ***Note:*** *The proxy server auth private key can be a newly-generated key. The Proxy Server Auth Public Key field can be left blank, as the public key is generated from the private key by the SSH proxy.*
8. Click **Add**.
9. When you are finished adding and editing rules, click **Commit Changes to System**.

The SSH proxy profile is saved to the system.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

### Authenticating SSH Proxy with the server private key

For this scenario, the SSH virtual server IP address to which you attach the SSH Proxy profile has the same IP address as the backend SSH server.

If your backend SSH server has the virtual server address, and clients connect directly to the backend SSH server address, using the SSH proxy in the middle, you can specify the private key from the backend server in the SSH proxy configuration.

1. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **SSH Proxy**.
   The Protocol Security: Security Profiles: SSH Proxy screen opens.
2. Click the name of the SSH proxy profile to edit.
   The SSH Profile screen opens.
3. Click the **Key Management** tab.
4. Click **Add New Auth Info**.
5. In the **Edit Auth Info Name** field, type a name for the authentication info settings.
   - To edit an existing rule, click the name of the rule. For example, click **Default Actions** to edit the default rule for a profile.
   - To add a new rule, click **Add New Rule**. A new line is added to the list of rules. Add a name to the rule to begin editing.
6. In the **Real Server Auth Public Key** field paste the public key from your backend server.

   The real server auth key must not be commented out in your sshd configuration. To make sure, on your backend SSH server, locate the file `etc/ssh/sshd_config`, and make sure the line `HostKey /etc/ssh/ssh_host_rsa_key` is not commented out.
7. Get the private key from the backend SSH server.

For example, on the SSH backend server, at the following prompt, the admin uses the specified command to get the SSH server private key:

```
admin@Ubuntu-VM3:~$ sudo cat /etc/ssh/ssh_host_rsa_key
```

The output of this command is the private key:

```
-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEAs4kusmrz6RbkYyz/Yc0YhAXFYCw8p6FqjTLsAqzkRJEog6lq
hUa8nRQhsumdVsMCbgzCMOYd7CLqrTqO/M3eqQWm16Y9EC1Mi7RsfNDnt7yJ6cMb
xtv2F/Smho6H5GrGSfrTqqDnuULHJ1GK+yMOghLqNnQVSGci/6NSMk7w3y/Pslzu
Lz82nZi9IL1dReen3kVbAhdB1K4VsHa0OgqSKV+mnLGNB2sq4Thj5lReKkc+3y8k
hyeV0M+SClyUTRyRG18drYldU7kJYc/IDjKjKdiIkqsig3FE5NjstHz2JDQFj5Yn
6uxqZWJIrfORC+VAoLR3+fea6omzkCVhQAMxxQIDAQABAoIBAHTx2cIMGr7s022q
hNtu3hY5MBz6E7RZV2+MCOGhPrtPFmXUt/cCYZ+r2luRApTeR7npg6CYdEs5X0Xh
S/xuGShd7xSvSz07VI33w2b2KMms/OSQ24oIA2ANU194fhoSVwEfajrNvsMVNWZu
HiqB5lRh/7/ik25rCAgemU79zraBdYC5FMzlMnl2TRrxlT0NjGtaniH+wpkZm1x6
S/evuvaJOYWhp8tarMQDcfPi0HNU4+agwRxrCcGNqei7nROTvXjVmsqxrcHGKCdF
4LdJyPJ6KYjtm0IcEYzKAFY3+haeX7ico3vRjSNSfMQwJbcJDMgoQpf44dFf9Jht
fEIuHUECgYEA4nwySeehTVftHxg3iv1Azy6FGT5q4KwXktA4G3fMjUmjjDQ2NAx0
VxlSEOU5sH2au8b19s/rOPsPjvYBYRAp8s+JD5BVVnfiJ/pcK8d+ws9gB65V0c3X
/ly3Gvz/He8B//CaaGCJOfzlmP4KKwfD3KzHw6+LJHEIdTHjQCMRnvUCgYEAyu60
WDEUpZf3dlOcfpTwaDdKtaHMOCQPH5LMD1vZAQdD1Gts20rEgDp8iKf/jXbo8/uA
HfR5jz89AgDygIlWO15an710W8DrhCBYvRP44X9KcQeZlqJswDiOc5tRApunrac1
fEPaJ7OTdLElyA7GuZlIJVkgCLfyDodohewb5ZECgYBfLVwgzLNvglTGrXGh+h2D
M4SBgEZ/1jIt40zA1k5izaBqKgLhSp6Vf7GKIhplPdOJt+njZ6rtDiySonUf6iAG
xwpNPRVvuf+TV1Xmm/Z8PZOYhr3P5lYvsZzNPaakWK2Zde4dkPv6H3oJGjEBtkir
8vwcEyhBDzNDtMxQRqyABQKBgQCmSsVuH4oTyFv4kruC3vnB7M1D2bpHpwTdkqW1
UEabGSD0SLODX9l2WncCZOh9PBvZExcBdPzH7cJIig4uVlxbeg45KD7ZkVVtiDQv
fNZNssmFpfyt+5uySKYzBet0f6kAHC0wD0oNjpIe5atYLQObw4fjUw11F4c7cKqu
U7TogQKBgFUu0Q5FLxaNNV1p9hNTCU+KDGN/kIe5K+8aJ08TpYhTSFSzgV2k47av
xCzTcSufjcZIpjNiGuwmT+spiwoPYqP+AdXKWWcxNfC4ahBfi7ROP6xSriCkzsYv
ZFhMHDfIjDAGDFmHI5v9Gcjxt+iFLdiDV9Pzv1XFDKd5yfJNfmGd
-----END RSA PRIVATE KEY-----
```

8. Paste the private key into the **Proxy Server Auth Private Key** field.
9. Click **Add**.
10. When you are finished adding and editing rules, click **Commit Changes to System**.

The SSH proxy profile is saved to the system.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

## Logging SSH Proxy actions

You can use a local logging profile and Proxy to secure SSH traffic on a virtual server, on a per-user basis. You do this by creating an SSH proxy profile and attaching it to a virtual server. You must also provide the server public and private keys for the encrypted traffic.

Task summary

*Create and associate a logging profile for SSH proxy events*

*Associating a logging profile with a security policy*

### Create and associate a logging profile for SSH proxy events

Create an SSH logging profile to specify the events that are logged for SSH proxy. Use a unique name for the log profile, and specify the log publisher you created for SSH Proxy events.

In tmsh, create the log profile and associate the log profile with SSH proxy events with the following command: `tmsh create sec log profile <`*`log_profile_name`*`> ssh-proxy add { ssh-`

```
log { log-publisher <log_publisher_name> allowed-channel-action enabled
disallowed-channel-action enabled ssh-timeout enabled non-ssh-traffic
enabled successful-server-side-auth enabled unsuccessful-client-side-auth
enabled unsuccessful-server-side-auth enabled }}
```

A logging profile named is created, which includes the SSH proxy events.

Associate this log profile with the SSH virtual server.

## Associating a logging profile with a security policy

A logging profile determines where events are logged and what details are included. By default, when you create a security policy, the system associates the Log Illegal Requests profile with the virtual server used by the policy. You can change which logging profile is associated with the security policy or assign a new one to the virtual server.

1. Click **Local Traffic** > **Virtual Servers**
2. Click the name of the virtual server used by the security policy.
   The system displays the general properties of the virtual server.
3. From the Security menu, choose Policies.
   The system displays the policy settings for the virtual server.
4. Ensure that the **Application Security Policy** setting is **Enabled**, and that **Policy** is set to the security policy you want.
5. For the **Log Profile** setting:
   a) Check that it is set to **Enabled**.
   b) From the **Available** list, select the profile to use for the security policy, and move it into the **Selected** list.

   You can assign only one local logging profile to a virtual server, but it can have multiple remote logging profiles.
6. Click **Update**.

Information related to traffic controlled by the security policy is logged using the logging profile or profiles specified in the virtual server.

# Example: Securing SSH traffic with the SSH Proxy

In this example, you create an SSH proxy configuration, create a virtual server for SSH traffic, and apply the SSH proxy to the virtual server. This example contains IP addresses and public and private keys that do not apply to your configuration, but are included for example purposes only.

In this configuration, password or keyboard interactive authentication is used, and the SSH proxy policy disallows SCP downloads and uploads, and terminates the tunnel connection on a REXEC command.

Task summary

*Example: proxying SSH traffic with an SSH Proxy profile*
*Example: defining SSH tunnel authentication keys in an SSH Proxy profile*
*Example: creating an SSH virtual server with SSH proxy security*

## Example: proxying SSH traffic with an SSH Proxy profile

Configure an SSH proxy security profile to allow or deny SSH channel actions to specific users on a virtual server. In this example, the proxy profile disallows SCP uploads and downloads, and terminates the channel on REXEC commands for the `root` user. All data entered in this screen is example data, and may not work on your system.

1. On the Main tab, click **Security** > **Protocol Security** > **Security Profiles** > **SSH Proxy**.
   The Protocol Security: Security Profiles: SSH Proxy screen opens.
2. Click **Create**.
   The New SSH Profile screen opens.
3. In the Profile Name field, type the name `ssh_no_scp_terminate_rexec`.
4. Click **Add New Rule** to add a rule for the profile.
5. In the Enter Rule Name field, type `root_rules` as the name for the rule.
6. In the Users column, in the **add new user** field, type `root`, and click **Add**.
7. From the **SCP Up** list, select **Disallow**.
8. From the **SCP Down** list, select **Disallow**.
9. From the **REXEC** list, select **Terminate**.
10. To enable logging for the SSH actions, select the **Log** check boxes.
11. Click **Add Rule**.
12. When you are finished adding and editing rules, click **Commit Changes to System**.

The SSH proxy profile is saved to the system.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

## Example: defining SSH tunnel authentication keys in an SSH Proxy profile

Working with the SSH proxy you defined earlier, add key management info to allow authentication.

1. In the same SSH proxy profile you previously created, click the **Key Management** tab.
2. Click **Add New Auth Info**.
3. In the **Edit Auth Info Name** field, type `root_auth` for the auth info name.
4. In the **Real Server Auth Public Key** field paste the public key from your backend server.

   The real server auth key must not be commented out in your sshd configuration. To make sure, on your backend SSH server, locate the file `etc/ssh/sshd_config`, and make sure the line `HostKey /etc/ssh/ssh_host_rsa_key` is not commented out.

   This is an example key.

```
AAAAB3NzaC1yc2EAAAADAQABAAABAQCziS6yavPpFuRjLP9hzRiEBcVgLDynoW
qNMuwCrOREkSiDqWqFRrydFCGy6Z1WwwJuDMIw5h3sIuqtOo78zd6pBabXpj0Q
LUyLtGx80Oe3vInpwxvG2/YX9KaGjofkasZJ+tOqoOe5QscnUYr7Iw6CEuo2dB
VIZyL/o1IyTvDfL8+yXO4vPzadmL0gvV1F56feRVsCF0HUrhWwdrQ6CpIpX6ac
sY0HayrhOGPmVF4qRz7fLySHJ5XQz5IKXJRNHJEbXx2tiV1TuQlhz8gOMqMp2I
iSqyKDcUTk2Oy0fPYkNAWPlifq7GplYkit85EL5UCgtHf595rqibOQJWFAAzHF
```

5. In the **Proxy Server Auth Private Key** field, add a private key.

   ***Tip:*** *The proxy server auth private key can be a newly-generated key. The Proxy Server Auth Public Key field can be left blank, as the public key is generated from the private key by the SSH proxy.*

   This is an example key.

```
-----BEGIN RSA PRIVATE KEY-----
MIIEogIBAAKCAQEAuncfRQM+yzcJW32r9DPKCzDP6cDhHbeTUlBOERUp27De+Vax
dojovwVi/tRiE/4tSbHViPF6BgS2Ar3W3tkxJySXLNczLkVV7WWkTEXCY+VrLB2I
BXA5YBWYVOjreZ/TYaJM+WxmxlDaFt1Rd2e7WVuegKjV1nVQyqdsW6vxY9GB93Pa
2v1VWUktInUAISwrT0nrE/rDkncAoKK2PUisP5u84HBIaT6QfXExNnreYHq8fWXk
0FOSOS8XlJugfumgdH9i9U5agAmG535f89O9eTDFUHSM2aaPkG+wbbLi2pxZiXR+
8n9graKVWTHl2zRvbIWB6wyfqae4zQoJVNgjdQIBIwKCAQBakaF5SrgZj8K3aO0e
11OBx0BqORzijF1/wJryWryPR0e675gGX8GBWmNIkwsRBm3EtXZYdUnlqoRKeXb+
hsAaU5nilGlQ/RsbiSPqiEh5qfI5/7cYlZg1+9xGf8LUrLcgyyyzqa5DEVP8eiBB
T6QkFo7QxwjHQEvQJW8lNkIL6JX5LP73hxvuQ3JwZizOR6cRmOyedIJHP0oNPsYS
w/nkpk15mL70S8asjWTF837vGcHS1M7TAko/r5KAd6FsbNWkk486iOhPtU2F3wJi
H9VO/Tvdl8MVSNzVzyjBjqigIU8nsMIvalYunM82w99+CA0RlWooZvEiPp5Qbv3v
TzOrAoGBAO5D8JAOuGCuWtU9cNJdtjWSeTP9ZsPYna6i4WHZYfOAGUlu5su4htY5
J26DygeHI6bm4Wew09t/ctq2Or60p6fIg/6XhEVrEkv6eZeCm7a+qajVVk77ZayT
cQdpbiDYrFI5rChTnzlSZ/QgWOFQ7klx66Qfd2nV/JAnU2K9J+CNAoGBAMhYJqdH
H7spzOTBXv6xWukRDld1/nsJC7mIIfjT2sVSLBAr5ZkyOdXwF5je6LNli3d7CpcS
tzv6YdMDEDsYNLlKFuMhgwmeCX0zwSzyfgRFFFXvIgaUUIW9RRjfLhuLFNzQ4/QB
BTmv98ltvjhorgsSonu0oydB3vHD4TJfstiJAoGBAKNhyYdajQ8YeMy8ap7hLHyB
sjJHXGkJkLJDzb9wfa5JNek2GppSpZo10eVhrxsa1p5VLNljT3Hw/kzUupFl7056
3irrjeZ1Tl/8Nh6/9b8jp4m23Bjm5qI5N5ANx9wCSkcC+bVAp7JHIrYHjWdNcDJc
vtbxAW0lBPUiR86tl6/rAoGBAJqNJSH1CdmGpWAC4uG8BE1k7c5w94N8AbsCnd01
t2UE4Cm7dprAWIB3Yqkg/KemGyGoD3vbPOUgPNX7DIVb0Oa1f17CFKEE4r+rlQVq
m7omqUmbN4FrGYu95NisKuIMNKpYAE6Ecb7Jk0OdzUF1Uw/bLOMWUfm2eMkiFB+L
pzlTAoGAQRAi+l/GHR3W6p9ahetItzPWn2tBJQnQiuM0ZFXEct41USPL4Sok8G28
Pu0C9Gf4u+bEi3BDFZMg7N6cnUYKeQjxTNmNtwgopjrGutXOM8ieiWp8oLG0zev/
pavXWCxdecuoyLtNeyTPR/GPpBqN3c5KjKnfsoid8mK59xfhic4=
-----END RSA PRIVATE KEY-----
```

6. Click **Add**.
7. When you are finished adding and editing rules, click **Commit Changes to System**.

The SSH proxy profile is saved to the system.

To use an SSH proxy profile with a virtual server, attach the profile to a virtual server on the Properties page, in the **Configuration (Basic)** settings.

## Example: creating an SSH virtual server with SSH proxy security

When you enable protocol security for an HTTP virtual server, the system scans any incoming HTTP traffic for vulnerabilities before the traffic reaches the HTTP servers.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type `ssh_root`.
4. In the **Destination Address/Mask** field, type the IP address in CIDR format.

   For example, `10.1.1.20`.
5. In the **Service Port** field, type `22` or select **SSH** from the list.
6. From the **SSH Proxy Profile** list, select `ssh_no_scp_terminate_rexec`.
7. For the **Default Pool** setting, either select an existing pool from the list, or click the Create (+) button and create a new pool.

   The pool you create or select should contain your backend SSH server.
8. Click **Finished**.

The SSH virtual server appears in the Virtual Servers list.

# Configuring High-Speed Remote Logging of SIP DoS Events

## Overview: Configuring SIP DoS Protection event logging

You can configure the BIG-IP® system to log information about BIG-IP system SIP protocol denial-of-service (DoS) events, and send the log messages to remote high-speed log servers.

***Important:*** *The Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure SIP DoS Protection event logging. Additionally, for high volume logging requirements, such as DoS, ensure that the BIG-IP system sends the event logs to a remote log server.*

When configuring remote high-speed logging of DoS Protection event logging, it is helpful to understand the objects you need to create and why, as described here:

| Object | Reason |
|---|---|
| Pool of remote log servers | Create a pool of remote log servers to which the BIG-IP system can send log messages. |
| Destination (unformatted) | Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers. |
| Destination (formatted) | If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. |
| Logging profile | Create a custom Logging profile to enable logging of user-specified data at a user-specified level, and associate a log publisher with the profile. |
| LTM® virtual server | Associate a custom Logging profile with a virtual server to define how the BIG-IP system logs security events on the traffic that the virtual server processes. |

This illustration shows the association of the configuration objects for remote high-speed logging of DoS Protection events.

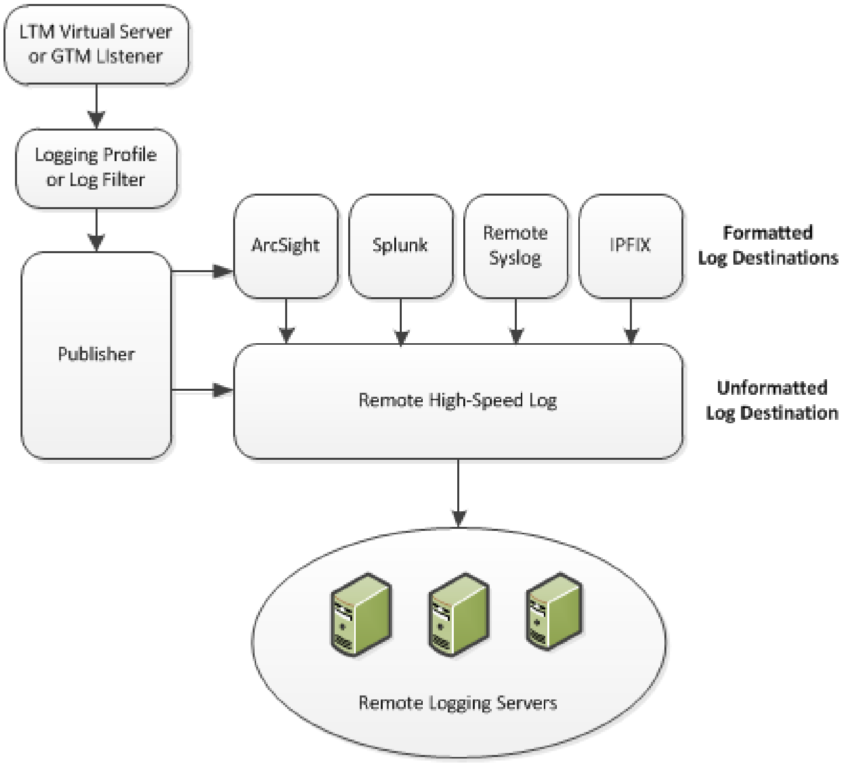
LTM Virtual Server or GTM Listener
Logging Profile or Log Filter
Publisher
ArcSight
Splunk
Remote Syslog
IPFIX
Formatted Log Destinations
Remote High-Speed Log
Unformatted Log Destination
Remote Logging Servers

**Figure 3: Association of remote high-speed logging configuration objects**

## Task summary

Perform these tasks to configure logging of SIP DoS Protection events on the BIG-IP® system.

***Note:*** *Enabling logging impacts BIG-IP system performance.*

*Creating a pool of remote logging servers*
*Creating a remote high-speed log destination*
*Creating a formatted remote high-speed log destination*
*Creating a publisher*
*Creating a custom SIP DoS Protection Logging profile*
*Logging DoS events on a virtual server*
*Disabling logging*

## Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
   - **DNS** > **Delivery** > **Load Balancing** > **Pools**
   - **Local Traffic** > **Pools**

   The Pool List screen opens.

2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type a service number in the **Service Port** field, or select a service name from the list.

      ***Note:*** *Typical remote logging servers require port* `514`*.*

   c) Click **Add**.
5. Click **Finished**.

## Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

   ***Important:*** *If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the* ***Remote High-Speed Log*** *type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

   The BIG-IP system is configured to send an unformatted string of text to the log servers.
5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

## Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

*__Important:__ ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.*

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

   *__Important:__ For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.*

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

## Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click **<<** to move the destination to the **Selected** list.

   *__Note:__ If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

## Creating a custom SIP DoS Protection Logging profile

Create a custom Logging profile to log SIP DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
4. In the SIP DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log SIP DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for DNS or Application DoS Protection.
5. Click **Finished**.

Assign this custom SIP DoS Protection Logging profile to a virtual server.

## Logging DoS events on a virtual server

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom logging profile to a virtual server when you want the system to log DoS protection events for the traffic the virtual server processes.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

## Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

***Note:*** *You can disable and re-enable logging for a specific resource based on your network administration needs.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

# Implementation result

You now have an implementation in which the BIG-IP® system logs specific DoS Protection events and sends the logs to a specific location.

# About Logging SIP DoS Events to IPFIX Collectors

## Overview: Configuring IPFIX logging for SIP DoS

You can configure the BIG-IP® system to log information about SIP denial-of-service (SIP DoS) events and send the log messages to remote IPFIX collectors.

IPFIX is a set of IETF standards. The BIG-IP system supports logging of SIP DoS events over the IPFIX protocol . IPFIX logs are raw, binary-encoded strings with their fields and field lengths defined by IPFIX templates. *IPFIX collectors* are external devices that can receive IPFIX templates and use them to interpret IPFIX logs.

The configuration process involves creating and connecting the following configuration objects:

| Object | Reason |
|---|---|
| Pool of IPFIX collectors | Create a pool of IPFIX collectors to which the BIG-IP system can send IPFIX log messages. |
| Destination | Create a log destination to format the logs in IPFIX templates, and forward the logs to the IPFIX collectors. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. |

### Task summary

Perform these tasks to configure IPFIX logging of SIP DoS events on the BIG-IP system.

***Note:*** *Enabling IPFIX logging impacts BIG-IP system performance.*

*Assembling a pool of IPFIX collectors*
*Creating an IPFIX log destination*
*Creating a publisher*
*Creating a custom DNS DoS Protection Logging profile*

## Assembling a pool of IPFIX collectors

Before creating a pool of IPFIX collectors, gather the IP addresses of the collectors that you want to include in the pool. Ensure that the remote IPFIX collectors are configured to listen to and receive log messages from the BIG-IP® system.

These are the steps for creating a pool of IPFIX collectors. The BIG-IP system can send IPFIX log messages to this pool.

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each IPFIX collector that you want to include in the pool:

a) Type the collector's IP address in the **Address** field, or select a node address from the **Node List**.
b) Type a port number in the **Service Port** field.

   By default, IPFIX collectors listen on UDP or TCP port `4739` and Netflow V9 devices listen on port `2055`, though the port is configurable at each collector.
c) Click **Add**.

5. Click **Finished**.

## Creating an IPFIX log destination

A log destination of the **IPFIX** type specifies that log messages are sent to a pool of IPFIX collectors. Use these steps to create a log destination for IPFIX collectors.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **IPFIX**.
5. From the **Protocol** list, select **IPFIX** or **Netflow V9**, depending on the type of collectors you have in the pool.
6. From the **Pool Name** list, select an LTM® pool of IPFIX collectors.
7. From the **Transport Profile** list, select **TCP**, **UDP**, or any customized profile derived from TCP or UDP.
8. The **Template Retransmit Interval** is the time between transmissions of IPFIX templates to the pool of collectors. The BIG-IP system only retransmits its templates if the **Transport Profile** is a **UDP** profile.

   An *IPFIX template* defines the field types and byte lengths of the binary IPFIX log messages. The logging destination sends the template for a given log type (for example, NAT44 logs or customized logs from an iRule) before sending any of those logs, so that the IPFIX collector can read the logs of that type. The logging destination assigns a template ID to each template, and places the template ID into each log that uses that template.

   The log destination periodically retransmits all of its IPFIX templates over a UDP connection. The retransmissions are helpful for UDP connections, which are lossy.
9. The **Template Delete Delay** is the time that the BIG-IP device should pause between deleting an obsolete template and re-using its template ID. This feature is helpful for systems that can create custom IPFIX templates with iRules.
10. The **Server SSL Profile** applies Secure Socket Layer (SSL) or Transport Layer Security (TLS) to TCP connections. You can only choose an SSL profile if the **Transport Profile** is a **TCP** profile. Choose an SSL profile that is appropriate for the IPFIX collectors' SSL/TLS configuration.

    SSL or TLS requires extra processing and therefore slows the connection, so we only recommend this for sites where the connections to the IPFIX collectors have a potential security risk.
11. Click **Finished**.

## Creating a publisher

A publisher specifies where the BIG-IP® system sends log messages for IPFIX logs.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.

4. Use the **Log Destinations** setting to select an existing IPFIX destination (perhaps along with other destinations for your logs): click any destination name in the **Available** list, and click **<<** to move it to the **Selected** list.

   ---

   ***Important:*** *If you configure a log publisher to use multiple logging destinations, then, by default, all logging destinations must be available in order to log to each destination. Unless all logging destinations are available, no logging will occur. If you want to log to the available logging destinations when one or more destinations become unavailable, you must set the* `logpublisher.atomic` *db variable to* `false`*.*

   ---

5. Click **Finished**.

## Creating a custom DNS DoS Protection Logging profile

Create a custom Logging profile to log DNS DoS Protection events and send the log messages to a specific location.

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a unique name for the logging profile.
4. In the Logging Profile Properties, select the **DoS Protection** check box.
   The DoS Protection tab opens.
5. In the DNS DoS Protection area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS DoS events.

   You can specify publishers for other DoS types in the same profile, for example, for SIP or Application DoS Protection.
6. Click **Finished**.

Assign this custom DNS DoS Protection Logging profile to a virtual server.

# Implementation result

Now you have an implementation in which the BIG-IP® system logs messages about SIP DoS events and sends the log messages to a pool of IPFIX collectors.

# Configuring High-Speed Remote Logging of Protocol Security Events

## Overview: Configuring Remote Protocol Security Event Logging

You can configure the BIG-IP® system to log information about BIG-IP system Protocol Security events and send the log messages to remote high-speed log servers.

***Important:*** *The Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure Protocol Security event logging.*

This illustration shows the association of the configuration objects for remote high-speed logging.

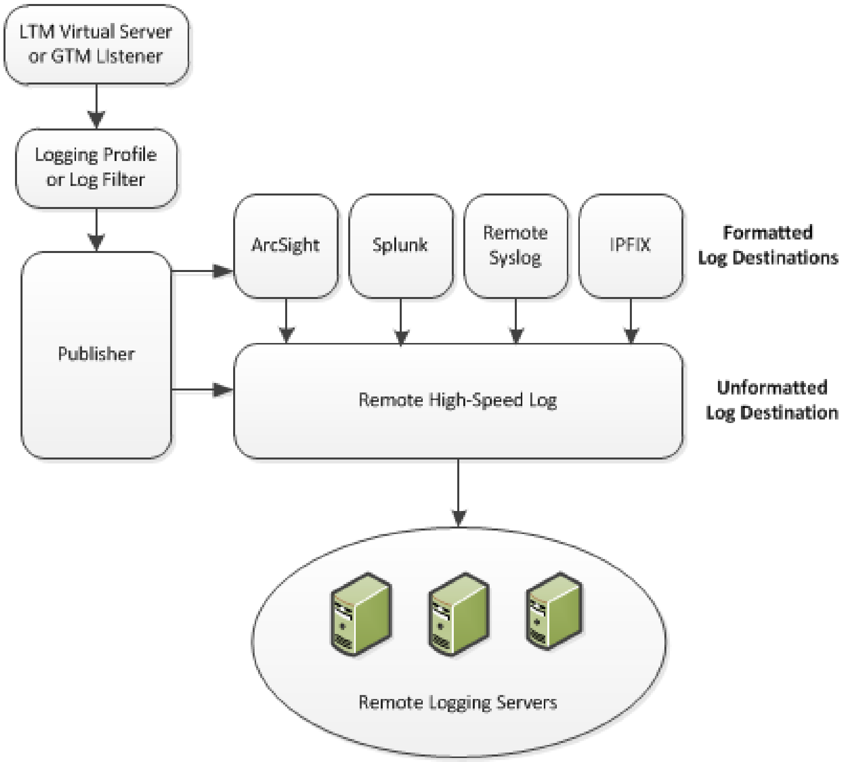
LTM Virtual Server or GTM Listener
Logging Profile or Log Filter
Publisher
ArcSight
Splunk
Remote Syslog
IPFIX
Formatted Log Destinations
Remote High-Speed Log
Unformatted Log Destination
Remote Logging Servers

**Figure 4: Association of remote high-speed logging configuration objects**

### Task summary

Perform these tasks to configure Protocol Security event logging on the BIG-IP® system.

***Note:*** *Enabling remote high-speed logging impacts BIG-IP system performance.*

*Creating a pool of remote logging servers*
*Creating a remote high-speed log destination*
*Creating a formatted remote high-speed log destination*
*Creating a publisher*
*Creating a custom Protocol Security Logging profile*
*Configuring a virtual server for Protocol Security event logging*

*Disabling logging*

## About the configuration objects of remote protocol security event logging

When configuring remote high-speed logging of Protocol Security events, it is helpful to understand the objects you need to create and why, as described here:

| Object | Reason | Applies to |
|---|---|---|
| Pool of remote log servers | Create a pool of remote log servers to which the BIG-IP® system can send log messages. | Creating a pool of remote logging servers. |
| Destination (unformatted) | Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers. | Creating a remote high-speed log destination. |
| Destination (formatted) | If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination. | Creating a formatted remote high-speed log destination. |
| Publisher | Create a log publisher to send logs to a set of specified log destinations. | Creating a publisher. |
| DNS Logging profile | Create a custom DNS Logging profile to define the data you want the BIG-IP system to include in the DNS logs and associate a log publisher with the profile. | Creating a custom Protocol Security Logging profile. |
| LTM® virtual server | Associate a custom DNS profile with a virtual server to define how the BIG-IP system logs the DNS traffic that the virtual server processes. | Configuring a virtual server for Protocol Security event logging. |

## Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
   - **DNS** > **Delivery** > **Load Balancing** > **Pools**
   - **Local Traffic** > **Pools**

   The Pool List screen opens.
2. Click **Create**.

The New Pool screen opens.

3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type a service number in the **Service Port** field, or select a service name from the list.

   ***Note:*** *Typical remote logging servers require port* `514`*.*

   c) Click **Add**.
5. Click **Finished**.

## Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

   ***Important:*** *If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the* ***Remote High-Speed Log*** *type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

   The BIG-IP system is configured to send an unformatted string of text to the log servers.
5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

## Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

   ***Important:*** *ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access*

*Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.*

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

   ***Important:*** *For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.*

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

## Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click **<<** to move the destination to the **Selected** list.

   ***Note:*** *If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

## Creating a custom Protocol Security Logging profile

Create a logging profile to log Protocol Security events for the traffic handled by the virtual server to which the profile is assigned.

***Note:*** *You can configure logging profiles for HTTP and DNS security events on Advanced Firewall Manager™, and FTP and SMTP security events on Application Security Manager™.*

1. On the Main tab, click **Security** > **Event Logs** > **Logging Profiles**.
   The Logging Profiles list screen opens.
2. Click **Create**.
   The Create New Logging Profile screen opens.
3. Select the **Protocol Security** check box.
4. In the HTTP, FTP, and SMTP Security area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log HTTP, FTP, and SMTP Security events.
5. In the DNS Security area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS Security events.
6. Select the **Log Dropped Requests** check box, to enable the BIG-IP system to log dropped DNS requests.

7. Select the **Log Filtered Dropped Requests** check box, to enable the BIG-IP system to log DNS requests dropped due to DNS query/header-opcode filtering.

   ***Note:*** *The system does not log DNS requests that are dropped due to errors in the way the system processes DNS packets.*

8. Select the **Log Malformed Requests** check box to enable the BIG-IP system to log malformed DNS requests.
9. Select the **Log Rejected Requests** check box to enable the BIG-IP system to log rejected DNS requests.
10. Select the **Log Malicious Requests** check box to enable the BIG-IP system to log malicious DNS requests.
11. From the **Storage Format** list, select how the BIG-IP system formats the log.

| Option | Description |
|---|---|
| **None** | Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example:<br>`"management_ip_address","bigip_hostname","context_type", "context_name","src_ip","dest_ip","src_port", "dest_port","vlan","protocol","route_domain", "acl_rule_name","action","drop_reason` |
| **Field-List** | Allows you to:<br>• Select, from a list, the fields to be included in the log.<br>• Specify the order the fields display in the log.<br>• Specify the delimiter that separates the content in the log. The default delimiter is the comma character. |
| **User-Defined** | Allows you to:<br>• Select, from a list, the fields to be included in the log.<br>• Cut and paste, in a string of text, the order the fields display in the log. |

12. Click **Finished**.

Assign this custom Protocol Security Logging profile to a virtual server.

## Configuring a virtual server for Protocol Security event logging

Ensure that at least one Log Publisher exists on the BIG-IP® system.

Assign a custom Protocol Security Logging profile to a virtual server when you want the BIG-IP system to log Protocol Security events on the traffic the virtual server processes.

***Note:*** *This task applies only to systems provisioned at a minimum level (or higher) for* ***Local Traffic (LTM)****. You can check the provisioning level on the* ***System*** *>* ***Resource Provisioning*** *screen.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

## Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

***Note:** You can disable and re-enable logging for a specific resource based on your network administration needs.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security** > **Policies**.
   The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

# Implementation result

You now have an implementation in which the BIG-IP® system logs specific Protocol Security events and sends the logs to a specific location.

# IPFIX Templates for AFM DNS Events

## Overview: IPFIX Templates for AFM DNS Events

The IP Flow Information Export (IPFIX) Protocol is a logging mechanism for IP events. This appendix defines the IPFIX Information Elements (IEs) and Templates used to log F5's Application Firewall Manager (AFM) DNS events. An *IE* is the smallest form of useful information in an IPFIX log message, such as an IP address or a timestamp for the event. An *IPFIX template* is an ordered collection of specific IEs used to record one IP event, such as the denial of a DNS query.

## About IPFIX Information Elements for AFM DNS events

Information Elements (IEs) are individual fields in an IPFIX template. An IPFIX template describes a single Advanced Firewall Manager™ (AFM™) DNS event.

### IANA-defined IPFIX Information Elements

IANA maintains a list of standard IPFIX Information Elements (IEs), each with a unique Element Identifier. The F5® AFM™ DNS IPFIX implementation uses a subset of these IEs to publish AFM DNS events. This subset is summarized in the table.

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| destinationIPv4Address | 12 | 4 |
| destinationIPv6Address | 28 | 16 |
| destinationTransportPort | 11 | 2 |
| ingressVRFID | 234 | 4 |
| observationTimeMilliseconds | 323 | 8 |
| sourceIPv4Address | 8 | 4 |
| sourceIPv6Address | 27 | 16 |
| sourceTransportPort | 7 | 2 |

### IPFIX enterprise Information Elements

IPFIX provides for enterprises to define their own Information Elements. F5® currently uses the following non-standard IEs for AFM™ DNS events:

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| action | 12276 - 39 | Variable |
| attackEvent | 12276 - 41 | Variable |
| attackId | 12276 - 20 | 4 |
| attackName | 12276 - 21 | Variable |
| bigipHostName | 12276 - 10 | Variable |

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| bigipMgmtIPv4Address | 12276 - 5 | 4 |
| bigipMgmtIPv6Address | 12276 - 6 | 16 |
| contextName | 12276 - 9 | Variable |
| deviceProduct | 12276 - 12 | Variable |
| deviceVendor | 12276 - 11 | Variable |
| deviceVersion | 12276 - 13 | Variable |
| dnsQueryType | 12276 - 8 | Variable |
| errdefsMsgNo | 12276 - 4 | 4 |
| flowId | 12276 - 3 | 8 |
| ipfixMsgNo | 12276 - 16 | 4 |
| messageSeverity | 12276 - 1 | 1 |
| msgName | 12276 - 14 | Variable |
| packetsDropped | 12276 - 23 | 4 |
| packetsReceived | 12276 - 22 | 4 |
| partitionName | 12276 - 2 | Variable |
| queryName | 12276 - 7 | Variable |
| vlanName | 12276 - 15 | Variable |

***Note:*** *IPFIX, unlike NetFlow v9, supports variable-length IEs, where the length is encoded within the field in the Data Record. NetFlow v9 collectors (and their variants) cannot correctly process variable-length IEs, so they are omitted from logs sent to those collector types.*

## About individual IPFIX Templates for each event

This section enumerates the IPFIX templates used by F5 to publish AFM DNS Events.

### IPFIX template for DNS security

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| action | 12276 - 39 | Variable | This IE is omitted for NetFlow v9. |
| bigipHostName | 12276 - 10 | Variable | This IE is omitted for NetFlow v9. |
| bigipMgmtIPv4Address | 12276 - 5 | 4 | |
| bigipMgmtIPv6Address | 12276 - 6 | 16 | |
| contextName | 12276 - 9 | Variable | This IE is omitted for NetFlow v9. |
| observationTimeMilliseconds | 323 | 8 | |
| destinationIPv4Address | 12 | 4 | |
| destinationIPv6Address | 28 | 16 | |

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| destinationTransportPort | 11 | 2 | |
| deviceProduct | 12276 - 12 | Variable | This IE is omitted for NetFlow v9. |
| deviceVendor | 12276 - 11 | Variable | This IE is omitted for NetFlow v9. |
| deviceVersion | 12276 - 13 | Variable | This IE is omitted for NetFlow v9. |
| queryName | 12276 - 7 | Variable | This IE is omitted for NetFlow v9. |
| dnsQueryType | 12276 - 8 | Variable | This IE is omitted for NetFlow v9. |
| errdefsMsgNo | 12276 - 4 | 4 | |
| flowId | 12276 - 3 | 8 | |
| ipfixMsgNo | 12276 - 16 | 4 | |
| messageSeverity | 12276 - 1 | 1 | |
| partitionName | 12276 - 2 | Variable | This IE is omitted for NetFlow v9. |
| ingressVRFID | 234 | 4 | |
| sourceIPv4Address | 8 | 4 | |
| sourceIPv6Address | 27 | 16 | |
| sourceTransportPort | 7 | 2 | |
| vlanName | 12276 - 15 | Variable | This IE is omitted for NetFlow v9. |
| msgName | 12276 - 14 | Variable | This IE is omitted for NetFlow v9. |

## IPFIX template for DNS DoS

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| action | 12276 - 39 | Variable | This IE is omitted for NetFlow v9. |
| attackEvent | 12276 - 41 | Variable | This IE is omitted for NetFlow v9. |
| attackId | 12276 - 20 | 4 | |
| attackName | 12276 - 21 | Variable | This IE is omitted for NetFlow v9. |
| bigipHostName | 12276 - 10 | Variable | This IE is omitted for NetFlow v9. |
| bigipMgmtIPv4Address | 12276 - 5 | 4 | |
| bigipMgmtIPv6Address | 12276 - 6 | 16 | |
| contextName | 12276 - 9 | Variable | This IE is omitted for NetFlow v9. |
| observationTimeMilliseconds | 323 | 8 | |
| destinationIPv4Address | 12 | 4 | |
| destinationIPv6Address | 28 | 16 | |
| destinationTransportPort | 11 | 2 | |
| deviceProduct | 12276 - 12 | Variable | This IE is omitted for NetFlow v9. |
| deviceVendor | 12276 - 11 | Variable | This IE is omitted for NetFlow v9. |
| deviceVersion | 12276 - 13 | Variable | This IE is omitted for NetFlow v9. |

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| queryName | 12276 - 7 | Variable | This IE is omitted for NetFlow v9. |
| dnsQueryType | 12276 - 8 | Variable | This IE is omitted for NetFlow v9. |
| errdefsMsgNo | 12276 - 4 | 4 | |
| flowId | 12276 - 3 | 8 | |
| ipfixMsgNo | 12276 - 16 | 4 | |
| messageSeverity | 12276 - 1 | 1 | |
| partitionName | 12276 - 2 | Variable | This IE is omitted for NetFlow v9. |
| ingressVRFID | 234 | 4 | |
| sourceIPv4Address | 8 | 4 | |
| sourceIPv6Address | 27 | 16 | |
| sourceTransportPort | 7 | 2 | |
| vlanName | 12276 - 15 | Variable | This IE is omitted for NetFlow v9. |
| msgName | 12276 - 14 | Variable | This IE is omitted for NetFlow v9. |
| packetsDropped | 12276 - 23 | 4 | |
| packetsReceived | 12276 - 22 | 4 | |

# IPFIX Templates for AFM SIP Events

## Overview: IPFIX Templates for AFM SIP Events

The IP Flow Information Export (IPFIX) Protocol is a logging mechanism for IP events. This appendix defines the IPFIX Information Elements (IEs) and Templates used to log F5's Application Firewall Manager (AFM) events related to the Session Initiation Protocol (SIP). An *IE* is the smallest form of useful information in an IPFIX log message, such as an IP address or a timestamp for the event. An *IPFIX template* is an ordered collection of specific IEs used to record one IP event, such as the acceptance of a SIP session.

## About IPFIX Information Elements for AFM SIP events

Information Elements (IEs) are individual fields in an IPFIX template. An IPFIX template describes a single Advanced Firewall Manager™ (AFM™) SIP event.

### IANA-defined IPFIX information elements

IANA maintains a list of standard IPFIX Information Elements (IEs), each with a unique Element Identifier. The F5® AFM™ SIP implementation uses a subset of these IEs to publish AFM SIP events. This subset is summarized in the table.

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| destinationIPv4Address | 12 | 4 |
| destinationIPv6Address | 28 | 16 |
| destinationTransportPort | 11 | 2 |
| ingressVRFID | 234 | 4 |
| observationTimeMilliseconds | 323 | 8 |
| sourceIPv4Address | 8 | 4 |
| sourceIPv6Address | 27 | 16 |
| sourceTransportPort | 7 | 2 |

### IPFIX enterprise Information Elements

IPFIX provides for enterprises to define their own Information Elements. F5® currently uses the following non-standard IEs for AFM™ events:

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| action | 12276 - 39 | Variable |
| attackEvent | 12276 - 41 | Variable |
| attackId | 12276 - 20 | 4 |
| attackName | 12276 - 21 | Variable |

| Information Element (IE) | ID | Size (Bytes) |
|---|---|---|
| bigipHostName | 12276 - 10 | Variable |
| bigipMgmtIPv4Address | 12276 - 5 | 4 |
| bigipMgmtIPv6Address | 12276 - 6 | 16 |
| contextName | 12276 - 9 | Variable |
| deviceProduct | 12276 - 12 | Variable |
| deviceVendor | 12276 - 11 | Variable |
| deviceVersion | 12276 - 13 | Variable |
| errdefsMsgNo | 12276 - 4 | 4 |
| flowId | 12276 - 3 | 8 |
| ipfixMsgNo | 12276 - 16 | 4 |
| messageSeverity | 12276 - 1 | 1 |
| msgName | 12276 - 14 | Variable |
| packetsDropped | 12276 - 23 | 4 |
| packetsReceived | 12276 - 22 | 4 |
| partitionName | 12276 - 2 | Variable |
| sipCallee | 12276 - 19 | Variable |
| sipCaller | 12276 - 18 | Variable |
| sipMethodName | 12276 - 17 | Variable |
| vlanName | 12276 - 15 | Variable |

***Note:** IPFIX, unlike NetFlow v9, supports variable-length IEs, where the length is encoded within the field in the Data Record. NetFlow v9 collectors (and their variants) cannot correctly process variable-length IEs, so they are omitted from logs sent to those collector types.*

## About individual IPFIX Templates for each event

This section enumerates the IPFIX templates used by F5 to publish AFM SIP Events.

### IPFIX template for SIP security

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| action | 12276 - 39 | Variable | This IE is omitted for NetFlow v9. |
| bigipHostName | 12276 - 10 | Variable | This IE is omitted for NetFlow v9. |
| bigipMgmtIPv4Address | 12276 - 5 | 4 | |
| bigipMgmtIPv6Address | 12276 - 6 | 16 | |
| contextName | 12276 - 9 | Variable | This IE is omitted for NetFlow v9. |
| observationTimeMilliseconds | 323 | 8 | |

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| destinationIPv4Address | 12 | 4 | |
| destinationIPv6Address | 28 | 16 | |
| destinationTransportPort | 11 | 2 | |
| deviceProduct | 12276 - 12 | Variable | This IE is omitted for NetFlow v9. |
| deviceVendor | 12276 - 11 | Variable | This IE is omitted for NetFlow v9. |
| deviceVersion | 12276 - 13 | Variable | This IE is omitted for NetFlow v9. |
| errdefsMsgNo | 12276 - 4 | 4 | |
| flowId | 12276 - 3 | 8 | |
| ipfixMsgNo | 12276 - 16 | 4 | |
| messageSeverity | 12276 - 1 | 1 | |
| partitionName | 12276 - 2 | Variable | This IE is omitted for NetFlow v9. |
| ingressVRFID | 234 | 4 | |
| sipCallee | 12276 - 19 | Variable | This IE is omitted for NetFlow v9. |
| sipCaller | 12276 - 18 | Variable | This IE is omitted for NetFlow v9. |
| sipMethodName | 12276 - 17 | Variable | This IE is omitted for NetFlow v9. |
| sourceIPv4Address | 8 | 4 | |
| sourceIPv6Address | 27 | 16 | |
| sourceTransportPort | 7 | 2 | |
| vlanName | 12276 - 15 | Variable | This IE is omitted for NetFlow v9. |
| msgName | 12276 - 14 | Variable | This IE is omitted for NetFlow v9. |

## IPFIX template for SIP DoS

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| action | 12276 - 39 | Variable | This IE is omitted for NetFlow v9. |
| attackEvent | 12276 - 41 | Variable | This IE is omitted for NetFlow v9. |
| attackId | 12276 - 20 | 4 | |
| attackName | 12276 - 21 | Variable | This IE is omitted for NetFlow v9. |
| bigipHostName | 12276 - 10 | Variable | This IE is omitted for NetFlow v9. |
| bigipMgmtIPv4Address | 12276 - 5 | 4 | |
| bigipMgmtIPv6Address | 12276 - 6 | 16 | |
| contextName | 12276 - 9 | Variable | This IE is omitted for NetFlow v9. |
| observationTimeMilliseconds | 323 | 8 | |
| destinationIPv4Address | 12 | 4 | |
| destinationIPv6Address | 28 | 16 | |
| destinationTransportPort | 11 | 2 | |

| Information Element (IE) | ID | Size (Bytes) | Notes |
|---|---|---|---|
| deviceProduct | 12276 - 12 | Variable | This IE is omitted for NetFlow v9. |
| deviceVendor | 12276 - 11 | Variable | This IE is omitted for NetFlow v9. |
| deviceVersion | 12276 - 13 | Variable | This IE is omitted for NetFlow v9. |
| errdefsMsgNo | 12276 - 4 | 4 | |
| flowId | 12276 - 3 | 8 | |
| ipfixMsgNo | 12276 - 16 | 4 | |
| messageSeverity | 12276 - 1 | 1 | |
| partitionName | 12276 - 2 | Variable | This IE is omitted for NetFlow v9. |
| ingressVRFID | 234 | 4 | |
| sipCallee | 12276 - 19 | Variable | This IE is omitted for NetFlow v9. |
| sipCaller | 12276 - 18 | Variable | This IE is omitted for NetFlow v9. |
| sipMethodName | 12276 - 17 | Variable | This IE is omitted for NetFlow v9. |
| sourceIPv4Address | 8 | 4 | |
| sourceIPv6Address | 27 | 16 | |
| sourceTransportPort | 7 | 2 | |
| vlanName | 12276 - 15 | Variable | This IE is omitted for NetFlow v9. |
| msgName | 12276 - 14 | Variable | This IE is omitted for NetFlow v9. |
| packetsDropped | 12276 - 23 | 4 | |
| packetsReceived | 12276 - 22 | 4 | |

# Legal Notices

## Legal notices

### Publication Date

This document was published on November 13, 2017.

### Publication Number

MAN-0440-07

### Copyright

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

### Trademarks

For a current list of F5 trademarks and service marks, see *http://www.f5.com/about/guidelines-policies/trademarks*.

All other product and company names herein may be trademarks of their respective owners.

### Patents

This product may be protected by one or more patents indicated at: *https://f5.com/about-us/policies/patents*.

### Link Controller Availability

This product is not currently available in the U.S.

### Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

### RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

### FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

### Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

### Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

# Index