

Characterizing and Tracing Packet Floods Using Cisco Routers

Introduction

Denial of service (DoS) attacks are common on the Internet. The first step in responding to such an attack is to find out exactly what sort of attack it is. Many of the commonly used DoS attacks are based on high-bandwidth packet floods, or on other repetitive streams of packets.

The packets in many DoS attack streams can be isolated by matching them against Cisco IOS software access list entries. This is obviously valuable for filtering out attacks, but is also useful for characterizing unknown attacks, and for tracing "spoofed" packet streams back to their real sources.

Cisco router features such as debug logging and IP accounting can sometimes be used for similar purposes, especially with new or unusual attacks. However, with recent versions of Cisco IOS software, access lists and access list logging are the premiere features for characterizing and tracing common attacks.

The Most Common DoS Attacks

A wide variety of DoS attacks are possible. Even if we ignore attacks that use software bugs to shut down systems with relatively little traffic, the fact remains that any IP packet that can be sent across the network can be used to execute a flooding DoS attack. When you are under attack, you must always consider the possibility that what you're seeing is something that does not fall into the usual categories.

Subject to that caveat, however, it's also good to remember that many attacks are similar. Attackers choose common exploits because they are particularly effective, particularly hard to trace, or because tools are available. Many DoS attackers lack the skill or motivation to create their own tools, and use programs found on the Internet; these tools tend to fall in and out of fashion.

At the time of this writing, in July 1999, most customer requests for Cisco assistance involve the "smurf" attack. This attack has two victims: an "ultimate target" and a "reflector." The attacker sends a stimulus stream of ICMP echo requests ("pings") to the broadcast address of the reflector subnet. The source addresses of these packets are falsified to be the address of the ultimate target. For each packet sent by the attacker, many hosts on the reflector subnet will respond, flooding the ultimate target and wasting bandwidth for both victims.

A similar attack, called "fraggle," uses directed broadcasts in the same way, but uses UDP echo requests instead of ICMP echo requests. Fraggle usually achieves a smaller amplification factor than smurf, and is much less popular.

Smurf attacks are usually noticed because a network link becomes overloaded. A complete description of these attacks, and of defense measures, is on the World Wide Web at <http://users.quadrunner.com/chuegen/smurf.cgi>.

Another common attack is the SYN flood, in which a target machine is flooded with TCP connection requests. The source addresses and source TCP ports of the connection request packets are randomized; the purpose is to force the target host to maintain state information for many connections that will never be completed.

SYN flood attacks are usually noticed because the target host (frequently an HTTP or SMTP server) becomes extremely slow, crashes, or hangs. It's also possible for the traffic returned from the target host to cause trouble on routers; because this return traffic goes to the randomized source addresses

Characterizing and Tracing Packet Floods Using Cisco Routers

of the original packets, it lacks the locality properties of "real" IP traffic, and may overflow route caches. On Cisco routers, this problem often manifests itself in the router running out of memory.

Together, smurf and SYN flood attacks account for the vast majority of the flooding DoS attacks reported to Cisco, and recognizing them quickly is very important. Luckily, both attacks (as well as some "second tier" attacks, such as ping floods) are easily recognized using Cisco access lists.

A DoS Characterization Access List

Imagine a router with two interfaces. Ethernet 0 is connected to an internal LAN at a business or small ISP. Serial 0 provides an Internet connection via an upstream ISP. The input packet rate on serial 0 is "pegged" at the full link bandwidth, and hosts on the LAN are running slowly, crashing, hanging, or showing other signs of a DoS attack. The small site at which the router is connected has no network analyzer, and the people there have little or no experience in reading analyzer traces even if the traces were available.

Now, suppose that we apply an access list as follows:

```
access-list 169 permit icmp any any echo
access-list 169 permit icmp any any echo-reply
access-list 169 permit udp any any eq echo
access-list 169 permit udp any eq echo any
access-list 169 permit tcp any any established
access-list 169 permit tcp any any
access-list 169 permit ip any any
```

```
interface serial 0
ip access-group 169 in
```

This list doesn't filter out any traffic at all; all the entries are permits. However, because it categorizes packets in useful ways, the list can be used to tentatively diagnose all three types of attacks: smurf, SYN floods, and fraggle.

Smurf Ultimate Target

If we issue the show access-list command, we'll see output similar to the following:

```
Extended IP access list 169
permit icmp any any echo (2 matches)
permit icmp any any echo-reply (21374 matches)
permit udp any any eq echo
permit udp any eq echo any
permit tcp any any established (150 matches)
permit tcp any any (15 matches)
permit ip any any (45 matches)
```

It's obvious that most of the traffic arriving on the serial interface consists of ICMP echo reply packets. This is probably the signature of a smurf attack, and our site is the ultimate target, rather than the reflector. We can easily gather more information about the attack by revising the access list, as shown below:

Characterizing and Tracing Packet Floods Using Cisco Routers

```
interface serial 0
no ip access-group 169 in

no access-list 169
access-list 169 permit icmp any any echo
access-list 169 permit icmp any any echo-reply log-input
access-list 169 permit udp any any eq echo
access-list 169 permit udp any eq echo any
access-list 169 permit tcp any any established
access-list 169 permit tcp any any
access-list 169 permit ip any any
```

```
interface serial 0
ip access-group 169 in
```

The change here is that we've added the log-input keyword to the access list entry that matches the suspect traffic. (Cisco IOS software earlier than version 11.2 lacks this keyword, and we would use the keyword "log" instead.) This will cause the router to log information about packets that match the list entry. Assuming that logging buffered is configured, we can see the resulting messages with the show log command (it may take a while for the messages to accumulate because of rate limiting). The messages might look something like this:

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.142 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.113 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.212.72 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 172.16.132.154 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.15 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.142 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 172.16.132.47 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.212.35 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.113 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 172.16.132.59 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.82 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.212.56 (Serial0 *HDLC*) -> 10.2.3.7
(0/0), 1 packet
```

Characterizing and Tracing Packet Floods Using Cisco Routers

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 172.16.132.84 (Serial0 *HDLC*) -> 10.2.3.7 (0/0), 1 packet
```

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.212.47 (Serial0 *HDLC*) -> 10.2.3.7 (0/0), 1 packet
```

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.45.35 (Serial0 *HDLC*) -> 10.2.3.7 (0/0), 1 packet
```

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 192.168.212.15 (Serial0 *HDLC*) -> 10.2.3.7 (0/0), 1 packet
```

```
%SEC-6-IPACCESSLOGDP: list 169 denied icmp 172.16.132.33 (Serial0 *HDLC*) -> 10.2.3.7 (0/0), 1 packet
```

We see that the source addresses of the echo reply packets are clustered in a few address prefixes: 192.168.212.0/24, 192.168.45.0/24, and 172.16.132.0/24. This is very characteristic of a smurf attack, and the source addresses are the addresses of the smurf reflectors. By looking up the owners of these address blocks in the appropriate Internet "whois" databases, we can find the administrators of these networks, and ask for their help in dealing with the attack.

It's important at this point in a smurf incident to remember that these reflectors are fellow victims, not attackers. It's extremely rare for attackers to use their own source addresses on IP packets in any DoS flood, and impossible for them to do so in a working smurf attack. Any address in a flood packet should be assumed to be either completely falsified, or the address of a victim of some sort. The most productive approach for the ultimate target of a smurf attack is to contact the reflectors, either to ask them to reconfigure their networks to shut down the attack, or to ask for their assistance in tracing the stimulus stream.

Because the damage to the ultimate target of a smurf attack is usually caused by overloading of the incoming link from the Internet, there's often no response other than to contact the reflectors; by the time the packets arrive at any machine under the target's control, most of the damage has already been done.

One stopgap measure is to ask the upstream network provider to filter out all ICMP echo replies, or all ICMP echo replies from specific reflectors. This sort of filter shouldn't usually be left in place permanently. Even for a temporary filter, only echo replies should be filtered, not all ICMP packets. Another possibility is to have the upstream provider use quality of service and rate limiting features to restrict the bandwidth available to echo replies; a reasonable bandwidth limitation can be left in place indefinitely. Both of these approaches depend on the upstream provider's equipment having the necessary capacity, and sometimes that capacity is not available.

Smurf Reflector

If the incoming traffic consisted of echo requests rather than echo replies (in other words, if the first access list entry, rather than the second, was counting many more matches than could reasonably be expected), we'd suspect a smurf attack in which our network was being used as a reflector, or possibly a simple ping flood. In either case, if the attack was successful, we would expect the outgoing side of the serial line to be swamped, as well as the incoming side. In fact, because of the amplification factor, we would expect the outgoing side to be even more overloaded than the incoming side.

Characterizing and Tracing Packet Floods Using Cisco Routers

We have several ways to distinguish the smurf attack from the simple ping flood:

Smurf stimulus packets will be sent to a directed broadcast address, rather than to a unicast address, whereas ordinary ping floods almost always use unicasts. We can see the addresses using the log-input keyword on the appropriate access list entry, as described in the previous section.

If you are being used as a smurf reflector, there will be a disproportionate number of output broadcasts in the show interface display on the Ethernet side of the system, and usually a disproportionate number of broadcasts sent in the show ip traffic display. A standard ping flood will not increase the background broadcast traffic.

If you are being used as a smurf reflector, there will be more traffic outgoing toward the Internet than traffic incoming from the Internet. In general, there will be more output packets than input packets on the serial interface. Even if the stimulus stream is completely filling the input interface, the response stream will be larger than the stimulus stream, and packet drops will be counted. A smurf reflector has more options than the ultimate target of a smurf attack. If a reflector chooses to shut down the attack, appropriate use of no ip directed-broadcast (or equivalent non-IOS commands) will usually suffice. These commands belong in every configuration, even if there is no active attack; for more information on preventing your Cisco equipment from being used in a smurf attack, see Improving Security on Cisco Routers. For more general information about smurf attacks in general, and for information about protecting non-Cisco equipment.

A smurf reflector is one step closer to the attacker than is the ultimate target, and is therefore in a better position to trace the attack. If you choose to trace the attack, you will need to work with the ISPs involved, and if you wish to have any action taken when you complete the trace, you will need to work with appropriate law enforcement agencies. If you seek to trace an attack, you should involve law enforcement as soon as possible. See the Tracing section for technical information on tracing flooding attacks.

Fraggle

The fraggle attack is analogous to the smurf attack, except that UDP echo requests are used for the stimulus stream instead of ICMP echo requests. The third and fourth lines of the access list identify fraggle attacks. The appropriate response for the victims is the same, except that UDP echo is a less important service in most networks than is ICMP echo, and can therefore be disabled completely with fewer negative consequences.

SYN Floods

The fifth and sixth lines of our access list are:

```
access-list 169 permit tcp any any established
access-list 169 permit tcp any any
```

The first of these lines matches any TCP packet with the ACK bit set. For our purposes, what this really means is that it matches any packet that is not a TCP SYN. The second line therefore matches only packets that are TCP SYNs. A SYN flood is easily identified from the counters on these list entries; in normal traffic, non-SYN TCP packets outnumber SYNs by at least a factor of two, and usually more like four or five. In a SYN flood, SYNs typically outnumber non-SYN TCP packets many times over.

Characterizing and Tracing Packet Floods Using Cisco Routers

The only non-attack condition that creates this signature is a massive overload of genuine connection requests. In general, such an overload will not come unexpectedly, and will not involve as many SYN packets as a real SYN flood. Also, SYN floods often contain packets with completely invalid source addresses; using the log-input keyword, it's possible to see if connection requests are coming from such addresses.

There is an attack usually called a "process table attack" which bears some similarity to the SYN flood. In the process table attack, the TCP connections are actually completed, then allowed to time out with no further protocol traffic, whereas in the SYN flood, only the initial connection requests are sent. Because a process table attack requires completing the TCP initial handshake, it must generally be launched using the IP address of a real machine to which the attacker has access (usually stolen access). Process table attacks are therefore easily distinguished from SYN floods using packet logging; all the SYNs in a process table attack will come from one or a few addresses, or at the most from one or a few subnets.

Unfortunately, response options for the victims of SYN floods are very limited. The system under attack is usually an important service, and blocking access to the system will usually accomplish what the attacker wants. Many router and firewall products, including Cisco's, have features that can be used to reduce the impact of SYN floods, but the effectiveness of these features depends on the environment. For more information, see the documentation for the Cisco IOS Firewall Feature Set, the documentation for the Cisco IOS TCP Intercept feature, and Improving Security on Cisco Routers.

It is possible to trace SYN floods, but the tracing process requires the assistance of each ISP along the path from the attacker to the victim. If you decide to try to trace a SYN flood, contact law enforcement early on, and work with your own upstream service provider. See the tracing section of this document for details on tracing using Cisco equipment.

Other Attacks

If you believe that you are under an attack, and if you can characterize that attack using IP source and destination addresses, protocol numbers, and port numbers, you can use access lists to test your hypothesis. Create an access list entry that matches the suspect traffic, apply it to an appropriate interface, and either watch the match counters or log the traffic.

Logging and Counter Caveats

Remember that the counter on an access list entry counts all matches against that entry. If you apply an access list to two interfaces, the counts you'll see will be aggregate counts.

Access list logging does not show every packet that matches an entry. Logging is rate-limited to avoid CPU overload. What logging shows you is a reasonably representative sample, but not a complete packet trace. Remember that there are packets you're not seeing.

In some software versions, access list logging works only in certain switching modes. If an access list entry is counting a lot of matches, but logging nothing, try clearing the route cache to force packets to be process switched. Be careful about doing this on heavily loaded routers with many interfaces; a lot of traffic can get dropped while the cache is rebuilt. Use CEF whenever possible.

Characterizing and Tracing Packet Floods Using Cisco Routers

Access lists and logging have a performance impact, but not a large one. Be careful on routers running at more than about 80 percent CPU load, or when applying access lists to very high-speed interfaces.

Tracing

The source addresses of DoS packets are almost always set to values that have nothing to do with the attackers themselves, and are therefore of no use in identifying the attackers. The only reliable way to identify the source of an attack is to trace it back hop-by-hop through the network. This process involves reconfiguring routers and examining log information, and therefore requires cooperation by all network operators along the path from the attacker to the victim. Securing that cooperation usually requires the involvement of law enforcement agencies, who must also be involved if any action is to be taken against the attacker.

The tracing process for DoS floods is relatively simple. Starting at a router (call it "A") that's known to be carrying flood traffic, one identifies the router (call it "B") from which A is receiving the traffic. One then logs into B, and finds the router (say "C") from which B is receiving the traffic. This continues until the ultimate source is found.

There are several complications in this method, which are described below:

The "ultimate source" may in fact be a computer which has been compromised by the attacker, but which is actually owned and operated by another victim. In this case, tracing the DoS flood will only be the first step.

Attackers know that they can be traced, and will usually continue their attacks only for a limited time; there may not be enough time to actually trace the flood.

Attacks may be coming from multiple sources, especially if the attacker is relatively sophisticated. It's important to try to identify as many sources as possible.

Communication problems slow down the tracing process. Frequently one or more of the network operators involved will not have appropriately skilled staff available.

Legal and political concerns may make it difficult to act against attackers even if one is found. The fact is that most efforts to trace DoS attacks fail. Because of this, many network operators will not even attempt to trace an attack unless placed under pressure. Many others will trace only "severe" attacks, with differing definitions of what is "severe." Some will assist with a trace only if law enforcement is involved.

Tracing With "log-input"

If you choose to trace an attack passing through a Cisco router, the most effective way of doing so is to construct an access list entry that matches the attack traffic, attach the log-input keyword to it, and apply the access list outbound on the interface through which the attack stream is being sent toward its ultimate target. The log entries produced by the access list will identify the router interface through which the traffic is arriving, and, if the interface is a multipoint connection, will give the layer 2 address of the device from which it is being received. The layer 2 address can then be used to identify the next router in the chain, using, for example, the show ip arp mac-address command.

SYN Flood

Characterizing and Tracing Packet Floods Using Cisco Routers

To trace a SYN flood, you might create an access list similar to the following:

```
access-list 169 permit tcp any any established
access-list 169 permit tcp any host victim-host log-input
access-list 169 permit ip any any
```

This will log all SYN packets destined for the target host, including legitimate SYNs. To identify the most likely actual path toward the attacker, examine the log entries in detail. In general, the source of the flood will be the source from which the largest number of matching packets are arriving. Remember that the source IP addresses themselves mean nothing; you're looking for source interfaces and source MAC addresses. Sometimes it's possible to distinguish flood packets from legitimate packets because flood packets may have invalid source addresses; any packet whose source address is not valid is likely to be part of the flood.

Remember that the flood may be coming from multiple sources, although this is relatively unusual for SYN floods.

Smurf Stimulus

To trace a smurf stimulus stream, use an access list like this:

```
access-list 169 permit icmp any any echo log-input
access-list 169 permit ip any any
```

Note that the first entry doesn't restrict itself to packets destined for the reflector address. The reason for this is that most smurf attacks use multiple reflector networks. If you're not in contact with the ultimate target, you may not know all the reflector addresses. As your trace gets closer to the source of the attack, you may begin to see echo requests going to more and more destinations; this is a good sign.

However, if you're dealing with a great deal of ICMP traffic, this may generate too much logging information for you to read easily. If this happens, you can restrict the destination address to be one of the reflectors that's known to be used. Another useful tactic is to use an entry that takes advantage of the fact that netmasks of 255.255.255.0 are very common in the Internet. And, because of the way that attackers find smurf reflectors, the reflector addresses actually used for smurf attacks are even more likely to match that mask. Host addresses ending in .0 or .255 are very uncommon in the Internet, so you can build a relatively specific recognizer for smurf stimulus streams like this:

```
access-list 169 permit icmp any host known-reflector echo log-input
access-list 169 permit icmp any 0.0.0.255 255.255.255.0 echo log-input
access-list 169 permit icmp any 0.0.0.0 255.255.255.0 echo log-input
access-list 169 permit ip any any
```

With this list, you can eliminate many of the "noise" packets from your log, while still having a good chance of noticing additional stimulus streams as you get closer to the attacker.

Tracing Without "log-input"

The log-input keyword exists in Cisco IOS software versions 11.2 and later, and in certain 11.1-based software created specifically for the service provider market. Older software does not support this keyword. If you're using a router with older software, you have three viable options:

Characterizing and Tracing Packet Floods Using Cisco Routers

Create an access list without logging, but with entries that match the suspect traffic. Apply the list on the input side of each interface in turn, and watch the counters. Look for interfaces with high match rates. This method has a very small performance overhead, and is good for identifying source interfaces. Its biggest drawback is that it doesn't give link-layer source addresses, and is therefore useful mostly for point-to-point lines.

Create access list entries with the log keyword (as opposed to log-input). Once again, apply the list to the incoming side of each interface in turn. This method still doesn't give source MAC addresses, but can be useful for seeing IP data, for instance to verify that a packet stream really is part of an attack. Performance impact can be moderate to high; newer software performs better than older software.

Use debug ip packet detail to collect information about packets. This method gives MAC addresses, but can have serious performance impact. It's easy to make a mistake with this method and make a router unusable. If you use this method, make sure that the router is switching the attack traffic in fast, autonomous, or optimum mode. Use an access list to restrict debugging to only the information you really need. Log debugging information to the local log buffer, but turn off logging of debug information to Telnet sessions and to the console. If possible, arrange for someone to be physically near the router, so that it can be power cycled as necessary.

Remember that debug ip packet will not display information about fast-switched packets; you will need to do clear ip cache to capture information. Each clear command will give you one or two packets of debug output.