

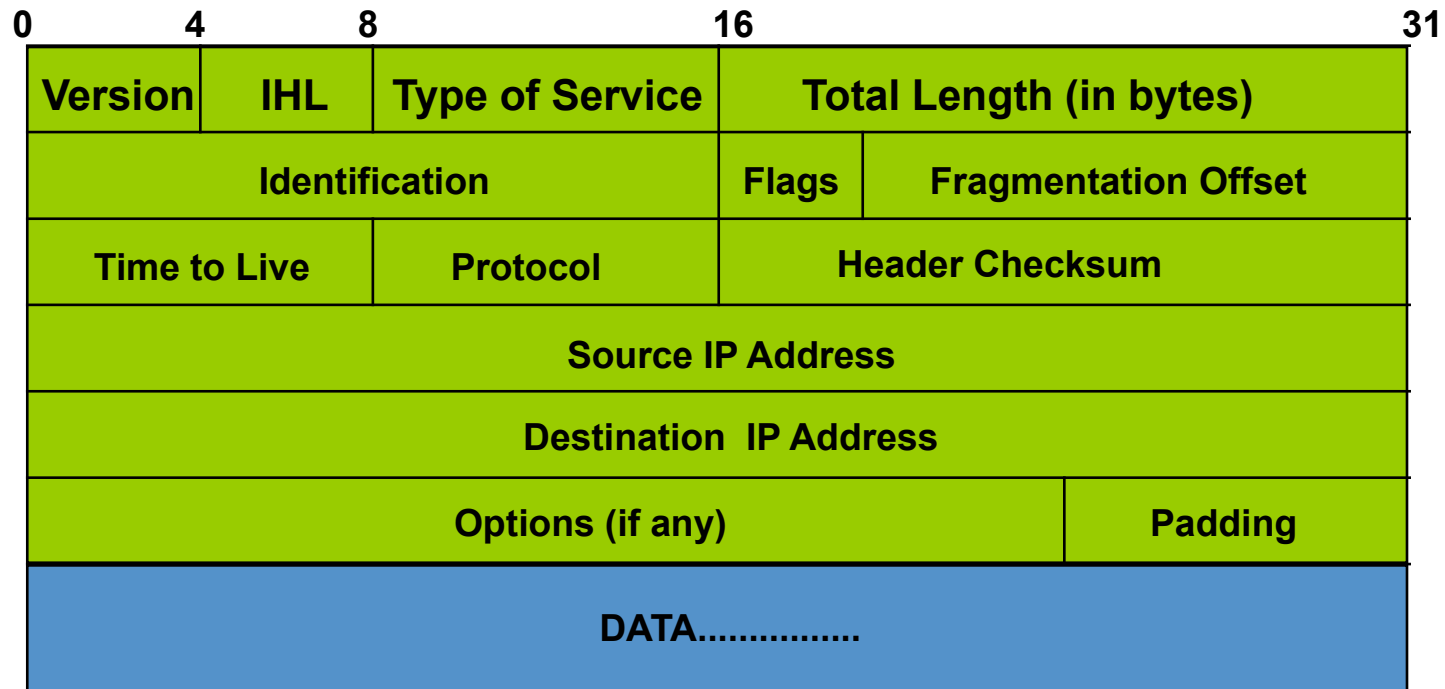
Filtering and Firewalls

Sorting Through FUD to get Sanity

Merike Kaeo

merike@doubleshotsecurity.com

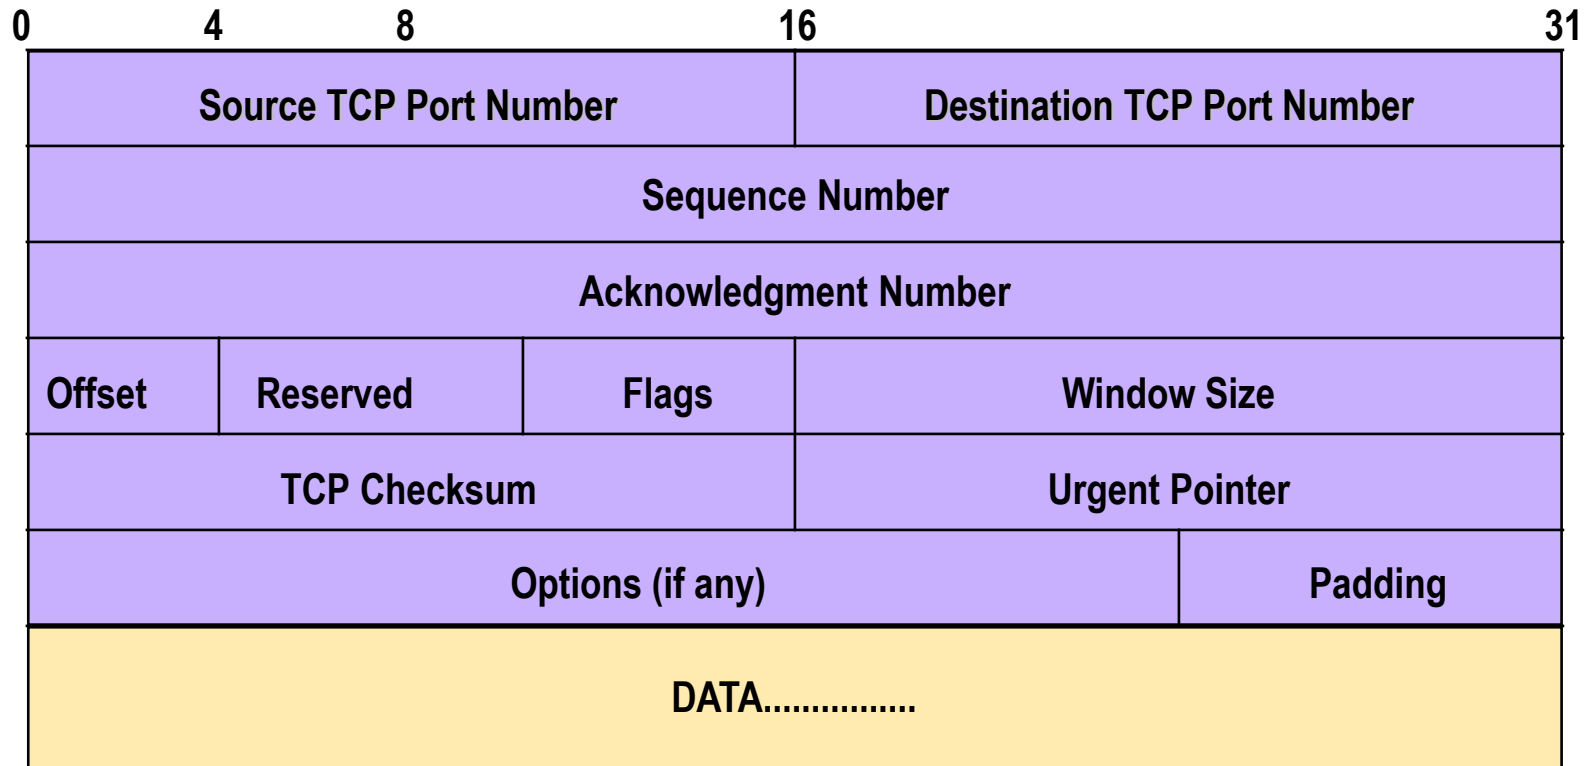
IPv4 Protocol Header



TCP (Transport Control Protocol)

- Provides reliable virtual circuits to user processes
- Lost or damaged packets are resent
- Sequence numbers maintain ordering
- All packets except first contain ACK #
(ACK# = sequence number of last sequential byte successfully received)

TCP Header Format



TCP Control Flags



URG	ACK	PSH	RST	SYN	FIN
------------	------------	------------	------------	------------	------------

- **URG**: indicates urgent data in data stream
- **ACK**: acknowledgement of earlier packet
- **PSH**: flush packet and not queue for later delivery
- **RST**: reset connection due to error or other interruption
- **SYN**: used during session establishment to synchronize sequence numbers
- **FIN**: used to tear down a session


TCP Session



Connection
Initialization



Connection
Termination



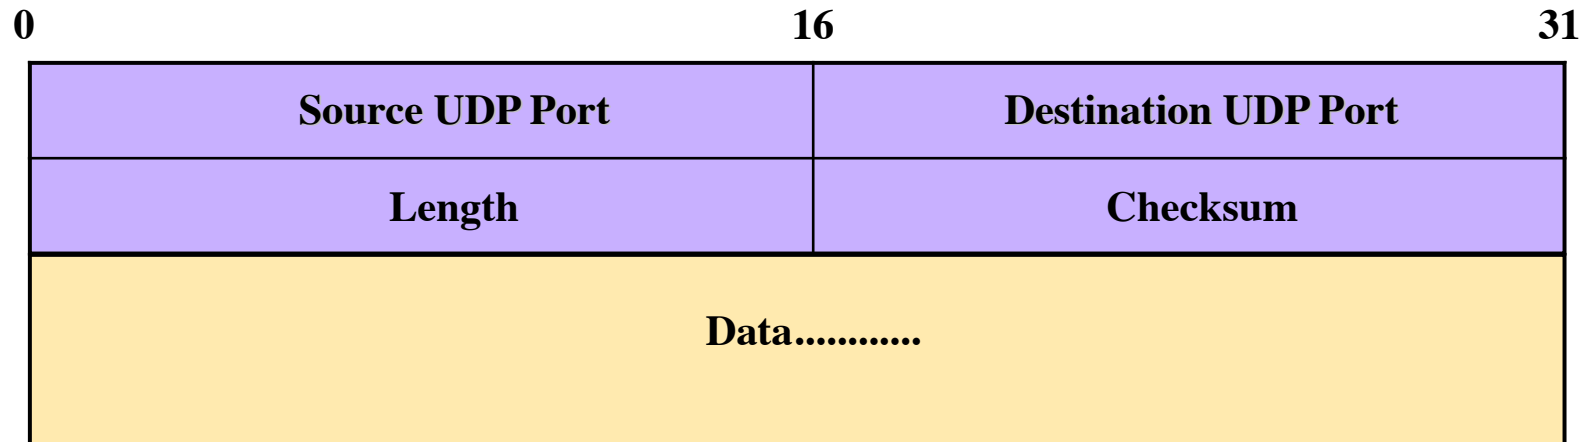
TCP Port Numbers

- Port numbers < 1024 are privileged ports
- Destination port is fixed
- Source port is randomly generated

UDP (User Datagram Protocol)

- **Delivery is on a best-effort basis**
 - No error correction
 - No retransmission
 - No lost, duplicate, re-ordered packet detection
- **Easier to spoof than TCP packets**
 - no handshake
 - no sequence numbers

UDP Header Format



ICMP

- **Transmits command and control information**
 - ICMP Echo
 - determines whether another system is alive
 - ICMP Destination Unreachable
 - No route to destination
 - ICMP Source Quench
 - Slow down number of packets sent

ICMP

- **IP Hdr and first 64 bits of transport header**
 - Included in ICMP Message
 - Limits scope of changes dictated by ICMP
 - Older implementations do not use this info
 - Destination Unreachable messages can affect all connections between a pair of hosts
 - Redirect messages should only be obeyed by hosts (from router or directly connected network)

ICMP Message Types

Message Type	Value	Description
Echo Reply	0	Ping response if system alive
Destination Unreachable	3	Earlier IP message not deliverable
Source Quench	4	Packets received too fast to process
Redirect	5	Traffic should be directed to another router
Echo	8	Send a ping
Time Exceeded	11	Max # of hops in TTL field is exceeded
Parameter Problem	12	Bad parameter in header field
Timestamp	13	Includes time on sending machine and requests time on destination machine
Timestamp Reply	14	Timestamp response
Information Request	15	Used by host to determine which network it is on
Information Reply	16	Contains response to information request

IP Fragments

- Only first fragmented packet contains port number information
- Firewall should have capability of fragment reassembly

How Do We Control Traffic ?

- **Firewalls**
 - Simple Rule-Based
 - Proxy
 - Stateful
- **Which One Is Needed ?**
- **Where Do I Put It ?**
- **What Do I Configure ?**

Firewall Cost Tradeoff

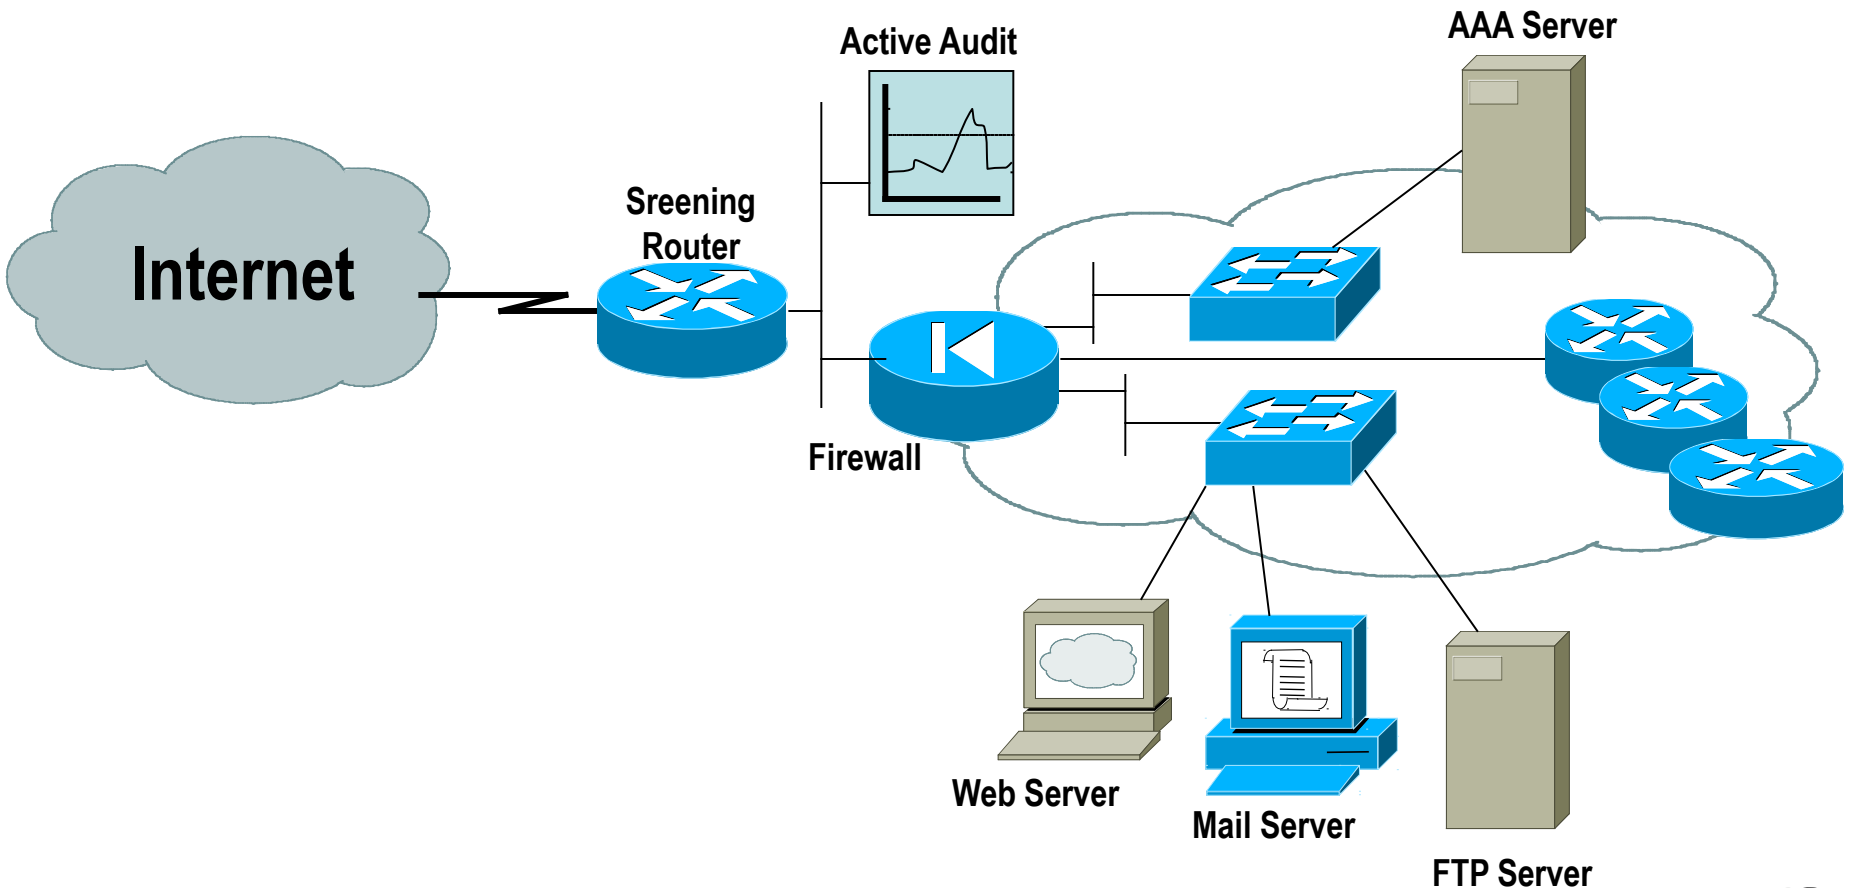
USING A FIREWALL

- Hardware cost and maintenance
- Software purchase and updates
- Administrative setup and training
- Lost business from blocked service
- Loss of some service

NOT USING A FIREWALL

- Effort spent dealing with break-ins
- Legal costs

Typical Secure Infrastructure Architecture



Filtering Recommendations

- Log filter port messages properly
- Allow only internal addresses to enter the router from the internal interface
- Block packets from outside (untrusted) that are obviously fake or commonly used for attacks
- Block packets that claim to have a source address of any internal (trusted) network.

Filtering Recommendations

- Block incoming loopback packets and RFC 1918 networks
 - 127.0.0.0
 - 10.0.0.0 – 10.255.255.255
 - 172.16.0.0 – 172.31.0.0
 - 192.168.0.0 – 192.168.255.255
- Block multicast packets (if NOT using multicast)
- Block broadcast packets (careful of DHCP and BOOTP users)
- Block incoming packets that claim to have same destination and source address

DoS Filtering (RFC 3330)

(* these networks may be reallocated)

Description	Network
default	0.0.0.0 /8
loopback	127.0.0.0 /8
RFC 1918	10.0.0.0 /8
RFC 1918	172.16.0.0 /12
RFC 1918	192.168.0.0 /16
Net Test	192.0.2.0 /24
Testing devices *	192.18.0.0 /15
IPv6 to IPv4 relay *	192.88.99.0 /24
RFC 1918 nameservers *	192.175.48.0 /24
End-node auto configuration *	169.254.0.0 /16

IPv4 DoS Filtering - Current

- <http://tools.ietf.org/html/rfc3330>

Obsoleted by:

- <http://tools.ietf.org/html/rfc5735>

Updated by:

- <http://tools.ietf.org/html/rfc6598>

Email Spam Sources

- Open relays and proxies
- Compromised machines
- Direct Spam sources
- Insecure Webmail interfaces / Perl scripts

Preventing Outbound SPAM

- Scan network for open relays and proxies
- Block compromised hosts until fixed
- Block outbound port 25 for dynamic IP addresses
- Filter inbound access to known proxy ports

Filtering Inbound SPAM

- Check SMTP headers
- Build DNS block lists (DNSBLs)
- HELO filtering
- Use SPAM filters (Spamassassin, Razor)
- Block routes to major spammers

EMAIL (SMTP) Filtering

- **Sample SMTP Filtering**
 - Permit outgoing traffic to port 25
 - Permit incoming traffic from port 25
 - Permit our trusted hosts with dst port 25
 - Permit all other traffic with src port 25 and ACK flag set (the reply)

Defining Filtering Rules (SMTP)

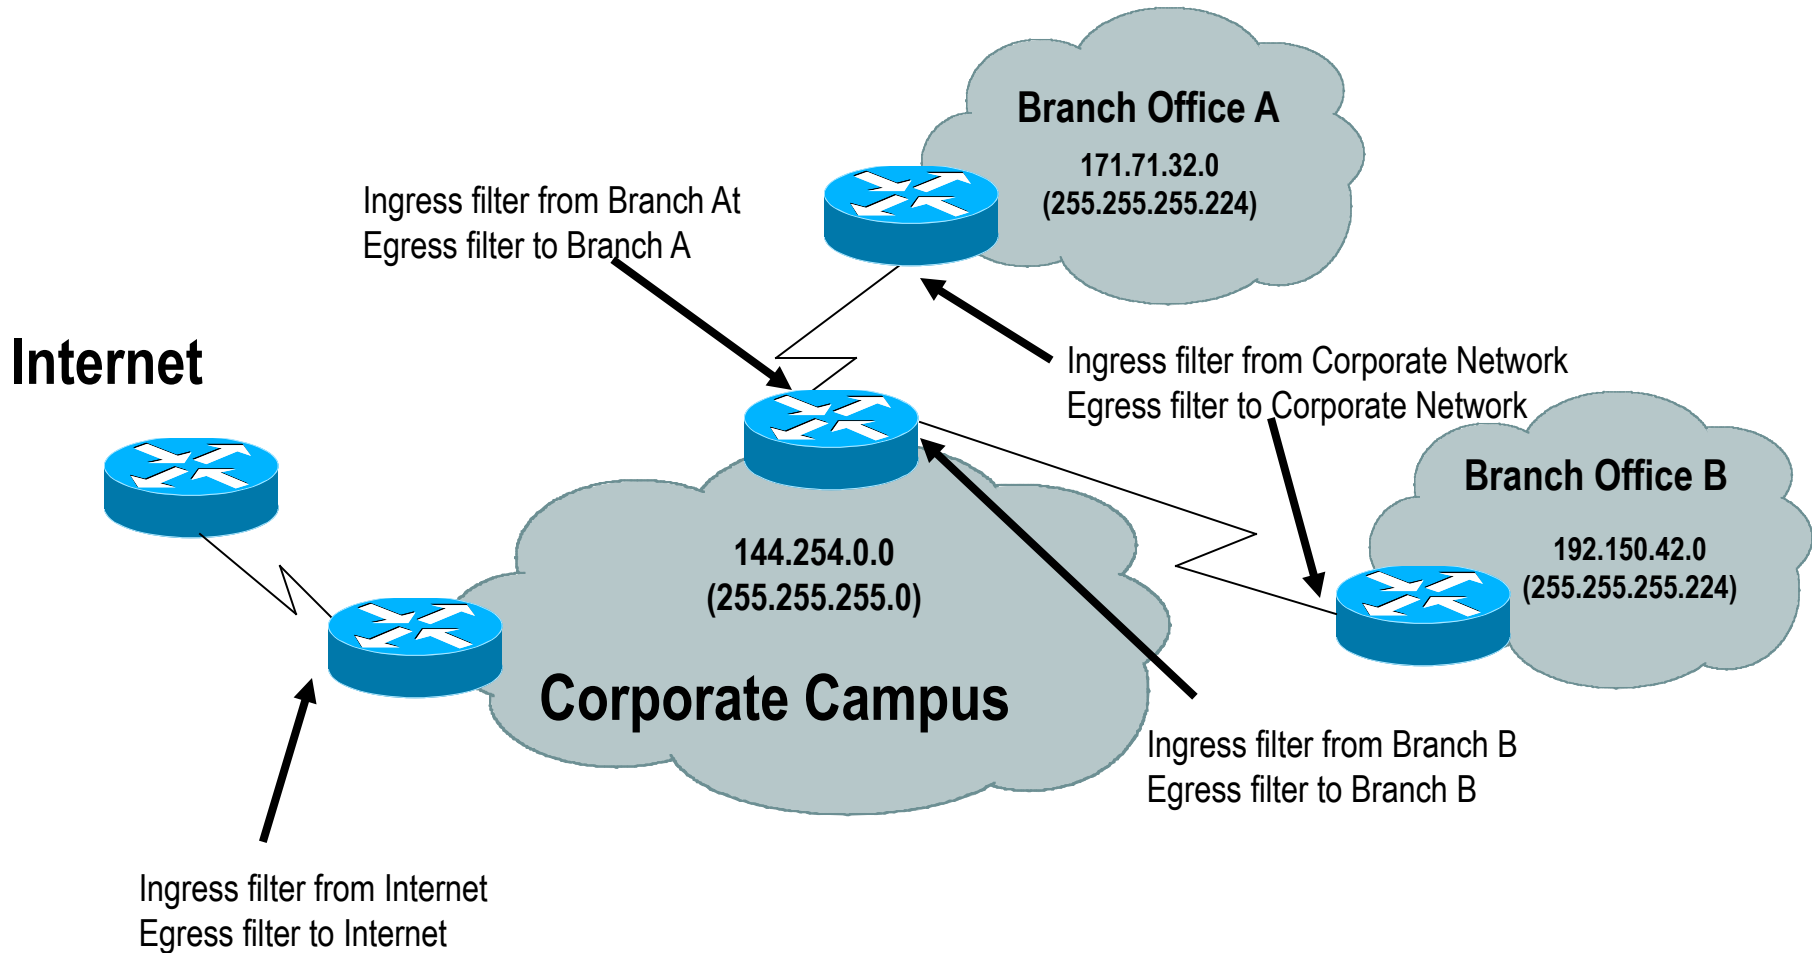
Direction	SRC IP address	DST IP address	Protocol	SRC Port	DST Port	ACK set	Description
In	External	Internal	TCP	>1023	25	*	Incoming mail Sender to recipient
Out	Internal	External	TCP	25	>1023	Yes	Incoming mail Recipient to sender
Out	Internal	External	TCP	>1023	25	*	Outgoing mail Sender to recipient
In	External	internal	TCP	25	>1023	Yes	Outgoing mail Recipient to sender

* ACK not set on first packet but set on all subsequent packets

Filtering Issues

- **Ordering**
 - What sequence is packet inspected in?
- **Performance**
 - Are there any limitations?
- **Logging**
 - Get appropriate information
 - Timestamps

Simple Filtering Example



Branch Router Policy

- **Ingress filtering:**
 - deny all rfc 1918 and special use addresses from entering the branch network
 - deny all traffic with an IP source address that matches the branch network address allocation
 - permit all other traffic
- **Egress filtering:**
 - permit only traffic with an IP source address that matches the branch network
 - deny all other traffic

Branch A Router Configuration

```
access-list 133 deny ip host 0.0.0.0 any
access-list 133 deny ip 127.0.0.0 0.255.255.255 any
access-list 133 deny ip 10.0.0.0 0.255.255.255 any
access-list 133 deny ip 172.16.0.0 0.15.255.255 any
access-list 133 deny ip 192.168.0.0 0.0.255.255 any
access-list 133 deny ip 192.0.2.0 0.0.0.255 any
access-list 133 deny ip 169.254.0.0 0.0.255.255 any
access-list 133 deny ip 240.0.0.0 15.255.255.255 any
access-list 133 deny ip 171.71.32.0 0.0.0.31 any
access-list 133 permit ip any any
```

```
access-list 144 permit ip 171.71.32.0 0.0.0.31 any
access-list 144 deny ip any any
```

```
interface BRI0
description To Corporate Network
ip access-group 133 in
ip access-group 144 out
```

- Modify according to current special use address standards
- You will periodically need to check whether allocations have been updated and/or modified

NAS Router Policy

- **Ingress filtering:**
 - permit only traffic with an IP source address of branch networks
 - deny all other traffic
- **Egress filtering:**
 - deny all rfc 1918 and special use addresses from propagating to branch networks
 - deny all traffic with an IP source address that matches the branch network address allocation
 - permit all other traffic

NAS Router Configuration

```
access-list 133 permit ip 171.71.32.0 0.0.0.31 any
access-list 133 permit ip 192.150.42.0 0.0.0.31 any
access-list 133 deny ip any any
```

```
access-list 144 deny ip host 0.0.0.0 any
access-list 144 deny ip 127.0.0.0 0.255.255.255 any
access-list 144 deny ip 10.0.0.0 0.255.255.255 any
access-list 144 deny ip 172.16.0.0 0.15.255.255 any
access-list 144 deny ip 192.168.0.0 0.0.255.255 any
access-list 144 deny ip 192.0.2.0 0.0.0.255 any
access-list 144 deny ip 169.254.0.0 0.0.255.255 any
access-list 144 deny ip 240.0.0.0 15.255.255.255 any
access-list 144 deny ip 171.71.32.0 0.0.0.31 any
access-list 144 deny ip 192.150.42.0 0.0.0.31 any
access-list 144 permit ip any any
```

```
interface Serial 0:23
description To Branch Offices
ip access-group 133 in
ip access-group 144 out
```

- Modify according to current special use address standards
- You will periodically need to check whether allocations have been updated and/or modified

Internet Router Policy

- **Ingress filtering:**
 - deny all rfc 1918 and special use addresses from entering the corporate network
 - deny all traffic with an IP source address of the corporate network or branch networks
 - permit all other traffic
- **Egress filtering:**
 - permit only traffic with an IP source address of the corporate network and branch networks
 - deny all other traffic

Internet Router Configuration

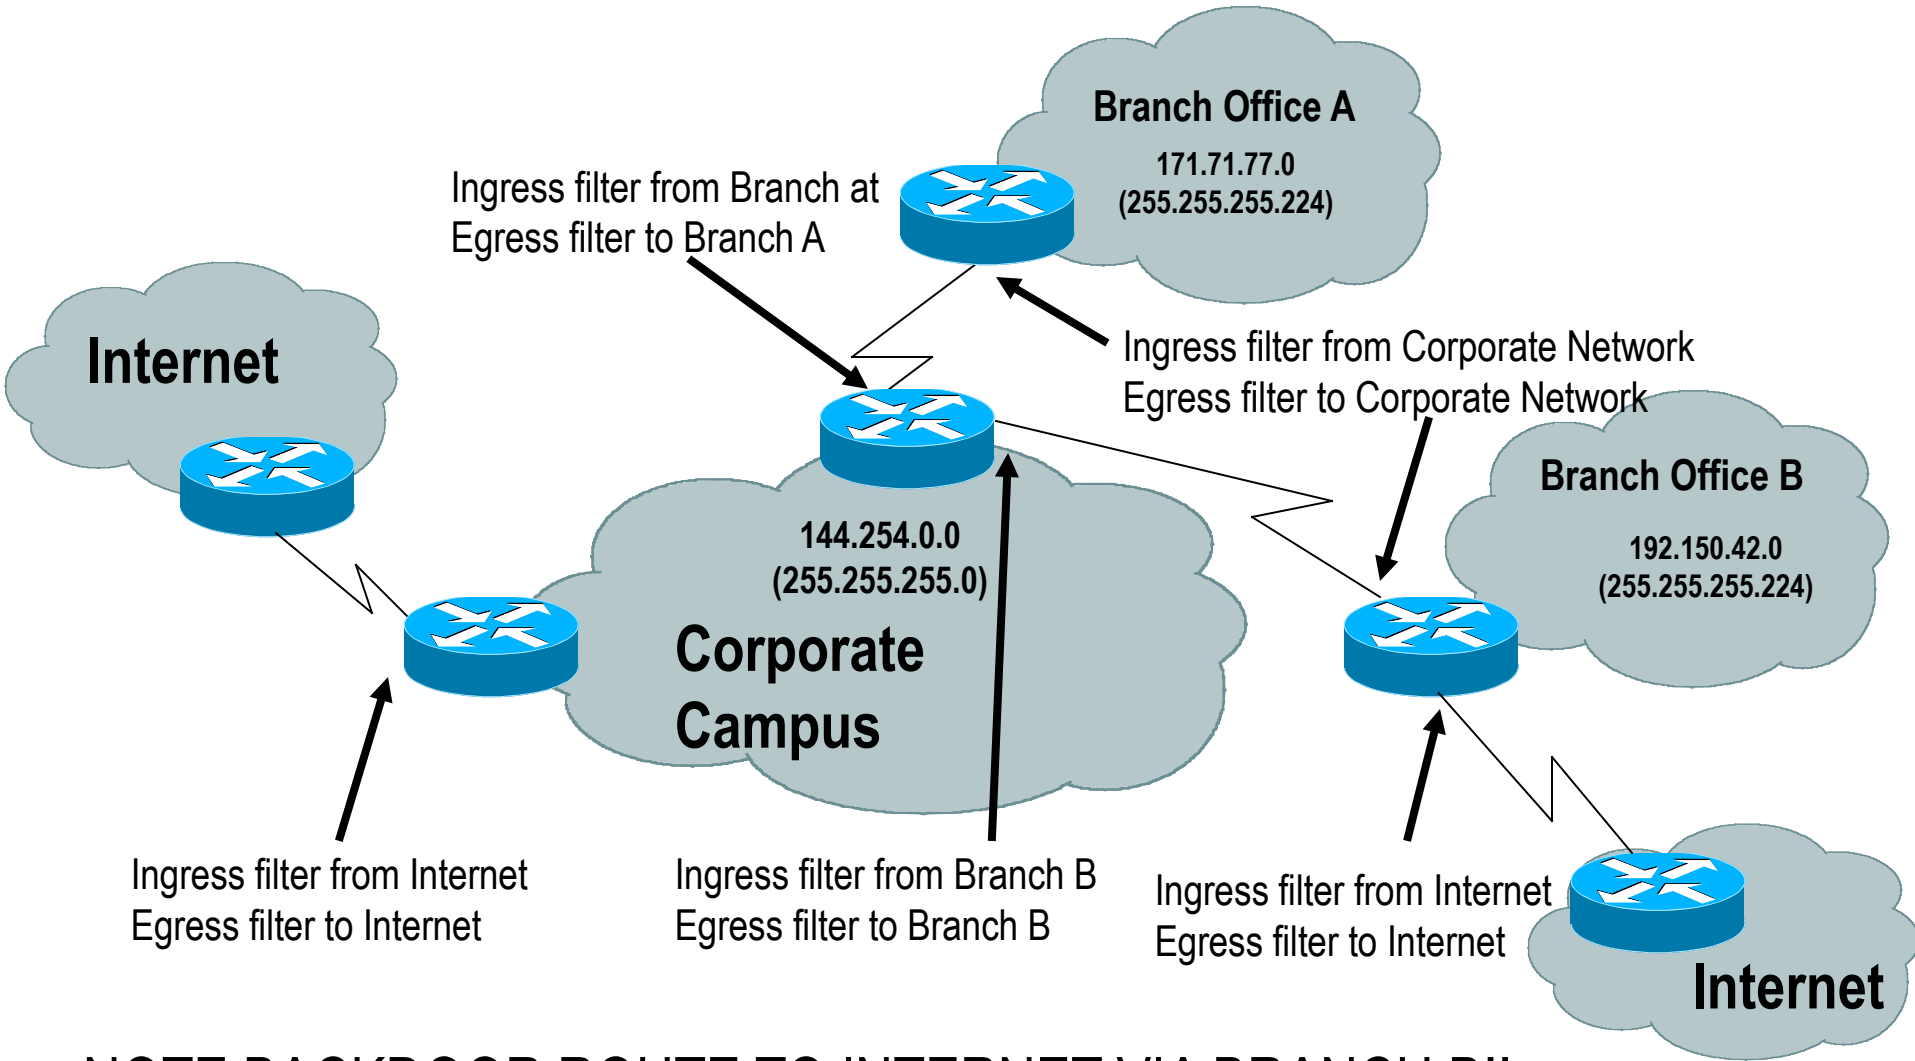
```
access-list 133 deny ip host 0.0.0.0 any
access-list 133 deny ip 127.0.0.0 0.255.255.255 any
access-list 133 deny ip 10.0.0.0 0.255.255.255 any
access-list 133 deny ip 172.16.0.0 0.15.255.255 any
access-list 133 deny ip 192.168.0.0 0.0.255.255 any
access-list 133 deny ip 192.0.2.0 0.0.0.255 any
access-list 133 deny ip 169.254.0.0 0.0.255.255 any
access-list 133 deny ip 240.0.0.0 15.255.255.255 any
access-list 133 deny ip 144.254.0.0 0.0.255.255 any
access-list 133 deny ip 171.71.32.0 0.0.0.31 any
access-list 133 deny ip 192.150.42.0 0.0.0.31 any
access-list 133 permit ip any any
```

```
access-list 144 permit ip 144.254.0.0 0.0.255.255 any
access-list 144 permit ip 171.71.32.0 0.0.0.31 any
access-list 144 permit ip 192.150.42.0 0.0.0.31 any
access-list 144 deny ip any any
```

```
interface Serial 0/0
description To Internet
ip access-group 133 in
ip access-group 144 out
```

- Modify according to current special use address standards
- You will periodically need to check whether allocations have been updated and/or modified


Advanced Filtering Example



NOTE BACKDOOR ROUTE TO INTERNET VIA BRANCH B!!


Packet Filtering Firewall

- Examines the source and destination address of the data packet and either allows or denies the packet from traveling the network
- Blocks access through the firewall to any packets, which try to access ports which have been declared "off-limits"




Application Layer Firewall

- Also known proxy firewalls, application gateway
- attempts to hide the configuration of the network behind the firewall by acting on behalf of that network/servers
- All requests for access are translated at the firewall so that all packets are sent to and from the firewall, rather than from the hosts behind the firewall



Stateful Inspection Firewall

- Examines the state and the context of the packets
- Remembers what outgoing requests have been sent and only allow responses to those requests back through the firewall
- Attempts to access the internal network that have not been requested by the internal network will be denied



Only allows reply packets for requests made out
Blocks other unregistered traffic

Firewall BCP


- Explicitly deny all traffic and only allow what you need
- The default policy should be that if the firewall doesn't know what to do with the packet, deny/drop it
- Don't rely only on your firewall for all protection of your network
- Implement what's called "defense in depth." - multiple layers of network protection
- Make sure all of the network traffic passes through the firewall
- Log all firewall exceptions (if possible)

Some NANOG List Threads

- <http://mailman.nanog.org/pipermail/nanog/>
- ISP Port Blocking Practice (Fall 2009)
- DDoS Mitigation HW/SW (Early 2010)
- I Don't Need No Stinking Firewall (Early 2010)
- Acknowledgment to all whose comments I am incorporating

ISP Port Blocking Practice

- **Started as question**
 - Have outbound rule to drop packets with specific destination port and inbound rule to drop packets with specific inbound ports, which rule mostly used?
 - Are rules based on SYN or SYN-ACK flags?



ISP Port Blocking Practices

1. Never allow traffic to egress (i.e. exit) any subnet unless its source IP address is within that subnet range.
 - uRPF should be used along with ACLs
 - uRPF works best on egress but does little on outside ingress (i.e. bogons)
 - Unless you have implemented an automated s/RTBH|sink and Cymru bogons(learnt via peering) on a trigger box, pushed in through a route-map tagged with the null-route community to the PE. Works magic.
2. Never allow traffic to egress any subnet, if that traffic claims to originate from the subnet's network number or broadcast address.

ISP Port Blocking Practices

4. Never allow traffic to ingress any subnet, if that traffic is directed to the subnet's network number or broadcast address.
5. Never allow traffic to ingress any network if the source address is bogus.
 - a. Never flag a source address as bogus unless you can verify it is bogus*today*, not when you installed the filter. Out of date bogon filters are evil.
6. Never allow traffic to ingress or egress any network if it has a protocol not "supported" by your network (e.g., allow only TCP, UDP, ICMP, ESP, AH, GRE, etc.).
7. Never allow traffic to ingress or egress any network if it has an invalid TCP flags configuration.

Filtering SNMP

- Use 587 (with SMTP Auth and SSL/TLS) - very few places block or proxy it.
- You only allow your customers to connect to your SMTP server, and if they attempt to connect to *any* other SMTP server, they are blocked or redirected to your SMTP server.

I Don't Need No Stinking Firewall

- Started as question asking when firewall vs stateful inspection useful
- Definition (s)
 - A ***firewall is anything that examines IP packets in-line*** for the purpose of discarding undesirable packets before they can be interpreted by the transport layer protocol (e.g. TCP) on the endpoint computer.
 - A ***stateful firewall performs bidirectional classification of communications between nodes***, and makes a pass/fail determination on each packet based on a) whether or not a bidirectional communications session is already open between the nodes and b) any policy rules configured on the firewall as to what ports/protocols should be allowed between said nodes.

Stateful Firewalls

- **Placement**
 - makes good sense in front of machines which are primarily clients; the stateful inspection part keeps unsolicited packets away from the clients.
 - makes ***absolutely no sense in front of servers***, given that by definition, every packet coming into the server is unsolicited (some protocols like ftp work a bit differently in that there're multiple bidirectional/omni directional communications sessions, but the key is that the initial connection is always unsolicited)

Truth or Myth of Stateful FWs?

(1) Security in depth. In an ideal world every packet arriving at a server would be for a port that is intended to be open and listening. Unfortunately ports can be opened unintentionally on servers in several ways: sysadmin error, package management systems pulling in an extra package which starts a service, etc. By having a firewall in front of the server we gain security in depth against errors like this.

- Stateless ACLs in router/switch hardware capable of handling large amounts of pps takes care of this.

Truth or Myth of Stateful FWs?

- (2) Centralized management of access controls. Many services should only be open to a certain set of source addresses. While this could be managed on each server we may find that some applications don't support this well, and management of access controls is then spread across any number of management interfaces. Using a firewall for network access control reduces the management overhead and chances of error. Even if network access control is managed on the server, doing it on the firewall offers additional security in depth.
- Stateless ACLs in router/switch hardware capable of handling large amounts of pps takes care of this.

Truth or Myth of Stateful FWs?

(3) Outbound access controls. In many cases we want to stop certain types of outbound traffic. This may contain an intrusion and prevent attacks against other hosts owned by the organisation or other organisations. Trying to do outbound access control on the server won't help as if the server is compromised the attacker can likely get around it.

- Stateless ACLs in router/switch hardware capable of handling large amounts of pps takes care of this.

Truth or Myth of Stateful FWs?

(4) Rate limiting. The ability to rate limit incoming and outgoing data can prevent certain sorts of DoSes.

- Rate-limiting during a DDoS - i.e., an attack against state and *capacity* is absolutely the *worst* thing one can possibly do, in almost all circumstances.

Truth or Myth of Stateful FWs?

(5) Signature based blocking. Modern firewalls can be tied to intrusion prevention systems which will 'raise the shields' in the face of certain attacks. Many exploits require repeated probing and this provides a way to stop the attack before it is successful.

- Signatures are obsolete before they're ever created; 15 years of firewalls and so-called IDS/'IPS', and the resultant hundreds of millions of botnet hosts prove this point

DoS Protection

- 100Mb/s will carry 148,800 pps worth of 64byte packets
 - A fairly fast firewall can support 100k new connections a second
 - Same firewall can probably forward 2-3mpps when it comes to small packets and will run out of state long before running out of forwarding horsepower.
- Pentium equivalent or low to mid-range mips might support a rate of 2-10k connections per second at which point the threshold for DoSing it based on session rate is quite a bit lower (quite a bit lower than that of a webserver or desktop pc for example)
- Fronting one's Web server farms/load-balancers with a tier of transparent reverse-proxy caches is a better way to scale TCP connection capacity

Rate Limiting Useless or Useful?

- It may be good practice to rate limit outgoing ICMP PING replies from your server to the real world. It's being a good neighbor in the event of certain types of attacks on other parties.
- This can be extended into more specific types of outgoing rate limits. For example, an ISP DNS recursive server that normally serves 1Mbps of traffic in aggregate but lives on a 1Gbps Ethernet might use a per-destination outgoing limit to restrict the amount of damage that could be inflicted on a remote DNS server (without affecting other destinations); things like FreeBSD ipfw/dummynet have these sorts of capabilities.

Value of Firewalls

- **The primary value of a firewall:**
 - It enables a network administrator to define his "edge", the interior of which he is responsible for.
 - It enables a network administrator to isolate his network from externally-originated traffic per a defined security policy.
- **A statefull firewall is most useful for *outbound* traffic, inbound traffic controls usually break things that depend on maintaining state.**
- **Simple ACLs can keep traffuc out or in. Stateful things are only needed when you want to keep track of things you sent outbound, so you can let (hopefull) the same thing back inbound.**
- **Remember to audit firewall exception rules**

Personal Observations

- A firewall by itself != security
- ISPs are not the Internet police
 - But they need to protect themselves from misbehaving upstream/downstream traffic
- Know your hardware/software limitations
- Don't create single point of failure