

Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks

Document ID: 13634

Introduction

Understanding the Basics of DDoS Attacks

Characteristics of Common Programs Used to Facilitate Attacks

Prevention

Capturing Evidence and Contacting Law Enforcement


Related Information

Introduction

This white paper contains information in order to help you understand how Distributed Denial of Service (DDoS) attacks are orchestrated, recognize programs used to facilitate DDoS attacks, apply measures to prevent the attacks, gather forensic information if you suspect an attack, and learn more about host security.

Understanding the Basics of DDoS Attacks

Refer to this illustration:



Behind a **Client** is a person that orchestrate an attack. A **Handler** is a compromised host with a special program running on it. Each handler is capable of controlling multiple agents. An **Agent** is a compromised host that runs a special program. Each agent is responsible for generating a stream of packets that is directed toward the intended victim.

Attackers have been known to use these four programs to launch DDoS attacks:

1. Trinoo
2. TFN
3. TFN2K
4. Stacheldraht

In order to facilitate DDoS, the attackers need to have several hundred to several thousand compromised hosts. The hosts are usually Linux and SUN computers; but, the tools can be ported to other platforms as well. The process of compromising a host and installing the tool is automated. The process can be divided into these steps, in which the attackers:

1. Initiate a scan phase in which a large number of hosts (on the order of 100,000 or more) are probed for a known vulnerability.
2. Compromise the vulnerable hosts to gain access.
3. Install the tool on each host.
4. Use the compromised hosts for further scanning and compromises.

Because an automated process is used, attackers can compromise and install the tool on a single host in under five seconds. In other words, several thousand hosts can be compromised in under an hour.

Characteristics of Common Programs Used to Facilitate Attacks

These are common programs that hackers use in order to facilitate distributed denial of services attacks:

- Trinoo

Communication between clients, handlers and agents use these ports:

```
1524 tcp
27665 tcp
27444 udp
31335 udp
```

Note: The ports listed above are the *default* ports for this tool. Use these ports for orientation and example only, because the port numbers can easily be changed.

- TFN

Communication between clients, handlers and agents use ICMP ECHO and ICMP ECHO REPLY packets.

- Stacheldraht

Communication between clients, handlers and agents use these ports:

```
16660 tcp
65000 tcp
ICMP ECHO
ICMP ECHO REPLY
```

Note: The ports previously listed are the default ports for this tool. Use these ports for orientation and example only, because the port numbers can easily be changed.

- TFN2K

Communication between clients, handlers and agents does not use any specific port, for example, it may be supplied on run time or it is chosen randomly by a program, but is a combination of UDP, ICMP and TCP packets.

For a detailed analysis of DDoS programs, read these articles.

Note: Theaw links point to external web sites not maintained by Cisco Systems.

The DoS Project's "trinoo" distributed denial of service attack tool

The "Tribe Flood Network" distributed denial of service attack tool

The "stacheldraht" distributed denial of service attack tool

Additional information regarding DDoS tools and their variants can be found at the Packet Storm web site's [Index of Distributed Attack Tools](#) .

Prevention

These are suggested methods to prevent distributed denial of service attacks.

1. Use the **ip verify unicast reverse-path** interface command on the input interface on the router at the upstream end of the connection.

This feature examines each packet received as input on that interface. If the source IP address does not have a route in the CEF tables that points back to the same interface on which the packet arrived, the router drops the packet.

The effect of Unicast RPF is that it stops SMURF attacks (and other attacks that depend on source IP address spoofing) at the ISP's POP (lease and dial-up). This protects your network and customers, as well as the rest of the Internet. To use unicast RPF, enable "CEF switching" or "CEF distributed switching" in the router. There is no need to configure the input interface for CEF switching. As long as CEF is running on the router, individual interfaces can be configured with other switching modes. RPF is an input side function that enabled on an interface or sub-interface and operates on packets received by the router.

It is very important for CEF to be turned on in the router. RPF does not work without CEF. Unicast RPF is not supported in any 11.2 or 11.3 images. Unicast RPF is included in 12.0 on platforms that support CEF, which includes the AS5800. Hence, unicast RPF can be configured on the PSTN/ISDN dial-up interfaces on the AS5800.

2. Filter all RFC-1918 address space using Access Control Lists (ACLs).

Refer to this example:

```
access-list 101 deny ip 10.0.0.0    0.255.255.255 any
access-list 101 deny ip 192.168.0.0 0.0.255.255 any
access-list 101 deny ip 172.16.0.0  0.15.255.255 any
access-list 101 permit ip any any

interface xy
 ip access-group 101 in
```

Another source of information about special use IPv4 address space that can be filtered is the (now expired) IETF draft 'Documenting Special Use IPv4 Address Blocks that have been registered with IANA.'

3. Apply ingress and egress filtering (see RFC-2267) using ACLs.

Refer to this example:

```
{ ISP Core } -- ISP Edge Router -- Customer Edge Router -- { Customer network }
```

The ISP edge router should only accept traffic with source addresses belonging to the customer network. The customer network should only accept traffic with source addresses other than the

customer network block. This is a sample ACL for an ISP edge router:

```
access-list 190 permit ip {customer network} {customer network mask} any
access-list 190 deny ip any any [log]

interface {ingress interface} {interface #}
    ip access-group 190 in
```

This is a sample ACL for a customer edge router:

```
access-list 187 deny ip {customer network} {customer network mask} any
access-list 187 permit ip any any

access-list 188 permit ip {customer network} {customer network mask} any
access-list 188 deny ip any any

interface {egress interface} {interface #}
    ip access-group 187 in
    ip access-group 188 out
```

If you are able to turn on Cisco Express Forwarding (CEF), the length on the ACLs can be substantially reduced and thus increase performance by enabling unicast reverse path forwarding. In order to support unicast reverse path forwarding, you only need to be able to enable CEF on the router as a whole; the interface on which the feature is enabled does not need to be a CEF switched interface.

4. Use CAR to rate limit ICMP packets.

Refer to this example:

```
interface xy
    rate-limit output access-group 2020 3000000 512000 786000 conform-action
    transmit exceed-action drop

access-list 2020 permit icmp any any echo-reply
```

5. Configure rate limiting for SYN packets.

Refer to this example:

```
access-list 152 permit tcp any host eq www
access-list 153 permit tcp any host eq www established

interface {int}
    rate-limit output access-group 153 45000000 100000 100000
    conform-action transmit exceed-action drop
    rate-limit output access-group 152 1000000 100000 100000
    conform-action transmit exceed-action drop
```

In the previous example, replace:

- ◆ **45000000** with the maximum link bandwidth
- ◆ **1000000** with a value that is between 50% and 30% of the SYN flood rate
- ◆ *burst normal* and *burst max* rates with accurate values

Note that if you set the burst rate greater than 30%, many legitimate SYNs may be dropped. In order to get an idea of where to set the burst rate, use the **show interfaces rate-limit** command in order to display the conformed and exceeded rates for the interface. Your objective is to rate-limit the SYNs as little as necessary to get things working again.



Warning: It is recommended that you first measure amount of SYN packets during normal state (before attacks occur) and use those values to limit. Review the numbers carefully before you

deploy this measure.

If an SYN attack is aimed against a particular host, consider installing an IP filtering package on that host. One such package is IP Filter . Refer to IP Filter Examples for implementation details.

Capturing Evidence and Contacting Law Enforcement

If possible, obtain an attack traffic sample for posterior analysis (commonly known as a 'packet capture'). Use a Solaris or Linux workstation with enough processing power to keep up with the flow of packets. For obtaining such a packet capture, use either the tcpdump program (available for Windows, Solaris and Linux operating systems) or the snoop program (available for the Solaris OS only). This is a basic example of how to use those programs:

```
tcpdump -i interface -s 1500 -w capture_file
snoop -d interface -o capture_file -s 1500
```

The MTU size in this example is 1500; change this parameter if the MTU is greater than 1500.

If you want to involve law enforcement and you are within the United States, contact your local FBI field office. More information is available at the National Infrastructure Protection Center web site. If you are located in Europe, no single point of contact exists. Contact your local law enforcement agency and ask for assistance.

CISCO CANNOT CONTACT LAW ENFORCEMENT AGENCIES ON YOUR BEHALF. The Cisco PSIRT team can work with law enforcement once you have made the initial contacts.

For general host security material, visit the CERT/CC web page.

Related Information

- [Characterizing and Tracing Packet Floods Using Cisco Routers](#)
- [Worm Mitigation Technical Details](#)
- [Improving Security on Cisco Routers](#)
- [Cisco Product Security Incident Response](#)
- [Security @ Cisco](#)
- [Technical Support & Documentation – Cisco Systems](#)

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

© 2009 – 2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Apr 22, 2008

Document ID: 13634
