

# Improving Security on Cisco Routers

Document ID: 13608

---

***Interactive:*** This document offers customized analysis of your Cisco device.

---

## **Introduction**

### **Prerequisites**

- Requirements
- Components Used
- Conventions

### **Background Information**

#### **Password Management**

- enable secret
- service password-encryption (and limitations)

#### **Control Interactive Access**

- Console Ports
- General Interactive Access
- Warning Banners

#### **Commonly Configured Management Services**

- SNMP
- HTTP

#### **Management and Interactive Access via the Internet (and Other Untrusted Networks)**

- Packet Sniffers
- Other Internet Access Dangers

#### **Logging**

- Save Log Information
- Record Access List Violations

#### **Secure IP Routing**

- Anti-Spoofing
- Control Directed Broadcasts
- Path Integrity

#### **Flood Management**

- Transit Floods
- Router Self-Protection

#### **Possibly Unnecessary Services**

- TCP and UDP Small Services
- Finger
- NTP
- CDP

#### **Stay Up To Date**

#### **Command List**

#### **NetPro Discussion Forums – Featured Conversations**

#### **Related Information**

---

# Introduction

This document is an informal discussion of some Cisco configuration settings that network administrators should consider changing on their routers, especially on their border routers, in order to improve security. This document is about basic boilerplate configuration items that are almost universally applicable in IP networks, and about a few unexpected items of which you should be aware.

If you have the output of a **show running-configuration** command from your Cisco device, you can use to display potential issues and fixes. You must be a registered customer, be logged in, and have JavaScript enabled in order to use.

You can use Output Interpreter to display potential issues and fixes. You must be a registered customer, be logged in, and have JavaScript enabled in order to use Output Interpreter.

## Prerequisites

### Requirements

There are no specific requirements for this document.

### Components Used

This document is not restricted to specific software and hardware versions.

### Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

## Background Information

This is not an exhaustive list, nor can it be substituted for understanding on the part of the network administrator. This document is a reminder of some of the things that are sometimes forgotten. Only commands that are important in IP networks are mentioned. Many of the services that are enabled in Cisco routers require careful security configuration. However, this document concerns itself mainly with services that are enabled by default, or that are almost always enabled by users, and that might need to be disabled or reconfigured.

This is particularly important because some of the default settings in Cisco IOS® software are there for historical reasons. The settings made sense when chosen, but might be different if new defaults are chosen today. Other defaults make sense for most systems, but can create security exposures if they are used in devices that form part of a network perimeter defense. Other defaults are actually required by standards, but are not always desirable from a security point of view.

Cisco IOS software has many security-specific features, such as packet-filtering access lists, the Cisco IOS Firewall Feature Set, TCP Intercept, AAA, and encryption. Many other features, such as packet logging and quality of service (QoS) features, can be used to increase network security against various attacks. None of these are discussed, except in passing. This is not a document about firewall configuration. For the most part, this is a document about how to secure the router itself, and ignores the equally important issue of the protection of other network devices.

# Password Management

Passwords and similar secrets, such as Simple Network Management Protocol (SNMP) community strings, are the primary defense against unauthorized access to your router. The best way to handle most passwords is to maintain them on a TACACS+ or RADIUS authentication server. However, almost every router still has a locally configured password for privileged access, and can also have other password information in its configuration file.

## enable secret

The **enable secret** command is used to set the password that grants privileged administrative access to the IOS system. An enable secret password must always be set. Use the **enable secret** command, *not* the older **enable password** command. The **enable password** command uses a weak encryption algorithm. See the service password–encryption section of this document for more information.

If no enable secret is set, and a password is configured for the console TTY line, the console password can be used to receive privileged access, even from a remote VTY session. This is almost certainly not what you want, and is another reason to be certain to configure an enable secret.

## service password–encryption (and limitations)

The **service password–encryption** command directs the IOS software to encrypt the passwords, CHAP secrets, and similar data that are saved in its configuration file. This is useful to prevent casual observers from reading passwords, such as when they look at the screen over the shoulder of an administrator.

However, the algorithm used by the **service password–encryption** command is a simple Vigenere cipher. Any competent amateur cryptographer can easily reverse it in a few hours. The algorithm is not designed to protect configuration files against serious analysis by even slightly sophisticated attackers, and should not be used for this purpose. Any Cisco configuration file that contains encrypted passwords must be treated with the same care used for a cleartext list of those same passwords.

This weak encryption warning does not apply to passwords set with the **enable secret** command, but it does apply to passwords set with the **enable password** command.

The **enable secret** command uses MD5 for password hashing. The algorithm has had considerable public review, and is not reversible as far as Cisco knows. It is, however, subject to dictionary attacks. A dictionary attack is when a computer tries every word in a dictionary or other list of candidate passwords. Therefore, remember to keep your configuration file out of the hands of untrusted people, especially if you are not sure your passwords are well chosen.

# Control Interactive Access

Anyone who can log in to a Cisco router can display information which you probably do not want to make available to the general public. A user who can log in to the router might be able to use it as a relay for further network attacks. Anyone who can get privileged access to the router can reconfigure it. You need to control interactive logins to the router in order to prevent inappropriate access.

Although most interactive access is disabled by default, there are exceptions. The most obvious exception is the interactive sessions that are from directly connected asynchronous terminals, such as the console terminal, and from integrated modem lines.

## Console Ports

It is important to remember that the console port of a Cisco IOS device has special privileges. In particular, if a BREAK signal is sent to the console port during the first few seconds after a reboot, the password recovery procedure can easily be used to take control of the system. This means that attackers who interrupt power or induce a system crash, and who have access to the console port via a hardwired terminal, a modem, a terminal server, or some other network device, can take control of the system, even if they do not have physical access to it or the ability to log in to it normally.

Any modem or network device that gives access to the Cisco console port must be secured to a standard comparable to the security used for privileged access to the router. At a bare minimum, any console modem should be of a type that can require the dialup user to supply a password for access, and the modem password must be carefully managed.

## General Interactive Access

There are more ways to get interactive connections to routers than users realize. Cisco IOS software, which depends on the configuration and software version, can support these connections:

- via Telnet
- rlogin
- SSH
- non IP-based network protocols, such as LAT, MOP, X.29, and V.120
- possibly other protocols
- via local asynchronous connections and modem dial-ins

More protocols for interactive access are always being added. Interactive Telnet access is available not only on the standard Telnet TCP port (port 23), but on a variety of higher-numbered ports as well.

All interactive access mechanisms use the IOS TTY abstraction (in other words, they all involve sessions on lines of one sort or another). Local asynchronous terminals and dialup modems use standard lines, known as TTYs. Remote network connections, regardless of the protocol, use virtual TTYs (VTYs). The best way to protect a system is to make certain that appropriate controls are applied on all lines, which includes both VTY lines and TTY lines.

Because it is difficult to make certain that all possible modes of access have been blocked, administrators should use some sort of authentication mechanism in order to make sure that logins on all lines are controlled, even on machines that are supposed to be inaccessible from untrusted networks. This is especially important for VTY lines and for lines connected to modems or other remote access devices.

The **login** and **no password** commands can be configured in order to completely prevent interactive logins. This is the default configuration for VTYs, but not for TTYs. There are many ways to configure passwords and other forms of user authentication for TTY and VTY lines. Refer to the Cisco IOS software documentation for more information.

## Control TTYs

Local asynchronous terminals are less common than they once were, but they still exist in some installations. Unless the terminals are physically secured, and usually even if they are, the router should be configured to require users on local asynchronous terminals to log in before they use the system. Most TTY ports in modern routers are either connected to external modems, or are implemented by integrated modems. The security of these ports is obviously even more important than securing local terminal ports.

By default, a remote user can establish a connection to a TTY line over the network. This is known as reverse Telnet. This allows the remote user to interact with the terminal or modem connected to the TTY line. It is possible to apply password protection for such connections. Often, it is desirable to allow users to make connections to modem lines, so that they can make outgoing calls. However, this feature can allow a remote user to connect to a local asynchronous terminal port, or even to a dial-in modem port, and simulate the login prompt of the router to steal passwords. This feature can also do other things that can trick local users or interfere with their work.

Issue the **transport input none** configuration command in order to disable this reverse Telnet feature on any asynchronous or modem line that should not receive connections from network users. If possible, do not use the same modems for both dial-in and dial-out, and do not allow reverse Telnet connections to the lines you use for dial-in.

## Control VTYS and Ensure VTY Availability

Any VTY must be configured to accept connections only with the protocols actually needed. This is performed with the **transport input** command. For example, a VTY that is expected to receive only Telnet sessions is configured with the **transport input telnet** command, while a VTY that permits both Telnet and SSH sessions has the **transport input telnet ssh** command. If your software supports an encrypted access protocol such as SSH, then enable only that protocol, and disable cleartext Telnet. Also, issue the **ip access-class** command in order to restrict the IP addresses from which the VTY accepts connections.

A Cisco IOS device has a limited number, usually five, of VTY lines. When all of the VTYS are in use, no more remote interactive connections can be established. This creates the opportunity for a denial-of-service attack. If an attacker can open remote sessions to all the VTYS on the system, the legitimate administrator might not be able to log in. The attacker does not have to log in to do this. The sessions can simply be left at the login prompt.

One way to reduce this exposure is to configure a more restrictive **ip access-class** command on the last VTY in the system than on the other VTYS. The last VTY, usually VTY 4, can be restricted to accept connections only from a single, specific administrative workstation, whereas the other VTYS can accept connections from any address in a corporate network.

Another useful tactic is to issue the **exec-timeout** command in order to configure VTY timeouts. This prevents an idle session from consuming a VTY indefinitely. Although its effectiveness against deliberate attacks is relatively limited, it also provides some protection against sessions accidentally left idle. Similarly, if you enable TCP keepalives on incoming connections with the **service tcp-keepalives-in** command, this can help to guard against both malicious attacks and orphaned sessions caused by remote system crashes.

You can disable all non IP-based remote access protocols and use IPSec encryption for all remote interactive connections to the router in order to provide complete VTY protection. IPSec is an extra-cost option, and its configuration is beyond the scope of this document.

## Warning Banners

In some jurisdictions, civil and criminal prosecution of crackers who break into your systems is made much easier if you provide a banner that informs unauthorized users that their use is unauthorized. In other jurisdictions, you can be forbidden to monitor the activities of even unauthorized users unless you have taken steps to notify them of your intent. One method to provide this notification is to put it into a banner message configured with the Cisco IOS **banner login** command.

Legal notification requirements are complex, and vary in each jurisdiction and situation. Even within jurisdictions, legal opinions vary, and this issue should be discussed with your own legal counsel. In cooperation with counsel, you must consider what information is put into your banner:

- A notice that the system is to be logged in to or used only by specifically authorized personnel, and perhaps information about who can authorize use.
- A notice that any unauthorized use of the system is unlawful, and can be subject to civil and/or criminal penalties.
- A notice that any use of the system can be logged or monitored without further notice, and that the resulting logs can be used as evidence in court.
- Specific notices required by specific local laws.

From a security, rather than a legal point of view, your login banner must not contain any specific information about your router, its name, its model, what software it runs, or who owns it. This information can be abused by crackers.

## Commonly Configured Management Services

Many users use protocols other than interactive remote login in order to manage their networks. The most common protocols for this purpose are SNMP and HTTP.

Neither of these protocols is enabled by default, and, as for any other service, the most secure option is to not enable them at all. However, if they are enabled, they must be secured as described in this section.

### SNMP

SNMP is very widely used for router monitoring, and frequently for router configuration changes. Unfortunately, version 1 of the SNMP protocol, which is the most commonly used, uses a very weak authentication scheme based on a community string. This amounts to a fixed password transmitted over the network without encryption. If possible, use SNMP version 2, which supports an MD5-based digest authentication scheme and allows for restricted access to various management data.

If you must use SNMP version 1, choose inobvious community strings. Do not choose, for example, "public" or "private". If possible, avoid the use of the same community strings for all network devices. Use a different string or strings for each device, or at least for each area of the network. Do not make a read-only string the same as a read-write string. If possible, periodic SNMP version 1 polling should be done with a read-only community string. Read-write strings should be used only for actual write operations.

SNMP version 1 is not suited to use across the public Internet for these reasons:

- It uses cleartext authentication strings.
- Most SNMP implementations send those strings repeatedly as part of periodic polling.
- It is an easily spoofable, datagram-based transaction protocol.

You must carefully consider the implications before you use it that way.

In most networks, legitimate SNMP messages come only from certain management stations. If this is true in your network, you should probably use the access list number option on the **snmp-server community** command in order to restrict SNMP version 1 access to only the IP addresses of the management stations. Do not use the **snmp-server community** command for any purpose in a pure SNMP version 2 environment. This command implicitly enables SNMP version 1.

For SNMP version 2, configure digest authentication with the **authentication** and **md5** keywords of the **snmp-server party** configuration command. If possible, use a different MD5 secret value for each router.

SNMP management stations often have large databases of authentication information, such as community strings. This information can provide access to many routers and other network devices. This concentration of information makes the SNMP management station a natural target for attack, and it must be secured accordingly.

## HTTP

Most recent Cisco IOS software versions use the World Wide Web HTTP protocol in order to support remote configuration and monitoring. In general, HTTP access is equivalent to interactive access to the router. The authentication protocol used for HTTP is equivalent to sending a cleartext password across the network. Unfortunately, there is no effective provision in HTTP for challenge-based or one-time passwords. This makes HTTP a relatively risky choice for use across the public Internet.

If you choose to use HTTP for management, issue the **ip http access-class** command in order to restrict access to appropriate IP addresses. Also, issue the **ip http authentication** command in order to configure authentication. As with interactive logins, the best choice for HTTP authentication is to use a TACACS+ or RADIUS server. Avoid the use of the enable password as an HTTP password.

## Management and Interactive Access via the Internet (and Other Untrusted Networks)

Many users manage their routers remotely, and sometimes this is done over the Internet. Any unencrypted remote access carries some risk, but access over a public network such as the Internet is especially dangerous. All remote management schemes, which includes interactive access, HTTP, and SNMP, are vulnerable.

The attacks discussed in this section are relatively sophisticated ones, but they are not out of the reach of crackers today. These attacks can often be thwarted if the public network providers involved have taken proper security measures. You need to evaluate your level of trust in the security measures used by all the providers that carry your management traffic. Even if you trust your providers, it is recommended to take at least some steps to protect yourself from the results of any mistakes that might occur.

All the cautions here apply as much to hosts as to routers. This document discusses the protection of router login sessions, but you should use analogous mechanisms to protect your hosts if you administer those hosts remotely.

Remote Internet administration is useful, but requires careful attention to security.

## Packet Sniffers

Crackers frequently break into computers owned by Internet service providers (ISPs), or into computers on other large networks, and install packet sniffer programs. These programs monitor the traffic that passes through the network and steal data, such as passwords and SNMP community strings. Although this has become more difficult as network operators improve their security, it is still relatively common. In addition to the risk from outside crackers, it is not unheard of for rogue ISP personnel to install sniffers. Any password sent over an unencrypted channel is at risk. This includes the login and enable passwords for your routers.

If possible, avoid logging in to your router that uses any unencrypted protocol over any untrusted network. If your router software supports it, use an encrypted login protocol such as SSH or Kerberized Telnet. Another

possibility is to use IPSec encryption for all router management traffic, which includes Telnet, SNMP, and HTTP. All of these encryption features are subject to certain export restrictions imposed by the United States Government, and are special-order, extra-cost items on Cisco routers.

If you do not have access to an encrypted remote access protocol, another possibility is to use a one-time password system such as S/KEY or OPIE, together with a TACACS+ or RADIUS server. This controls both interactive logins and privileged access to your router. The advantage here is that a stolen password is of no use, because it is made invalid by the very session in which it is stolen. Non-password data transmitted in the session remains available to eavesdroppers, but many sniffer programs are set up to concentrate on passwords.

If you absolutely must send passwords over cleartext Telnet sessions, change your passwords frequently, and pay close attention to the path traversed by your sessions.

## Other Internet Access Dangers

In addition to packet sniffers, remote Internet management of routers presents these security risks:

- In order to manage a router over the Internet, you must permit at least some Internet hosts to have access to the router. It is possible that these hosts can be compromised, or that their addresses can be spoofed. By permitting interactive access from the Internet, you make your security dependent not only on your own anti-spoofing measures, but on those of the service providers involved.

You can make sure that all the hosts that are permitted to log into your router are under your own control in order to reduce dangers. Also, use encrypted login protocols with strong authentication.

- It is sometimes possible to hijack an unencrypted TCP connection (such as a Telnet session), and actually take control away from a user who is logged in. Although such hijack attacks are not as common as simple packet sniffing and can be complex to mount, these attacks are possible, and might be used by an attacker who has your network specifically in mind as a target. The only real solution to the problem of session hijack is to use a strongly authenticated encrypted management protocol.
- Denial of service attacks are relatively common on the Internet. If your network is subjected to a denial of service attack, you might not be able to reach your router to collect information or take defensive action. Even an attack on a network of another person can impair your management access to your own network. Although you can take steps to make your network more resistant to denial of service attacks, the only real defense against this risk is to have a separate, out-of-band management channel, such as a dialup modem, for use in emergencies.

## Logging

Cisco routers can record information about a variety of events, many of which have security significance. Logs can be invaluable to characterize and respond to security incidents. These are the main types of logging used by Cisco routers:

- **AAA logging** Collects information about user dial-in connections, logins, logouts, HTTP accesses, privilege level changes, commands executed, and similar events. AAA log entries are sent to authentication servers that use the TACACS+ and/or RADIUS protocols, and are recorded locally by those servers, typically in disk files. If you use a TACACS+ or RADIUS server, you can enable AAA logging of various sorts. Issue AAA configuration commands, such as **aaa accounting**, in order to enable this. Detailed description of AAA configuration is beyond the scope of this document.
- **SNMP trap logging** Sends notifications of significant changes in system status to SNMP management stations. Use SNMP traps only if you have an SNMP management infrastructure that already exists.

- System logging Records a large variety of events, which depends on the system configuration. System logging events can be reported to a variety of destinations, which include these:
  - ◆ The system console port (**logging console**).
  - ◆ Servers that use the UNIX syslog protocol (**logging ip-address, logging trap**).
  - ◆ Remote sessions on VTYS and local sessions on TTYs (**logging monitor, terminal monitor**).
  - ◆ A local logging buffer in router RAM (**logging buffered**).

From a security point of view, the most important events usually recorded by system logging are interface status changes, changes to the system configuration, access list matches, and events detected by the optional firewall and intrusion detection features.

Each system logging event is tagged with an urgency level. The levels range from debugging information (at the lowest urgency), to major system emergencies. Each logging destination can be configured with a threshold urgency, and receives logging events only at or above that threshold.

## Save Log Information

By default, system logging information is sent only to the asynchronous console port. Because many console ports are unmonitored, or are connected to terminals without historical memory and with relatively small displays, this information might not be available when it is needed, especially when a problem is debugged over the network.

Almost every router must save system logging information to a local RAM buffer. The logging buffer is of a fixed size, and retains only the newest information. The contents of the buffer are lost whenever the router is reloaded. Even so, a moderately-sized logging buffer is often of great value. On low-end routers, a reasonable buffer size might be 16384 or 32768 bytes. On high-end routers with lots of memory (and many logged events), even 262144 bytes might be appropriate. You can issue the **show memory** command to make sure that your router has enough free memory to support a logging buffer. Issue the **logging buffered buffer-size** configuration command in order to create the buffer.

Most larger installations have syslog servers. You can send **logging** information to a server with the **logging server-ip-address**, and you can control the urgency threshold for logging to the server with the **logging trap urgency** command. Even if you have a syslog server, you should still enable local logging.

If your router has a real-time clock or runs NTP, issue the **service timestamps log datetime msec** command in order to time-stamp log entries.

## Record Access List Violations

If you use access lists to filter traffic, you might want to log packets that violate your filtering criteria. Earlier Cisco IOS software versions use the **log** keyword in order to support logging. This causes logging of the IP addresses and port numbers associated with packets that match an access list entry. Later versions provide the **log-input** keyword, which adds information about the interface from which the packet was received, and the MAC address of the host that sent it.

It is not a good idea to configure logging for access list entries that match very large numbers of packets. This causes log files to grow excessively large, and can cut into system performance. However, access list log messages are rate-limited, so the impact is not catastrophic.

Access list logging can also be used to log the suspect traffic in order to characterize traffic associated with network attacks.

# Secure IP Routing

This section discusses some basic security measures related to the way in which the router forwards IP packets. Refer to essential IOS features for more information about these issues.

## Anti-Spoofing

Many network attacks rely on an attacker that falsifies, or spoofs, the source addresses of IP datagrams. Some attacks rely on spoofing to work at all, and other attacks are much harder to trace if the attacker can use the address of someone else instead of his or her own. Therefore, it is valuable for network administrators to prevent spoofing wherever feasible.

Anti-spoofing must be done at every point in the network where it is practical. It is usually both easiest and most effective at the borders between large address blocks, or between domains of network administration. It is usually impractical to perform anti-spoofing on every router in a network, because of the difficulty to determine which source addresses might legitimately appear on any given interface.

If you are an ISP, you might find that effective anti-spoofing along with other effective security measures, causes expensive, annoyed problem subscribers to take their business to other providers. ISPs must apply anti-spoofing controls at dialup pools and other end-user connection points (refer to RFC 2267 ).

Administrators of corporate firewalls or perimeter routers sometimes install anti-spoofing measures to prevent hosts on the Internet from assuming the addresses of internal hosts, but do not take steps to prevent internal hosts from assuming the addresses of hosts on the Internet. Try to prevent spoofing in both directions. There are at least three good reasons to perform anti-spoofing in both directions at an organizational firewall:

1. Internal users are less tempted to launch network attacks and less likely to succeed if they do try.
2. Accidentally misconfigured internal hosts are less likely to cause trouble for remote sites. Therefore, these are less likely to generate angry telephone calls or damage the reputation of your organization.
3. Outside crackers often break into networks as launching pads for further attacks. These crackers might be less interested in a network with outgoing spoofing protection.

## Anti-Spoofing with Access Lists

Unfortunately, it is not practical to give a simple list of commands that provide appropriate spoofing protection. The access list configuration depends too much on the individual network. The basic goal is to discard packets that arrive on interfaces that are not viable paths from the supposed source addresses of those packets. For example, on a two-interface router that connects a corporate network to the Internet, any datagram that arrives on the Internet interface, but whose source address field claims that it came from a machine on the corporate network, should be discarded.

Similarly, any datagram that arrives on the interface connected to the corporate network, but whose source address field claims that it came from a machine outside the corporate network, should be discarded. If CPU resources allow it, anti-spoofing should be applied on any interface where it is feasible to determine what traffic can legitimately arrive.

ISPs that carry transit traffic can have limited opportunities to configure anti-spoofing access lists, but such an ISP can usually at least filter outside traffic that claims to originate within the address space of the ISP.

In general, anti-spoofing filters must be built with input access lists. This means that packets must be filtered at the interfaces through which they arrive at the router, not at the interfaces through which they leave the router. This is configured with the **ip access-group list in** interface configuration command. You can use

output access lists in some two-port configurations in order to anti-spoof, but input lists are usually easier to understand even in those cases. Furthermore, an input list protects the router itself from spoofing attacks, whereas an output list protects only devices behind the router.

When anti-spoofing access lists exist, they should always reject datagrams with broadcast or multicast source addresses, and datagrams with the reserved loopback address as a source address. It is usually appropriate for an anti-spoofing access list to filter out all ICMP redirects, regardless of source or destination address. These are the appropriate commands:

```
access-list number deny icmp any any redirect
access-list number deny ip 127.0.0.0 0.255.255.255 any
access-list number deny ip 224.0.0.0 31.255.255.255 any
access-list number deny ip host 0.0.0.0 any
```

The fourth command filters out packets from many BOOTP/DHCP clients. Therefore, it is not appropriate in all environments.

### Anti-Spoofing with RPF Checks

In almost all Cisco IOS software versions that support Cisco Express Forwarding (CEF), it is possible to have the router check the source address of any packet against the interface through which the packet entered the router. If the input interface is not a feasible path to the source address according to the routing table, the packet is dropped.

This works only when routing is symmetric. If the network is designed in such a way that traffic from host A to host B normally takes a different path than traffic from host B to host A, the check always fails and communication between the two hosts is impossible. This sort of asymmetric routing is common in the Internet core. Make sure that your network does not use asymmetric routing before you enable this feature.

This feature is known as a reverse path forwarding (RPF) check, and is enabled with the **ip verify unicast rpf** command. It is available in Cisco IOS Software Releases 11.1CC, 11.1CT, 11.2GS, and all 12.0 and later versions, but requires that CEF be enabled in order to be effective.

### Control Directed Broadcasts

IP directed broadcasts are used in the extremely common and popular smurf denial of service attack, and can also be used in related attacks.

An IP directed broadcast is a datagram which is sent to the broadcast address of a subnet to which the sending machine is not directly attached. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet, where it is converted into a link-layer broadcast. Because of the nature of the IP addressing architecture, only the last router in the chain, the one that is connected directly to the target subnet, can conclusively identify a directed broadcast. Directed broadcasts are occasionally used for legitimate purposes, but such use is not common outside the financial services industry.

In a smurf attack, the attacker sends ICMP echo requests from a falsified source address to a directed broadcast address. This causes all the hosts on the target subnet to send replies to the falsified source. By sending a continuous stream of such requests, the attacker can create a much larger stream of replies. This can completely inundate the host, whose address is falsified.

If a Cisco interface is configured with the **no ip directed-broadcast** command, directed broadcasts that are otherwise exploded into link-layer broadcasts at that interface are dropped instead. This means that the **no ip directed-broadcast** command must be configured on every interface of every router that is connected to a

target subnet. It is not sufficient to configure only firewall routers. The **no ip directed-broadcast** command is the default in Cisco IOS Software Release 12.0 and later. In earlier releases, the command should be applied to every LAN interface that is not known to forward legitimate directed broadcasts.

For a strategy that blocks smurf attacks on some firewall routers, which depends on the network design, and for more general information on the smurf attack, refer to denial of service attacks .

## Path Integrity

Many attacks depend on the ability to influence the paths datagrams take through the network. If they control routing, crackers can spoof the address of another user machine and have the return traffic sent to them, or they can intercept and read data intended for someone else. Routing can also be disrupted purely for denial of service purposes.

## IP Source Routing

The IP protocol supports source routing options that allow the sender of an IP datagram to control the route that datagram takes toward its ultimate destination, and generally the route that any reply takes. These options are rarely used for legitimate purposes in real networks. Some older IP implementations do not process source-routed packets properly, and it is possible to send them datagrams with source routing options in order to crash machines that run these implementations.

A Cisco router with the **no ip source-route** command set never forwards an IP packet which carries a source routing option. You should use this command, unless your network needs source routing.

## ICMP Redirects

An ICMP redirect message instructs an end node to use a specific router as its path to a particular destination. In an IP network that functions properly, a router sends redirects only to hosts on its own local subnets. No end node ever sends a redirect, and no redirect is ever traversed more than one network hop. However, an attacker can violate these rules. Some attacks are based on this. Filter out incoming ICMP redirects at the input interfaces of any router that lies at a border between administrative domains. Also, it is not unreasonable for any access list that is applied on the input side of a Cisco router interface to filter out all ICMP redirects. This causes no operational impact in a correctly configured network.

This filter prevents only redirect attacks launched by remote attackers. It is still possible for attackers to cause significant trouble using redirects if their host is directly connected to the same segment as a host that is under attack.

## Routing Protocol Filter and Authentication

If you use a dynamic routing protocol that supports authentication, enable that authentication. This prevents malicious attacks on the routing infrastructure, and can also help to prevent damage caused by misconfigured rogue devices on the network.

For the same reasons, service providers and other operators of large networks are generally well advised to use route filtering (with the **distribute-list in** command) to prevent their routers from accepting clearly incorrect routing information. Although excessive use of route filtering can destroy the advantages of dynamic routing, judicious use often helps to prevent unpleasant results. For example, if you use a dynamic routing protocol to communicate with a stub customer network, you should not accept any routes from that customer other than routes to the address space you have actually delegated to the customer.

Detailed instruction on how to configure routing authentication and route filtering is beyond the scope of this document. Documentation is available on the Cisco website and elsewhere. Because of the complexity involved, novices are advised to seek experienced advice before configuring these features on important networks.

## **Flood Management**

Many denial of service attacks rely on floods of useless packets. These floods congest network links, slow down hosts, and can overload routers as well. Careful router configuration can reduce the impact of such floods.

An important part of flood management is to be aware of where performance bottlenecks lie. If a flood overloads a T1 line, then filtering out the flood on the router at the source end of the line is effective, whereas filtering at the destination end has little or no effect. If the router itself is the most overloaded network component, then filtering protections that place heavy demands on the router can make matters worse. Keep this in mind when you consider the implementation of the suggestions in this section.

### **Transit Floods**

It is possible to use Cisco QoS features to protect hosts and links against some kinds of floods. Unfortunately, a general treatment of this sort of flood management is beyond the scope of this document, and the protection depends heavily on the attack. The only simple, generally applicable advice is to use weighted fair queueing (WFQ) wherever CPU resources can support it. WFQ is the default for low-speed serial lines in recent versions of Cisco IOS software. Other features of possible interest include committed access rate (CAR), generalized traffic shaping (GTS), and custom queuing. It is sometimes possible to configure these features when under active attack.

If you do plan to use QoS features to control floods, it is important to understand how those features work, and how common flooding attacks work. For example, WFQ is much more effective against ping floods than against SYN floods. This is because the usual ping flood appears to WFQ as a single traffic flow, whereas each packet in a SYN flood generally appears as a separate flow. A smurf reply stream falls somewhere between the two. A great deal of information about Cisco QoS features is available on the Cisco World Wide Website, and information about common attacks is available at many websites maintained by other parties.

Cisco provides two different router features intended specifically to reduce the impact of SYN flooding attacks on hosts. The TCP Intercept feature is available in certain software versions for many routers with model numbers of 4000 or greater. The Cisco IOS Firewall Feature Set, which is now available on an increased number of Cisco routers, includes a different SYN flood protection feature. SYN flood protection can be complex, and results can vary. This depends on flood rate, router speed and memory size, and the hosts in use. If you configure either of these features, make sure to read the documentation on the Cisco World Wide Website. Also, if possible, test your configuration under an actual flood.

### **Router Self-Protection**

Before a router can protect other parts of the network from the effects of floods, the router itself must be protected from overload.

### **Switching Modes and Cisco Express Forwarding**

The CEF switching mode, available in Cisco IOS Software Releases 11.1CC, 11.1CT, 11.2GS, and 12.0, replaces the traditional Cisco routing cache with a data structure that mirrors the entire system routing table.

Because there is no need to build cache entries when traffic starts to arrive for new destinations, CEF behaves more predictably than other modes when presented with large volumes of traffic addressed to many destinations.

Although most flooding denial of service attacks send all of their traffic to one or a few targets and do not tax the traditional cache maintenance algorithm, many popular SYN flooding attacks use randomized source addresses. The host under attack replies to some fraction of the SYN flood packets, which creates traffic for a large number of destinations. Therefore, routers configured for CEF perform better under SYN floods (directed at hosts, not at the routers themselves) than routers that use the traditional cache. CEF is recommended when available.

## Scheduler Configuration

When a Cisco router is fast-switching a large number of packets, it is possible for the router to spend so much time in response to interrupts from the network interfaces that no other work is done. Some very fast packet floods can cause this condition. Issue the **scheduler interval** command, which instructs the router to stop handling interrupts and attend to other business at regular intervals, in order to reduce the effect. A typical configuration might include the **scheduler interval 500** command, which indicates that process-level tasks are to be handled no less frequently than every 500 milliseconds. This command rarely has any negative effects, and should be a part of your standard router configuration unless you know of a specific reason to leave it out.

Many newer Cisco platforms use the **scheduler allocate** command instead of the **scheduler interval** command. The **scheduler allocate** command takes two parameters: a period in microseconds for the system to run with interrupts enabled, and a period in microseconds for the system to run with interrupts masked. If your system does not recognize the **scheduler interval 500** command, issue the **scheduler allocate 3000 1000** command. These values were chosen to represent the midpoints of the ranges. The range for the first value is 400 to 60000, and the range for the second value is 100 to 4000. These parameters can be tuned.

## Possibly Unnecessary Services

As a general rule, any unnecessary service should be disabled in any router that is reachable from a potentially hostile network. The services listed in this section are sometimes useful, but should be disabled if they are not actively used.

### TCP and UDP Small Services

By default, Cisco devices from Cisco IOS versions 11.3 and earlier offer these small services:

- echo
- chargen
- discard

These services, especially their UDP versions, are infrequently used for legitimate purposes, but can be used to launch denial of service and other attacks that are otherwise prevented by packet filtering.

For example, an attacker might send a DNS packet, which falsifies both the source address to be a DNS server that is otherwise unreachable, and the source port to be the DNS service port (port 53). If such a packet is sent to the Cisco UDP echo port, the result is the Cisco that sends a DNS packet to the server in question. No outgoing access list checks is applied to this packet, because it is locally generated by the router itself.

Although most abuses of the small services can be avoided or made less dangerous by anti-spoofing access lists, the services should almost always be disabled in any router which is part of a firewall or lies in a security-critical part of the network. Because the services are rarely used, the best policy is usually to disable them on all routers of any description.

The small services are disabled by default in Cisco IOS Software Releases 12.0 and later. In earlier software, you can issue the **no service tcp-small-servers** and **no service udp-small-servers** commands in order to disable them.

## Finger

Cisco routers provide an implementation of the finger service, which is used to find out which users are logged into a network device. Although this information is not usually sensitive, it is sometimes useful to an attacker. The finger service can be disabled with the **no service finger** command.

## NTP

The Network Time Protocol (NTP) is not especially dangerous, but any unneeded service can represent a path for penetration. If NTP is actually used, it is important to explicitly configure trusted time source, and to use proper authentication. This is because the corruption of the time base is a good way to subvert certain security protocols. If NTP is not used on a particular router interface, it can be disabled with the **ntp disable** interface command.

## CDP

Cisco Discovery Protocol (CDP) is used for some network management functions, but is dangerous because it allows any system on a directly-connected segment to learn that the router is a Cisco device, and to determine the model number and the Cisco IOS software version that is run. This information can be used to design attacks against the router. CDP information is accessible only to directly-connected systems. The CDP protocol can be disabled with the **no cdp running** global configuration command. CDP can be disabled on a particular interface with the **no cdp enable** command.

## Stay Up To Date

Like all software, Cisco IOS software has bugs. Some of these bugs have security implications. In addition, new attacks are always invented, and behavior that might have been considered correct when a piece of software was written can have bad effects when deliberately exploited.

When a major new security vulnerability is found in a Cisco product, Cisco generally issues an advisory notice about the vulnerability. Refer to Cisco Product Security Incident Response for information about the process through which these notices are issued. Refer to Cisco Product Security Advisories and Notices for information on the notices.

Almost any unexpected behavior of any piece of software might create a security exposure somewhere, and only bugs with especially direct implications for system security are mentioned in advisories. Your security is enhanced if you keep your software up to date even in the absence of any security advisory.

Some security problems are not caused by software bugs, and it is important for network administrators to stay aware of trends in attacks. A number of World Wide Websites, Internet mailing lists, and Usenet newsgroups are concerned with this.

## Command List

This section is intended to serve as a reminder of the configuration suggestions in the other sections of this document. Cisco IOS configuration command names are used in this table as mnemonic aids. Always read the documentation for any command before you use it.

Use	To
<b>enable secret</b>	Configure a password for privileged router access.
<b>service password-encryption</b>	Provide a minimum of protection for configured passwords.
<b>no service tcp-small-servers</b>	Prevent abuse of the small services for denial of service or other attacks. Avoid the release of user information to possible attackers.
<b>no service udp-small-servers</b>	
<b>no service finger</b>	
<b>no cdp running</b>	Avoid the release of information about the router to directly-attacking devices.
<b>no cdp enable</b>	
<b>ntp disable</b>	Prevent attackers from using the router as a "smurf" amplifier. Control which protocols can be used by remote users to connect interactively to the VTYs of the router or to access its TTY ports.
<b>no ip directed-broadcast</b>	Control which IP addresses can connect to TTYs or VTYs. Reserve one VTY for access from an administrative workstation.
<b>exec-timeout</b>	Prevent an idle session from tying up a VTY indefinitely.
<b>service tcp-keepalives-in</b>	Detect and delete dead interactive sessions, which prevents them from tying up VTYs.
<b>logging buffered buffer-size</b>	Save logging information in a local RAM buffer on the router. With newer software, the buffer size can be followed with an urgency threshold.
<b>ip access-group list in</b>	Discard spoofed IP packets. Discard incoming ICMP redirects.

<b>ip verify unicast rpf</b>	Discard spoofed IP packets in <i>symmetric routing environments</i> with CEF only.
<b>no ip source-route</b>	Prevent IP source routing options from being used to spoof traffic.
<b>access-list number action criteria log</b>	
<b>access-list number action criteria log-input</b>	Enable logging of packets that match specific access list entries. Use <b>log-input</b> if it is available in your software version.
<b>scheduler-interval</b>	
<b>scheduler allocate</b>	Prevent fast floods from shutting down important processing.
<b>ip route 0.0.0.0 0.0.0.0 null 0 244</b>	Rapidly discard packets with invalid destination addresses.
<b>distribute-list list in</b>	Filter routing information to prevent accepting invalid routes.
<b>snmp-server community something-inobvious ro list</b>	
<b>snmp-server community something-inobvious rw list</b>	Enable SNMP version 1, configure authentication, and restrict access to certain IP addresses. Use SNMP version 1 only if version 2 is unavailable, and watch for sniffers.
<b>snmp-server party... authentication md5 secret ...</b>	Enable SNMP only if it is needed in your network. Configure MD5-based SNMP authentication. Enable read-write access unless you need it.
<b>ip http authentication method</b>	Authenticate HTTP connection requests (if you have enabled HTTP on your router).
<b>ip http access-class list</b>	Further control HTTP access by restricting it to certain host addresses (if you have enabled HTTP on your router).
<b>banner login</b>	Establish a warning banner to be displayed to users who try to log into the router.

## NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums – Featured Conversations for Security

Security: Intrusion Detection [Systems]

Security: AAA

Security: General

Security: Firewalling

---

## Related Information

- [Essential IOS Features Every ISP Should Consider](#)
- [The Latest in Denial of Service Attacks: "Smurfing"](#)
- [Cisco Product Security Incident Response](#)
- [Cisco Security Advisories](#)
- [RFC 2267](#)
- [Technical Support & Documentation – Cisco Systems](#)

---

All contents are Copyright © 1992–2006 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

---

Updated: Oct 23, 2006

Document ID: 13608

---