

Security Filtering

- Two reasons to be concerned about security:
 - Increasing role of computers in the workplace
 - Increasing amount of computer connectivity
- Awareness is the end goal.
 - Systems security is an on-going project.

Computers have become a major part of the workplace in the 90s. There are more businesses that depend on computer technology for survival. Typewriters are a thing of the past and have been replaced with computers, creating a paperless office. In a paperless office, data is stored electronically. Networking provides the ability to share the data between computers. Once a computer is connected to a network, the security perimeter becomes extended. A computer that is not connected to a network is contained within a single **box**, which is simple to secure. This is not the case in a network.

With the growing interest in the Internet, more and more people are coming on-line. The Internet has become a huge market for advertising and a major means of communication. Businesses are using the Internet to communicate between satellite offices along with customers. Connecting to the Internet extends the perimeter of a single computer almost indefinitely. Actions must be taken to protect the data on exposed machines.

The rapidly changing technologies of the Internet are advancing at an alarming rate. More programs are being developed to gain access to otherwise restricted data/areas.


Systems security protects machines and machine based resources.

Network security deals with connectivity and data flow between hosts.

Poor network security permits attacks on hosts. Poor host security permits attacks on networks.

PSINet *Protecting Resources*

- ❑ Careless systems administration allows access through weak software or misconfiguration. Uninformed users may choose weak passwords.



Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T11.3

Potential areas of concern:

- System misconfigurations could allow for back doors to system resources and go unnoticed.
- System software could have bugs that create a vulnerable point in a system.
- Inactive accounts could be probed and used without being noticed.
- System logs should be monitored and abnormalities investigated.
- An uninformed user may choose a password that is a pet's name or a dictionary word. They could also give their password to a friend so that the friend could use the account. Some programs allow access and a user could run such a program without knowing the consequences.

- Data that might be at risk:
 - Personal information and e-mail.
 - Corporate data and system software.
 - System configuration and data files.




Corporate secrets, patents or just sensitive data that is exposed to unprivileged users may be seen by anyone who looks. This could result in a monetary loss to the company or worse.

System data that is exposed could provide information about passwords, system configurations, and network configurations that would allow easy access to system resources. Access to system software could allow for the infection of viruses or planting a Trojan Horse. System software could be modified to conceal the evidence of a break-in. Modifying a startup routine could result in the machine being useless or even allow for future attack.

An attacker isn't always a person trying to get information for monetary gain. The attacker could be an ex-employee out for revenge or just a curious high school student trying to impress friends.

- Encryption of data
 - A key is used to encrypt cleartext into cyphertext.
 - Private key (symmetric) encryption
 - Public key (asymmetric) encryption
- Maintaining data integrity
 - Digital signatures or checksums



Encrypting sensitive data restricts its usefulness if the data is compromised. Changing the cyphertext file replaces the original cleartext and yields garbage when decrypted. Swapping cyphertext with other cyphertext is detectable. Place a date stamp or secret in the file before encryption for a method to detect the swapping of cyphertext. You need a **key** to decode the cyphertext file to a cleartext file and presumably only trustworthy people know the key.

A **private key (symmetric)** encryption uses the same key for decryption. Both parties need the key to have access to the cleartext. This poses a problem. How do you get the key to the other party?

A **public key (asymmetric)** encryption uses a different key for decryption. Encryption key is called the **public** key and the decryption key is called the **private** key.

A **signature** or **checksum** is generated from the data. Different data yield different checksums. It is used when integrity is more important than privacy and could be used to perform periodic audits in check data integrity. Again, changes to the original text are permanent and without good backups there is no method to recover the data.



Maintaining Data Integrity

- System access is not necessary to delete data.
- User privileges and access are still important.
 - A user can alter files their account has access to.
 - Password security is key to data integrity.

Catastrophic systems failure could cause a massive loss of data. A hard drive crash could result in the loss of data. A loss of power could cause damage to a file system, corrupting file integrity.

Removable media can be corrupted. A backup tape that is exposed to a magnetic field is useless.

If a file does not have the correct read/write permission, it is possible that unaware or malicious users could delete or alter that file to their benefit.

Password security is key!

- What is the medium? Consider Ethernet media:
 - Twisted pair and coax (10base2 and 10base5)
 - Fiber
- Where is it run?
- System access is not necessary to stop a machine. Again, user access and privilege is important.

Twisted pair and coax are easy to tap or interfere with. There are many tapping and emission monitoring tools that can be used to reconstruct packets that are traversing the wire.

Fiber, on the other hand, is hard to tap or interfere with and is non-emissive. However, it is costly in comparison to twisted pair or coax.

Does the cable physically pass through a public area where tapping would be possible? Could the cabling be broken by a janitor cleaning a vent? These things have to be considered to achieve a secure network.

An **active tap** permits an intruder to alter data in transit as well as to read passing data. A message's sender or recipient's address may be altered. False acknowledgments or queries may be sent causing TCP streams to be corrupt.

A **passive tap** allows an intruder to read data. Corporate data can be read. Passwords or other system information can be read for later use.

A hammer or can of Pepsi could be used to "**break**" a system.

A user can write a program to consume resources reducing the productivity level of others.

Again, **password security is key!**

- Kinds of services that may be attacked:
 - Network services
 - Machine availability
 - Access to servers on a machine
 - Information access

Corruption of data in transit makes the network unusable. The information in the header of a packet is essential for the packet to reach its destination. If the header of the packet were to be destroyed or altered, the results could be hazardous. A routing announcement that is altered could result in a network becoming unreachable.

Networks are vulnerable to flooding from hosts. When bandwidth is scarce and valuable, this becomes a potential problem.

A host could cause a router to never reach an idle-out timer. This may result in a dialup connection being connected 24x7 or much more than desired, driving connection charges higher than expected.

Protecting data flow into the network and accessing control to hosts is very vital to a system's security. To control what packets can traverse the network, you must design a secure network topology using routing techniques and packet filters.



Access Control Methods

- Filtering
- Network topology
- Network routing
- Proxy servers
- Bastion host
- Host routing

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

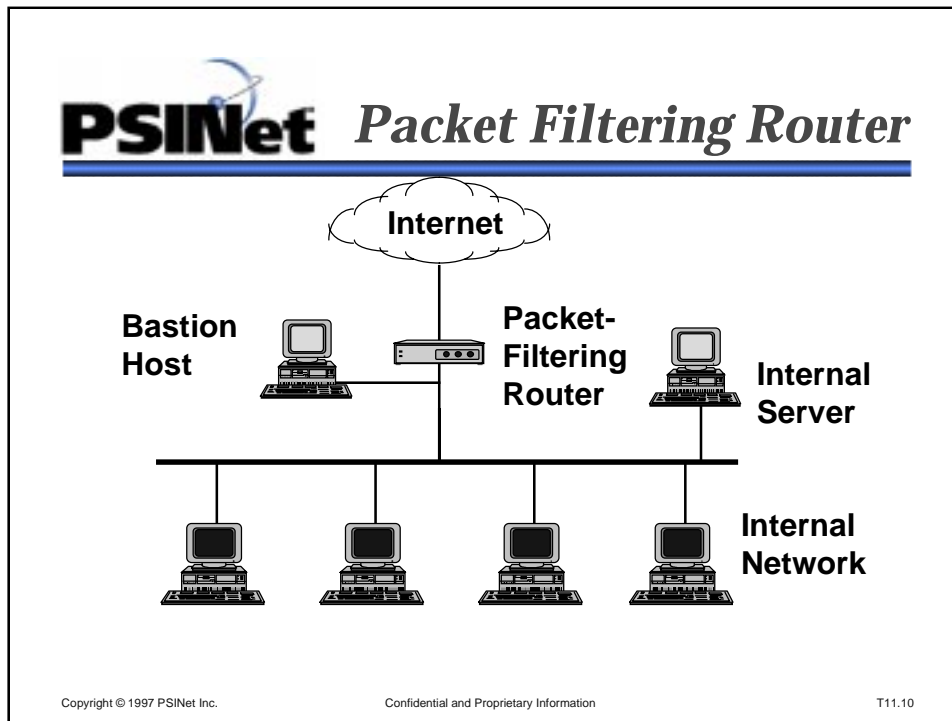
T11.9

Networks need to be designed around security. The network topology could be used to improve security.

Access to IP networks can be controlled with the use of packet filters, proxy servers, and bastion hosts.

Control network access by controlling routing. Restrict routing announcements so that internal networks are not announced to the outside world. Design a scheme where users on internal networks have to use external network machines, bastion hosts, or proxy servers to get to the outside world.

The same scheme could apply to hosts. TCP connections require routing in both directions. A host could be configured with routes only for trusted hosts, thereby restricting access. There is a fault to this scheme; if any trusted machine is compromised, the restricted host will be reachable.



Security can be provided by a packet filtering router. Incoming connections can only be allowed to the bastion host. These connections will be tightly controlled. The filters can be static or dynamic.

If the filters are dynamic, service on the internal network can be accessed by authenticating to the router. The router then opens a hole only for the machine that authenticates. This hole may be allowed for a limited amount of time.

Depending on security policy, the filter can allow outgoing connections directly from internal hosts or outgoing connections only from the bastion host. In the latter case, users will use a proxy server on the bastion host to open an outgoing connection. This is very similar to the dual-homed architecture which will be discussed in a later slide.

Disadvantages:

- The router is a single point of failure.

- If the bastion host is compromised, there is nothing to protect the rest of the LAN. Allows access to the LAN backbone.

- Less data flow control.


- Is limited to the filtering ability of the router.

Advantages:

- Flexible access controls.

- Depending on security policy could allow direct connectivity from any host.

- Less configuration and hardware requirements which makes for maintainability.



A dual-homed host has at least two network interfaces. It does not act as a router in the firewall architecture. Internal hosts can communicate with a dual-homed host, but not directly to Internet. Hosts on the Internet can only interact with the dual-homed host. Internal users use proxies on the dual-homed host or can log directly into the dual-homed host. (Direct login to the dual-homed host is **NOT** recommended.)

Advantages:

- It keeps Internet access off of interior LAN.

- Allows authorized users to get into the interior LAN via a proxy.

- Allows for a fair amount of functionality when proxies are configured correctly.

- Attacks become focused on the dual-homed host.

Disadvantages:

- It may become a performance bottleneck.

- Relatively complex host configuration.

- Not easy to maintain.

- Two approaches to proxying
 - Circuit level proxy
 - Application level proxy

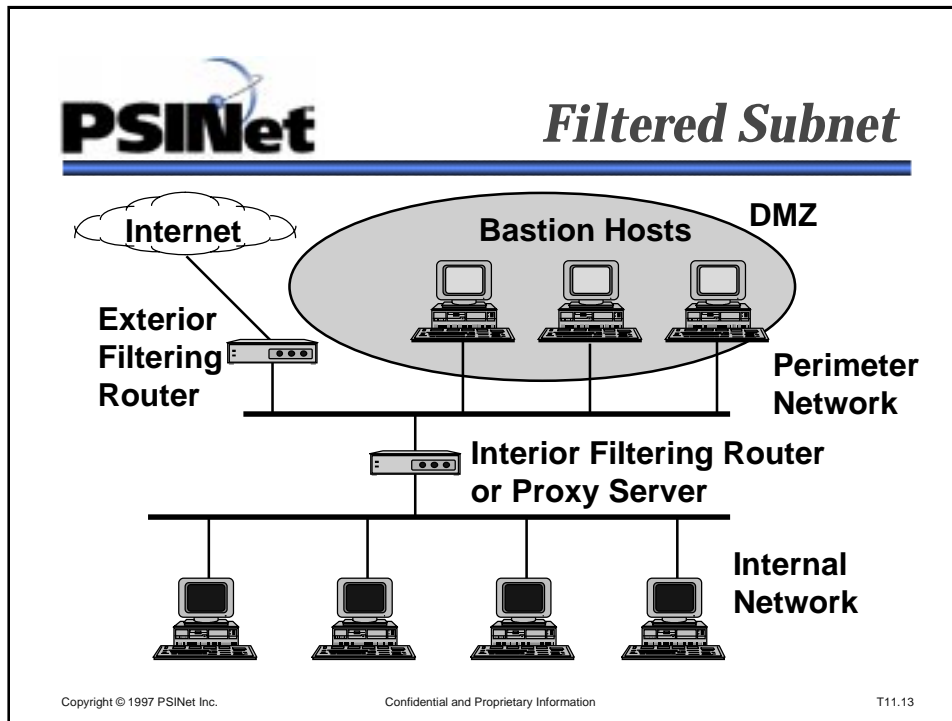
A **circuit level proxy** creates a connection between a client and a server without interpreting the application protocol. A circuit level proxy acts as a relay; a user connects to a TCP port on the proxy which connects to some destination on the other side. Users need to be educated about how to use each client, and not all services can be proxied. This also allows better logging.

Example:

A user wishing to connect to an anonymous FTP server (ftp.cert.org):

1. Using any FTP client, connect to proxy server.
2. At user name prompt, the user needs to specify the name of the real server to connect to: **anonymous@ftp.cert.org**

An **application level proxy** is knowledgeable about the application it is providing proxy services for. It does more than just relay the packet. It examines the packet content and can interpret commands in the application protocol. Special code could be used for each service. In this case, the client software knows how to contact the proxy server and how to tell it what real server to connect to. This method is transparent to the users, but requires that each client be correctly configured.



The filtered subnet architecture has the advantage of bastion hosts that are not on the interior LAN. The bastion hosts are on the perimeter network (also called the demilitarized zone or DMZ). In this configuration, the risk is reduced if a bastion host is compromised since it isn't on the internal network. Packet filtering or proxying could be used on the internal router. This would increase the security of the internal network. An intruder would have to defeat the external router, then the internal router or proxy server. For even more added security, don't route the internal networks.

Additional advantages:

- More hosts spread out the load to reduce bottlenecks.
- Simpler host configuration, although more to configure.
- Bastion hosts are not on the interior LAN.
- Information leakage risk is reduced.

Disadvantages:

- More equipment to buy, maintain, and configure.

A variation of this architecture is to use multiple bastion hosts on a perimeter LAN for different services, merge the interior and exterior router (dual-Ethernet router), and configure filters on each Ethernet interface. Yet another variation is to substitute an application level proxy for the internal router; this would be the most secure and most costly.

- What is packet filtering?
- Why use packet filters?
- What kind of packet filtering is available?
 - Static packet filters
 - Dynamic packet filters

A packet-filtering router examines all IP packets (incoming and/or outgoing) and decides whether or not to pass the packet based on a set of predefined rules. This gives the packet-filtering router the ability to block traffic between specific networks and/or specific hosts, and block specific services to a network and/or host.

Static packet filters have one set of rules that are not easily changed. They either will allow traffic of a certain type between two hosts or will not. If the filters need to be changed, someone needs log into the router and change the filter set.

Dynamic packet filters do everything a static filter will but allow filter rules to be added or deleted based on other criteria. For example, web access can be denied during the work day, but will be permitted off hours. If an off site user needs access to a protected server, the user can authenticate to the router and the router will then install a filter giving the user the necessary access to the internal server.

Packet-filtering is a method of protection against intrusions but by no means guarantees security of a LAN or internal data. Static packet filtering is an inexpensive method of security since most routers have packet-filtering software. Packet-filtering is a perimeter defense that would prevent many direct attacks to internal resources.

Three Steps to implement packet-filtering:

- Step 1
 - Design a security policy.
Know what to permit and what to restrict.
- Step 2
 - Formally define packets that should be permitted.
- Step 3
 - Translate formal definitions to router vendor syntax.

Most services can be identified by specific characteristics of an IP packet. The characteristic most often used is the port number associated with the service. A packet-filtering router examines the characteristics of each IP packet; therefore, specific services can be allowed or denied based on the characteristics of the IP packet.



1. Design a security policy.

- Filtering is based on TCP and UDP port numbers (RFC 1700).
- A security policy defines what traffic will be permitted through the router.
- Good policy
 - Allow a few select services in.
 - Deny all other incoming connections.
 - Allow everything out.

Most policies deal with what service an outsider will be able to access on your LAN. The primary use of filters is to control what outside machines can talk to your machines and to control what services an outsider can access on your machines.

An example of a service that might need to be restricted is SMTP (e-mail). A decision has to be made in accepting mail: should incoming mail be allowed to every host or just a particular host (i.e., SMTP gateway).

For people accessing your LAN use the philosophy:

“All that is not expressly permitted is prohibited.”

For your internal users the above policy is best, but the internal users will want more access to external resources.



Security Policy

- What service do you want to offer to the Internet?
- Are the services you want to offer and access "SAFE?"

If you will be receiving e-mail, you will be offering a e-mail server, and every machine on the Internet needs to have access to your server on the mail port.

It is impossible of have a list of "unsafe" services. There are some services that are known to be problematic to run across the Internet. For example, Xwindows, TFTP, NFS, and the "R" commands are known to be unsafe. It is possible to run "unsafe" services in a safe way if proper precautions are taken. Precautions could be an isolated subnet or a bastion host.



Security Policy

- Define allowable services.
 - Services to allow in (mail, news, WWW, DNS, etc.)
 - Services to allow out (all)
- Determine the interface.
 - Serial or Ethernet
- Design the filter

The services allowed through the filter are those you wish to offer. If you are not running an FTP server, don't allow FTP connections to your internal LAN.

Filters can be put on any interface and have a direction associated with the filter. Filters can be put on the serial interface for traffic coming in from the Internet and for traffic going out to the Internet.

The same thing can be done on the Ethernet interface, but the directions are reversed. When putting filters on multiple interfaces, be very careful. Interaction between filter sets can be complicated, and it is easy to filter yourself out of the router.

Sample Security Policy

- The following security policy is fairly restrictive and a good template for the average site's public services.
- It restricts most packets but allows some of the commonly offered services.
- It also takes into consideration PSINet's SNMP polling feature.
- The policy takes into account known security problems.

Sample Security Policy

- Deny address spoofing, XWindows inbound, source routed packets, TFTP, R commands.
- Allow SMTP inbound to one host.
- Allow WWW inbound to one host.
- Allow FTP inbound to one host.
- Allow TELNET inbound to one host.
- Allow NNTP inbound to one host.
- Allow DNS inbound for name server queries.
- Allow ICMP (Ping).
- Allow SNMP to router from PSINet.
- Allow all established TCP connections and all outbound connections.
- Deny all else.

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T11.20

The scheme of this security policy is to deny all known problems, allow the most commonly used services, and then deny all else.

Address spoofing is when packets that are destined to a network have the source address of that network. The source address is of the same network of the destination address.

XWindows or X11 is the most used windowing system for UNIX machines. It allows for applications to run on a server but to be displayed on another host/terminal.

NFS (Network File System) provides the functionality for hosts to share disks. This merges the remote file systems with the local directory tree to make the remote file system appear local.

Portmapper and RPC (Remote Procedure Call) are procedure calls that are used by distributed (client-server) computing. These calls are requested by a client and executed on the server, but the results are returned to the client over the network. Portmapper processes the request and defines the ports for the RPC request to use.

“R” commands (rlogin, rcp, rsh, rexec) use the BSD authentication mechanism and usually do not require a password to enter a remote system.

Source routed packets are packets that are destined for a host via some other host where the other host is not the local machine.

TFTP (Trivial File Transfer Protocol) is a version of FTP that uses UDP, and does not use any authentication protocol



PSINet *2. Formally define rules.*

- Once the filter is designed, define the rules in terms of IP addresses and port numbers. This is done in step 2.

- A good format for defining rules is:

Action	Source	Port	Destination	Port	Type	Comment
deny	<address>	###	<address>	###	<type>	...
allow	<address>	###	<address>	###	<type>	..

It is a good idea to plan out the filters on paper. The above form is often used. It allows you to check the order of the filters.

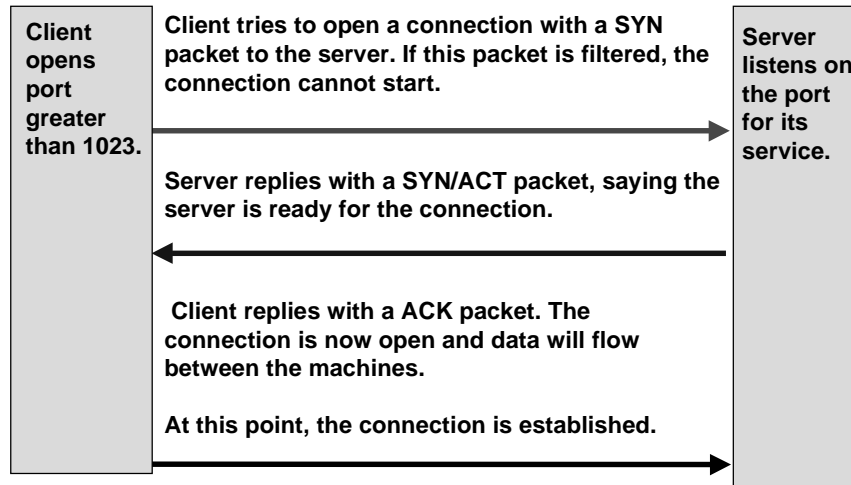
- Things to be aware of when defining rules:
 - Order is important.
 - When a packet meets a rule, the rest of the rules are ignored.
 - Packets flow in both directions!
 - Even if all the data is flowing one way, acknowledgment

Filter expanded

- To deny address spoofing, deny all incoming packets with a source address of the LAN.
- XWindows uses ports 6000 - 6100 (TCP protocol)
- WWW uses port 80 (TCP protocol).
- FTP is complex and is explained on a future
- TELNET uses port 23 (TCP protocol.)
- NNTP uses port 119 (TCP protocol).
- DNS is complex and is explained on a future slide.
- SNMP uses port 161 (UDP protocol).
- All other TCP packets that are acknowledgments to packets that originated within site's LAN.

These are well known port numbers and can be found in RFC 1700. As new services are added to the Internet, port numbers are associated with those services. If you wish to allow a new service through the firewall and the service is not listed in RFC 1700, contact the vendor or look at their web page. Many vendors have a section on running their service through a firewall.

PSINet *How a TCP connection works*



Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T11.25

Filters can be configured to look at the SYN bit. For example, you could choose not to allow any incoming packets with SYN set on port 80, but allow packets with the SYN bit set to port 25 (mail) to your mail server. Filtering on the SYN bit allows you to block incoming connections without blocking your internal machines from connecting to other machines on the Internet.

Most TCP connections follow the diagram above. There are some notable exceptions. The well known exceptions are FTP, Real Audio, iPhone, CUCM, and any application that uses streaming. Real Audio, iPhone, and CUCM all use multiple TCP and UDP ports to send data. FTP uses two (2) TCP connections, one connection for data and one for the control channel.



TCP Based Services

- To allow e-mail connections into your e-mail server:
 - This assumes that you are allowing all packets labeled established in and out.

Action	Source	Port	Destination	Port	Type	Comment
allow	any	any	mail server	25	TCP	incoming e-mail connections

This is an example of what is needed to allow connections to a mail server on the internal LAN. It assumes that all packets from established connections are passed. If packets from established connections are not permitted, more lines would be needed for mail to be delivered.



UDP Based Services

- ❑ Unlike TCP services, UDP does not make use of a SYN bit in its packets.
- ❑ Nearly all UDP based services operate by accepting requests on the service port and responding to some port above 1023 as specified by the request.
- ❑ To allow customers to use a UDP service outside their site, you need to allow the return packets back in from the service port.

UDP is a simpler kind of protocol. There is no SYN bit and traffic is simply sent out. There is no method to make sure that a packet reaches its destination. This makes services harder to block. Sites should be careful about what services they offer on UDP ports and what services they allow their users to access.



UDP Based Services

- ❑ To allow customers to initiate “talk” sessions to other sites, you must allow the returning packets back in. “Talk” uses port 517.

Action	Source	Port	Destination	Port	Type	Comment
allow	any	517	internal net	+1023	UDP	returning “talk” packets

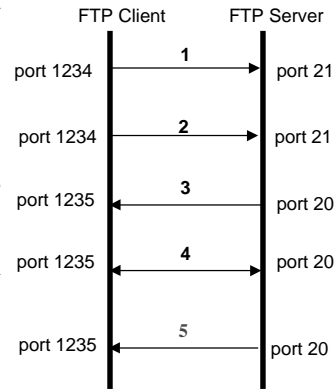



UDP Based Services

- ❑ If you were not allowing all UDP packets outbound, you would have to allow “talk” specifically:

Action	Source	Port	Destination	Port	Type	Comment
allow	internal net	any	any	517	UDP	outbound “talk” packets

1. Client opens a TCP connection from a port >1024 to port 21 on the server. This is the “command channel.”
2. User issues a get or put command.
3. Server opens a connection from port 20 to the next available port > 1023.
4. Data is passed on this channel, so it is known as the “data channel.”
5. Once data is finished transmitting, the channel is closed. The user can still send commands on the connection to port 21.



FTP is a special application. FTP uses two separate TCP connections.

The first connection on port 21 controls the connection.

The second channel on port 20 sends and receives the files.

FTP is the only “old” TCP application that uses multiple ports. In the future such applications will become more common.



FTP Rules

- ❑ These FTP rules allow connections to an internal FTP server and connections to an external FTP server.

To allow full FTP access to the outside world and the outside world access to the local FTP machine requires four rules, one rule for each port in each direction.

- Line 1 allows connections to an internal FTP server.

Action	Source	Port	Destination	Port	Type	Comment
allow	any	any	internal FTP server	21	TCP	connection to server (1) IN

- ❑ Line 2 allows data in from a connection to an external FTP server.

Action	Source	Port	Destination	Port	Type	Comment
allow	any	20	internal net	+1023	TCP	data channel (2) IN

- ❑ Line 3 allows a connection from internal hosts to an external FTP server.

Action	Source	Port	Destination	Port	Type	Comment
allow	internal net	any	any	21	TCP	connection to server (3) OUT

- ❑ Line 4 allows a data channel outbound.

Action	Source	Port	Destination	Port	Type	Comment
allow	internal FTP server	20	any	+1023	TCP	data channel (4) OUT

- ❑ The DNS service uses both TCP and UDP protocols on port 53.
- ❑ UDP is for client-to-server queries. Lookups.
 - ❑ Therefore, UDP to port 53 should always be passed.
- ❑ TCP is for server-to-server transfers. Zone transfers.
 - ❑ If the site is running an authoritative name server, then TCP on port 53 should be passed.

Action	Source	Port	Destination	Port	Type	Comment
allow	any	53	any internal	+1023	UDP	returning DNS requests
allow	secondary	any	name server	53	TCP	zone transfer requests
allow	any	any	name server	53	UDP	name queries

If you are setting up a primary name server, it is often helpful to allow zone transfers to any outside machine. This makes debugging much faster. After the server is up and running, the filters should only allow zone transfers from the necessary servers.



3. Convert to Vendor Syntax

- Router vendor manual
- Try PSINet Filter Builder
 - <http://www.support.psi.net/fbuilder>

- Computer Emergency Response Team (CERT)
 - cert@cert.org (412)268-7090
 - <http://info.cert.org/>
- COAST
Computer Operations, Audit, and Security Technology
 - <http://www.cs.purdue.edu/coast/coast.html>
- Be prepared!
 - Have contingency plans.
 - What stance will you take to an intrusion?
 - Open contacts with proper authorities.
 - It can happen. If you don't think so, it will happen.

Books to get:

Firewalls and Internet Security by Cheswick & Bellovin

Building Internet Firewalls by Chapman & Zwicky

Software to check out:

Screend

Filtering capabilities to a UNIX based system.
<ftp://coast.cs.purdue.edu/pub/tools/unix/screend>

Drawbridge

Turns a DOS-PC into a packet filtering machine.
<ftp://net.tamu.edu/pub/security/TAMU>

TCP Wrapper

Monitor and filter requests for servers started in inetd.
<ftp://ftp.win.tue.nl/pub/security/>

