CS 419: Computer Security

Weeks 10-11: Network Security
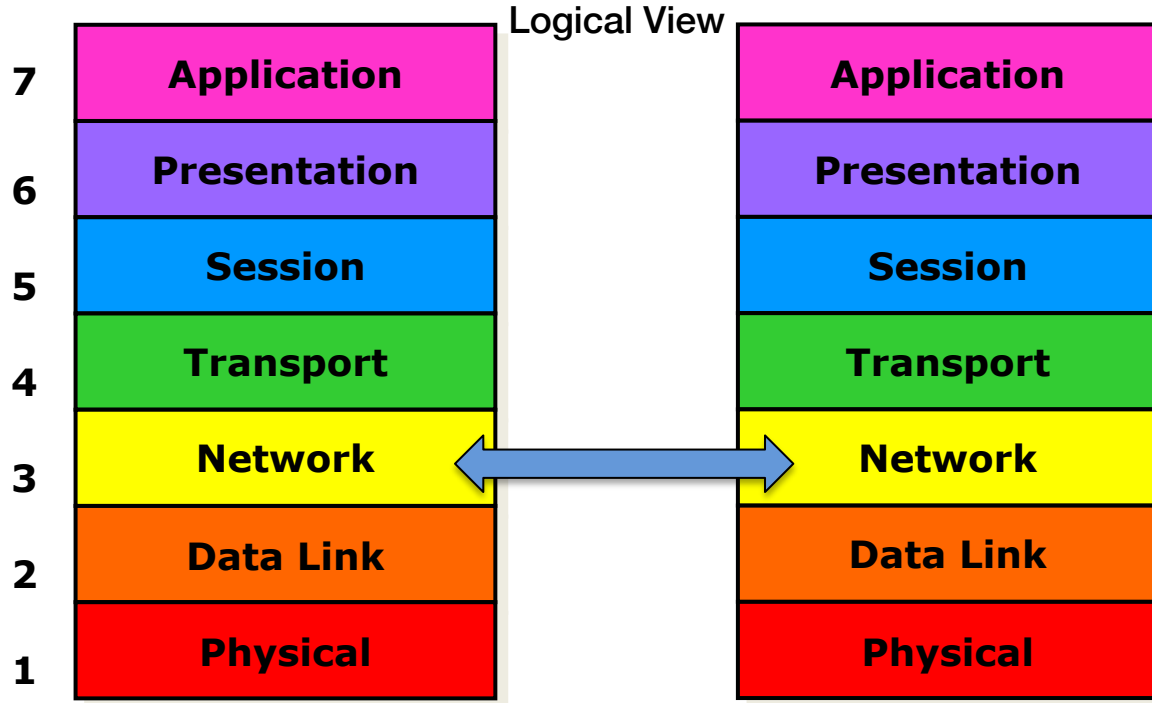
Lecture Notes
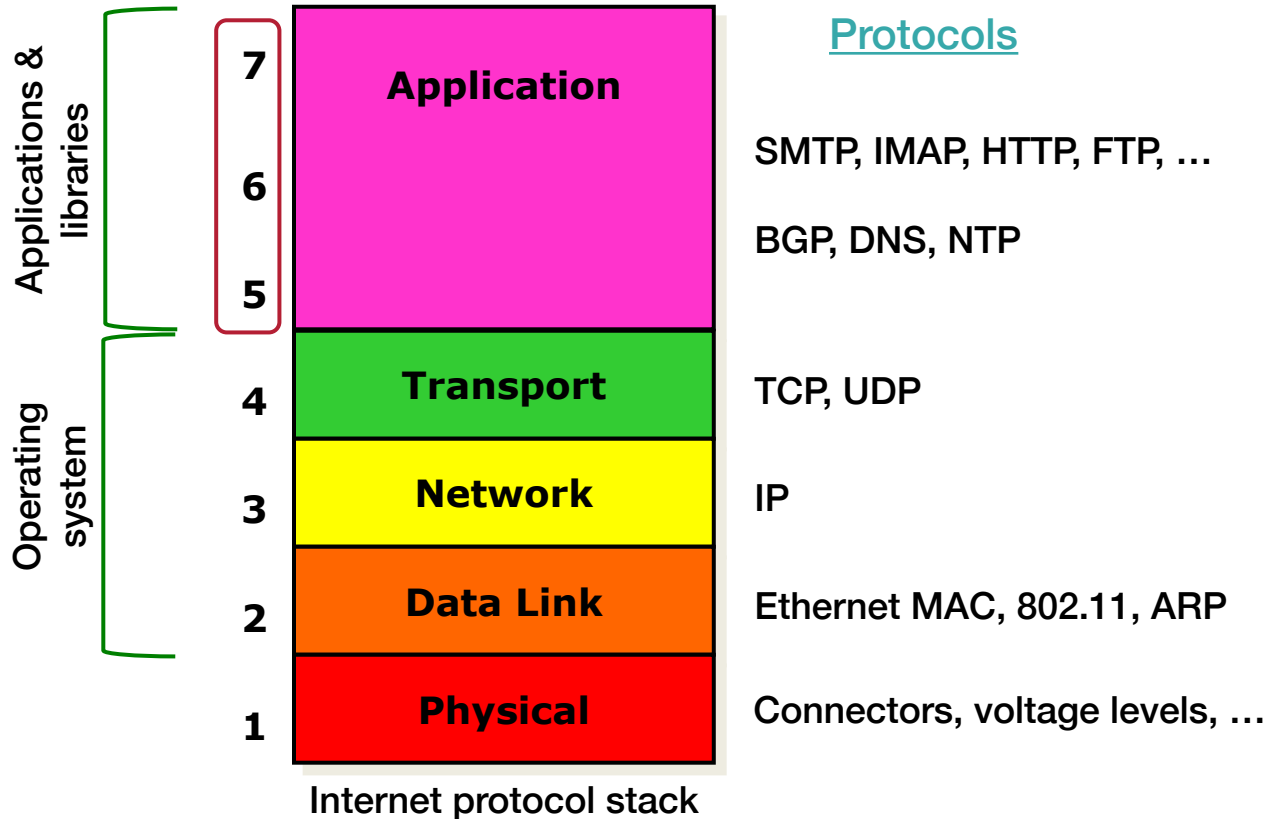
Paul Krzyzanowski

Network protocol layers

Networks are modular. Protocol layers communicate with their counterparts. Low-level attacks can affect higher levels.

Logical View

	Left		Right
7	**Application**		**Application**
6	**Presentation**		**Presentation**
5	**Session**		**Session**
4	**Transport**		**Transport**
3	**Network**	⟷	**Network**
2	**Data Link**		**Data Link**
1	**Physical**		**Physical**

IP Protocol Stack

		Protocols
7	**Application**	
		SMTP, IMAP, HTTP, FTP, …
6		
		BGP, DNS, NTP
5		
4	**Transport**	TCP, UDP
3	**Network**	IP
2	**Data Link**	Ethernet MAC, 802.11, ARP
1	**Physical**	Connectors, voltage levels, …

Applications & libraries — layers 5, 6, 7

Operating system — layers 2, 3, 4

Internet protocol stack

Data Link Layer

Data Link Layer (Layer 2)

Layer 2 (Ethernet/Wi-Fi switches) generally has weak security

- CAM overflow

- VLAN Hopping

- ARP cache poisoning

- DHCP spoofing

Link Layer: CAM overflow
Monitor all traffic on a LAN

Layer 2: Ethernet Switches

TP-Link Switch
- 8 1-GbE ports

Cisco Nexus 9516 Switch
- 1/10/40 GbE
- 21-rack-unit chassis
- Up to 576 1/10 Gb ports

Ethernet MAC addresses

Ethernet frames are delivered based on their 48-bit MAC* address

- Top 24 bits: manufacturer code assigned by IEEE
- Bottom 24 bits: assigned by manufacturer
- `ff:ff:ff:ff:ff:ff` = broadcast address

Ethernet MAC address ≠ IP address

*MAC = Media Access Control address – used as a link-layer address by Ethernet, Wi-Fi, and Bluetooth
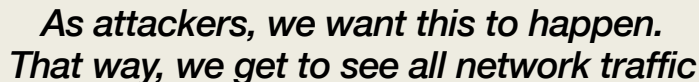
How does an Ethernet switch work?

A switch contains a switch table (MAC address table)

– Contains entries for known MAC addresses & their interface

Forwarding & filtering:

a frame arrives for some destination address D

1. Look up D in the switch table to find the interface

2. If found & the interface is the same as the one the frame arrived on
   Discard the frame (**filter**)

3. If found & D is on a different interface
   Forward the frame to that interface: queue if necessary

4. If not found
   - **Forward** to <u>ALL</u> interfaces

> *As attackers, we want this to happen.*
> *That way, we get to see all network traffic*

The switch table

A switch is self-learning

- **Switch table (MAC address → interface): initially empty**

- **Whenever a frame is received, associate the interface with the source MAC address in the frame**

- **Delete switch table entries if they have not been used for some time**

Switches must be fast: can't waste time doing lookups
- They use CAM – **Content Addressable Memory**
- Fixed size table

CAM overflow attack

Exploit size limit of CAM-based switch table

- **Send bogus Ethernet frames with random source MAC addresses**
  - Each new address will displace an entry in the switch table

- **With the CAM table full, legitimate traffic will be broadcast to all links**
  - A host on any port can now see all traffic
  - CAM overflow attack turns a switch into a hub

Countermeasures:

Port security
- Some managed switches let you limit # of addresses per switch port

802.1x support
- All traffic from a port is initially "unauthorized" and redirected to an authentication server

dsniff: collection of tools for network auditing and penetration testing
`https://monkey.org/~dugsong/dsniff/`

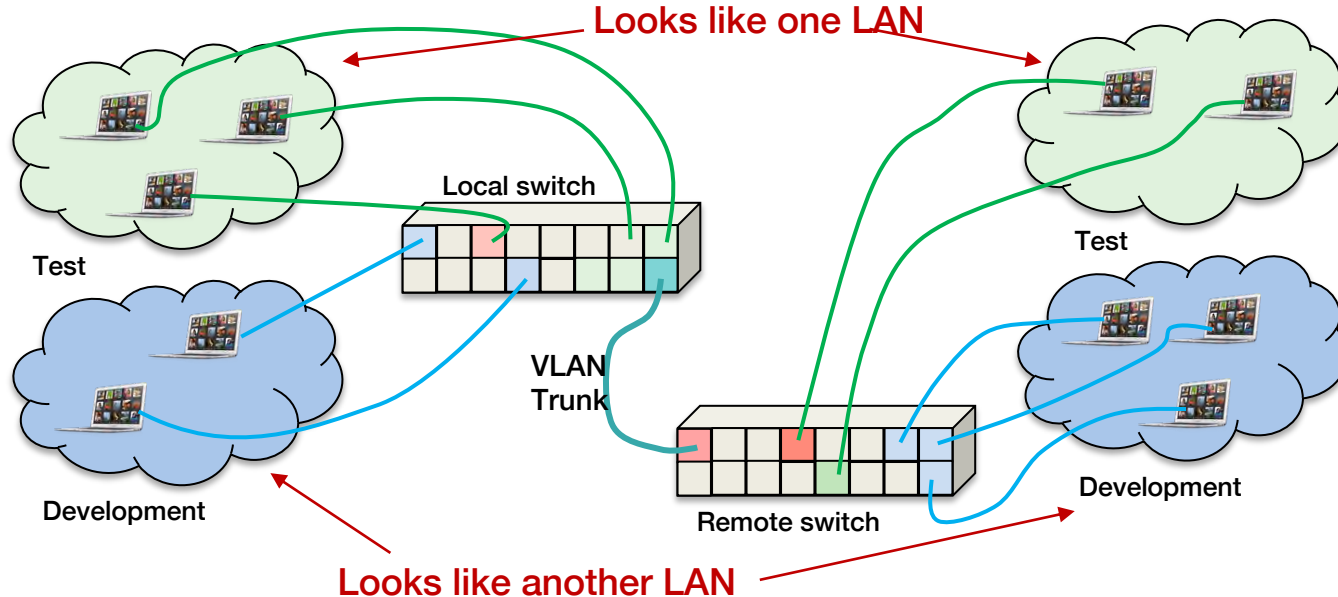# Link Layer: VLANs & VLAN hopping
Join VLANs you are not a member of

VLANs

- A switch & cables creates a local area network (LAN)

- We use LANs to
  - Isolate broadcast traffic from other groups of systems
  - Isolate users into groups
  - What if users move? What if switches are inefficiently used?

- Virtual Local Area Networks (VLANs)
  - Create multiple virtual LANs over one physical switch infrastructure
  - Network manager can assign a switch's ports to a specific VLAN
  - Each VLAN is a separate broadcast domain

VLAN Trunking

VLANs across multiple locations/switches

– VLAN Trunking: a single connection between two VLAN-enabled switches carries all traffic for all VLANs

Looks like one LAN

Local switch

Test

Development

VLAN Trunk

Test

Remote switch

Development

Looks like another LAN

VLAN Hopping Attack

- **VLAN trunk carries traffic for _all_ VLANs**

- **Extended Ethernet frame format**
  - 802.1Q for frames on an Ethernet trunk = Ethernet frame + VLAN tag
  - Sending switch adds VLAN tag for traffic on the trunk
  - Receiving switch removes VLAN tag and sends traffic to appropriate VLAN ports based on VLAN ID

Attack: **switch spoofing**

Devices can spoof themselves to look like a switch with a trunk connection and become a member of all VLANs

Local switch

VLAN Trunk

Remote switch
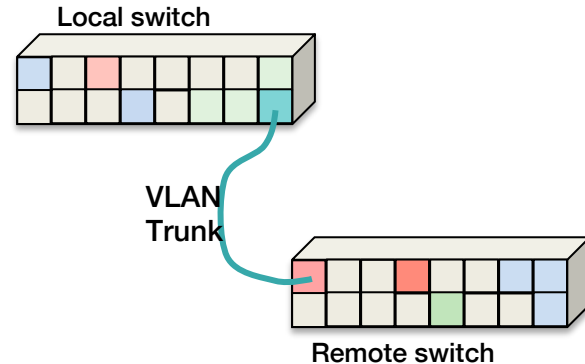
Avoiding VLAN Hopping

Disable	**Disable unused ports & assign them to an unused VLAN** • **Stops an attacker from plugging a device into an unused port**
Disable	**Disable auto-trunking** • **Stops an attacker from masquerading as a switch**
Configure	**Explicitly configure trunking on switch ports that are used for trunks** • **Allows legitimate connected switches to work**

ARP Cache Poisoning (ARP Spoofing)

Intercept traffic for other IP addresses

Find MAC address given an IP address

- We need to send a datagram to an IP address

- It is encapsulated in an Ethernet frame and a MAC address

MAC destination	MAC source	type	IP header	IP data	CRC

How do we know what MAC address to use?

ls.cs.rutgers.edu:
   IP address: `128.6.13.171`
   MAC address: `40:b0:34:f6:cd:0f`

ilab1.cs.rutgers.edu:
   IP address: `128.6.4.101`
   MAC address: `ee:4f:34:13:19:78`

Address Resolution Protocol (ARP)

ARP Table

– Kernel table mapping IP addresses & corresponding MAC addresses

– OS uses this to fill in the MAC header given an IP destination address

– *What if the IP address we want is not in the cache?*

ARP Messages

– A host creates an ARP query packet & broadcasts it on the LAN

  • Ethernet broadcast MAC address: `ff:ff:ff:ff:ff:ff`

– All adapters receive it

  • If an adapter's IP address matches the address in the query, it responds

  • Response is sent to the MAC address of the sender

HW Protocol (ethernet)	Protocol type (e.g., IPv4)	MAC addr length	query/ response	sender MAC addr	sender IP addr	target MAC addr	target IP addr

ARP packet structure

see the ***arp*** command on Linux/BSD/Windows/macOS

Address Resolution Protocol (ARP)

ARP Table

OS uses this to fill in the MAC header given an IP destination address

```
$ arp -a
? (169.254.169.254) at <incomplete> on enp1s0f0
vlan4-lcsr-gw.rutgers.edu (128.6.4.1) at 50:9a:4c:e2:d4:91 [ether] on enp1s0f0
weblogin.cs.rutgers.edu (128.6.4.52) at 52:54:00:4d:c8:a2 [ether] on enp1s0f0
rlab2.cs.rutgers.edu (128.6.4.149) at 00:25:90:e5:c3:c9 [ether] on enp1s0f0
services.cs.rutgers.edu (128.6.4.94) at 00:0c:29:34:76:8a [ether] on enp1s0f0
ilab1.cs.rutgers.edu (128.6.4.101) at a0:42:3f:38:bd:32 [ether] on enp1s0f0
krb2.cs.rutgers.edu (128.6.4.10) at 00:0c:29:e7:74:11 [ether] on enp1s0f0
rlab4.cs.rutgers.edu (128.6.4.151) at 00:25:90:ed:97:37 [ether] on enp1s0f0
ilab3.cs.rutgers.edu (128.6.4.103) at a0:42:3f:3d:07:4e [ether] on enp1s0f0
ilabu3.cs.rutgers.edu (128.6.4.133) at b4:96:91:61:de:4d [ether] on enp1s0f0
rlab1.cs.rutgers.edu (128.6.4.148) at 00:25:90:e5:c3:bd [ether] on enp1s0f0
krb1.cs.rutgers.edu (128.6.4.9) at 00:0c:29:db:d0:83 [ether] on enp1s0f0
rlab3.cs.rutgers.edu (128.6.4.150) at 00:25:90:ed:97:35 [ether] on enp1s0f0
ilab4.cs.rutgers.edu (128.6.4.28) at b4:96:91:78:c8:ac [ether] on enp1s0f0
```

What if the IP address we want is not in the cache?

ARP Cache Poisoning

- **Network hosts cache any ARP replies they see … even if they did not originate them … on the chance that they might have to use that IP address**

- **Any client is allowed to send an *unsolicited* ARP reply**
  - This is called a gratuitous ARP (a response that doesn't have a matching request)
  - Systems often do this when booting to announce themselves on the network

- **ARP replies will overwrite older entries in the ARP table**

- **An attacker can create fake ARP replies**
  - Containing the attacker's MAC address and the target's IP address
  - This will direct any traffic meant for the target to the attacker
  - Enables man-in-the-middle or denial of service attacks

See *Ettercap* – a multipurpose sniffer/interceptor/logger
https://github.com/Ettercap/ettercap

Defenses against ARP cache poisoning

- **Ignore replies that are not associated with requests**
  - But you have to hope that the reply you get is a legitimate one

- **Use static ARP entries**
  - But can be an administrative nightmare

- **Enable Dynamic ARP Inspection**
  - Validates ARP packets against DHCP Snooping database information or static ARP entries

DHCP Server Spoofing
Configure hosts with your chosen network settings

DHCP (Dynamic Host Configuration Protocol)

Computer joins a network – needs to be configured

– *Broadcasts* a *DHCP Discover* message

A DHCP server picks up this request and sends back a response

– IP address

– Subnet mask

– Default router (gateway)

– DNS servers

– Lease time

Attack:

Spoof responses that would be sent by a valid DHCP server

DHCP Spoofing

- **Anybody can pretend to be a DHCP server**
  - Spoof responses that would be sent by a valid DHCP server
  - Provide:
    - False gateway address
    - False DNS server address

- **Attacker can now direct traffic from the client to go anywhere**

- **The real server may reply too**
  - If the attacker responds first, he wins
  - Attack the server first – delay or disable the real server: denial of service attack

Defenses

## Some switches (Cisco, Juniper) support DHCP snooping

– Switch ports can be configured as "trusted" or "untrusted"

– Only specific machines are allowed to send DHCP responses

– The switch will use DHCP data to track client behavior

- Ensure hosts use only the IP address assigned to them
- Ensure hosts do not fake ARP responses

The Network (IP)
and Transport (TCP, UDP) Layers

CS 419 © 2025 Paul Krzyzanowski

The Internet

Packet switching: store-and-forward routing across multiple physical networks
... across multiple organizations

ISP

Backbone

ISP

ISP
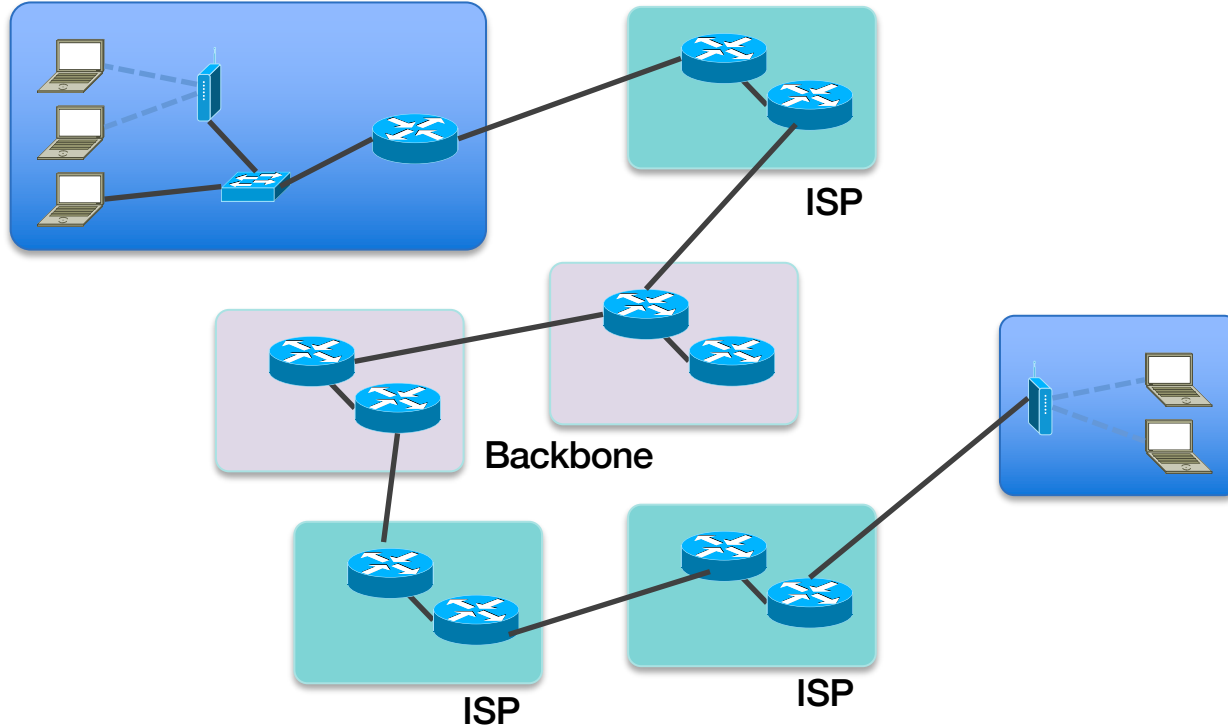
The Internet: Key Design Principles

1. **Use packet switching**
   – Break data into packets

2. **Support the interconnection of networks with routers**
   – Connect diverse networks into one cohesive system: IP is a *logical network*
   – Each router provides store & forward delivery of packets

3. **No centralized control and no central hub**
   – Each node makes its own decisions on the best next hop; no central hub for traffic

4. **Intelligence is handled at the edges: end-to-end principle**
   – Assume unreliable communication
   – Endpoints are responsible for implementing confidentiality, authentication, integrity, prioritization, reliability, sequencing, compression

5. **Protocols are organized into layers**
   – Each layer is responsible for a different aspect of communication

The Internet Introduces Risks

"The internet was designed to be open, transparent, and interoperable. Security and identity management were secondary objectives in system design. This lower emphasis on security in the internet's initial design not only gives attackers a built-in advantage. It can also make intrusions difficult to attribute, especially in real time. This structural property of the current architecture of cyberspace means that we cannot rely on the threat of retaliation alone to deter potential attackers. Some adversaries might gamble that they could attack us and escape detection."

– William J. Lynn III, Deputy Defense Secretary, 2010

http://archive.defense.gov/speeches/speech.aspx?speechid=1593

The Internet Makes It Easier To Attack

- **Security was not a design consideration**
  - This is not a bug but a design decision

- **Intelligence is at the edges of the network – distributed among many players**
  - Reliability, authentication, authorization, encryption, congestion notification, and quality of service are the responsibility of endpoints

- **Access and routing are not centrally managed**
  - Routing decisions distributed
  - DNS (domain name system) service is distributed too
  - No access control: any system can be added to the Internet

- **Bad actors can hide!**

How the Internet Creates Vulnerabilities

- **Action at a distance**

- **Asymmetric force**

- **Actors can be anonymous**

- **No borders or checkpoints**

- **No distinction**
  - Hard to distinguish valid data from attacks
  - Can't tell what code will be harmful until it's executed

Network Layer (IP) vulnerabilities

Network Layer: IP

Responsible for end-to-end delivery of packets

- **No guarantees on message ordering or delivery**

- **Key functions**
  - Routing
    - Each host knows the address of one or more connected routers (gateways)
    - The router knows how to route to other networks
  - Fragmentation & reassembly
    - An IP fragment may be split if the MTU* size on a network is too small
    - Reassembled at its final destination
  - Error reporting
    - ICMP messages sent back to the sender (e.g., if packet is dropped)
  - Time-to-live (TTL)
    - Hop count avoids infinite loops; packet dropped when TTL = 0

*MTU = Maximum Transmission Unit = maximum packet size on a network link

*ICMP = Internet Control Message Protocol = An IP protocol used by devices to send error and diagnostic info

Source IP address

No source IP address authentication

- **Clients are *supposed* to use their own source IP address**
  - Can override with raw sockets
  - Responses will be sent to the forged source IP address

- **Enables denial of service (DoS) attacks**
  - Forged source address provides anonymity
  - Forged source address keeps the attacker from getting responses
  - Changing the address to a victim's address enables *reflection attacks*
    - Send lots of messages to a service
    - All responses will go to the forged source address (the victim being attacked)

Attacks on routers

- **Routers are just special-purpose computers**
  - People may keep default passwords or not use strong passwords
  - Router OS & supporting services may be buggy or not be kept up to date

- **Subject to attacks:**
  - Denial of Service (DOS)
    - Flood the router (e.g., lots of ICMP packets from lots of sources)
  - **Routing table poisoning**
    - Either by breaking into a router or by sending modified routing data update packets

Routers are vulnerable, like computers

The Hacker News

Mysterious Cyber Attack Took Down 600,000+ Routers in the U.S.

May 31, 2024   Ravie Lakshmanan

Logical Overview:Chalubo Infection
Ocotber 2023

More than 600,000 small o... taken offline following a d... access to the internet.

techradar pro

Hackers of all kinds are attacking routers across the world

News   By Sead Fadilpašić published May 2, 2024

What happens when multiple groups target the same router?

The Hacker News

Quad7 Botnet Expands to Target SOHO Routers and VPN Appliances

Sep 11, 2024   Ravie Lakshmanan

The operators of the mysterious Quad7 botnet are actively evolving by compromising several brands of SOHO routers and VPN appliances by leveraging a combination of both known and unknown security flaws.

Targets include devices from TP-LINK, Zyxel, Asus, Axentra, D-Link, and NETGEAR, according to a new report by French cybersecurity company Sekoia.

KRATIKAL

Kratikal Blogs
Information Hub For Cyber Security Experts

CYBER SECURITY / NETWORK SECURITY

60,000+ Vulnerabilities Exposed by the U.S. Router Cyber Attack

June 11, 2024 · 7 mins read

ars TECHNICA

Thousands of hacked TP-Link routers used in yearslong account takeover attacks

The botnet is being skillfully used to launch "highly evasive" password-spraying attacks.

gbhackers.

Home · Cyber Attack · Hackers Exploit Router Flaws in Ongoing Attacks on Enterprise Networks

Hackers Exploit Router Flaws in Ongoing Attacks on Enterprise Networks

Cyber Attack   Cyber Security   Cyber Security News   THREATS

PUBLISHED ON APRIL 11, 2025   BY AMAN MISHRA

	Router Brand	Default IP Address	Default Username	Default Password
1	3Com	http://192.168.1.1	admin	Admin
2	Belkin	http://192.168.2.1	admin	admin
3	BenQ	http://192.168.1.1	admin	Admin
4	D-Link	http://192.168.0.1	admin	Admin
5	Digicom	http://192.168.1.254	admin	Michelangelo
6	Linksys	http://192.168.1.1	admin	Admin
7	Netgear	http://192.168.0.1	admin	password
8	Sitecom	http://192.168.0.1	sitecom	Admin
9	Asus	http://192.168.1.1	admin	admin
10	Synology	http://192.168.1.1	admin	Admin
11	Arris	http://192.168.0.1	admin	password
12	Apple iphoneIOS4.X	http://10.0.1.1	root	alpine
13	DELL	http://192.168.1.1	admin	password
14	Huawei ADSL2+	http://192.168.0.1	admin	admin
15	Netcomm	http://192.168.1.1	admin	password

https://www.softwaretestinghelp.com/default-router-username-and-password-list/

Industrial Routers Compromised for DDoS Attacks

- **December 27, 2024:**
  - Attackers use default credentials on Four-Faith industrial routers to perform command injection attacks.

- **Command injections were carried out via an HTTP POST message that sets `adjust_sys_time` and passes it a `$(shell_command)` in the string.**

- **These compromised routers were injected with a variant of the Mirai botnet, which contains about 15,000 active IP addresses.**

https://vulncheck.com/blog/four-faith-cve-2024-12856?ref=blog.xlab.qianxin.com

Just a few recent headlines (`routersecurity.org`)

Mar 16, 2025	Mirai Bot now incorporating DrayTek Vigor Router Exploits
Feb 25, 2025	Advisory: Multiple vulnerabilities affecting Draytek routers
Feb 4, 2025	Netgear warns users to patch critical Wi-Fi router vulnerabilities
Jan 28, 2025	Active Exploitation of Zero-day Zyxel CPE Vulnerability
Jan 23, 2025	Palo Alto Networks firewalls have UEFI flaws, Secure Boot bypasses
Jan 14, 2025	Fortinet warns of auth bypass zero-day exploited to hijack firewalls
Dec 27, 2025	Hackers exploit DoS flaw to disable Palo Alto Networks firewalls
Dec 3, 2024	ASUS Router Improper Input Validation

Transport Layer (UDP, TCP) vulnerabilities
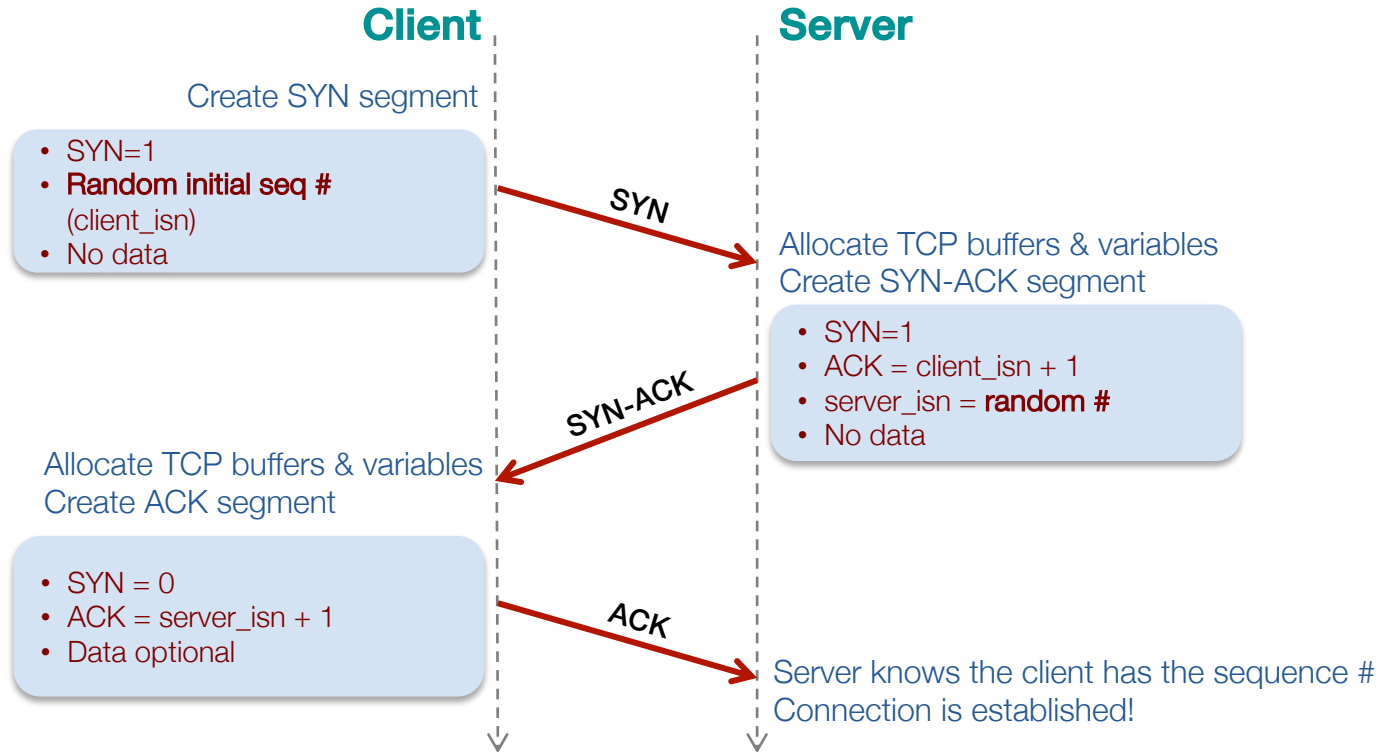
TCP & UDP

UDP: User Datagram Protocol

– Stateless, connectionless & unreliable

– Anyone can send forged UDP messages

TCP: Transmission Control Protocol

– Stateful, connection-oriented & reliable

– Every packet contains a sequence number (byte offset)

- Receiver assembles packets into correct order
- Sends acknowledgements
- Missing packets are retransmitted

TCP connection setup: three-way handshake

Client   **Server**

Create SYN segment

- SYN=1
- **Random initial seq #** (client_isn)
- No data

SYN →

Allocate TCP buffers & variables
Create SYN-ACK segment

- SYN=1
- ACK = client_isn + 1
- server_isn = **random #**
- No data

← *SYN-ACK*

Allocate TCP buffers & variables
Create ACK segment

- SYN = 0
- ACK = server_isn + 1
- Data optional

ACK →

Server knows the client has the sequence #
Connection is established!

Why random initial sequence numbers?

If predictable, an attacker can create a TCP session on behalf of a forged source IP address by guessing the ACK #

SYN
src_addr = **victim**

ACK
src_addr = victim
ack # = X+1

command
src_addr = victim

Attacker

Server

SYN/ACK
dest_addr = **victim**
server_isn = X

Huh?
Ignore.

Victim
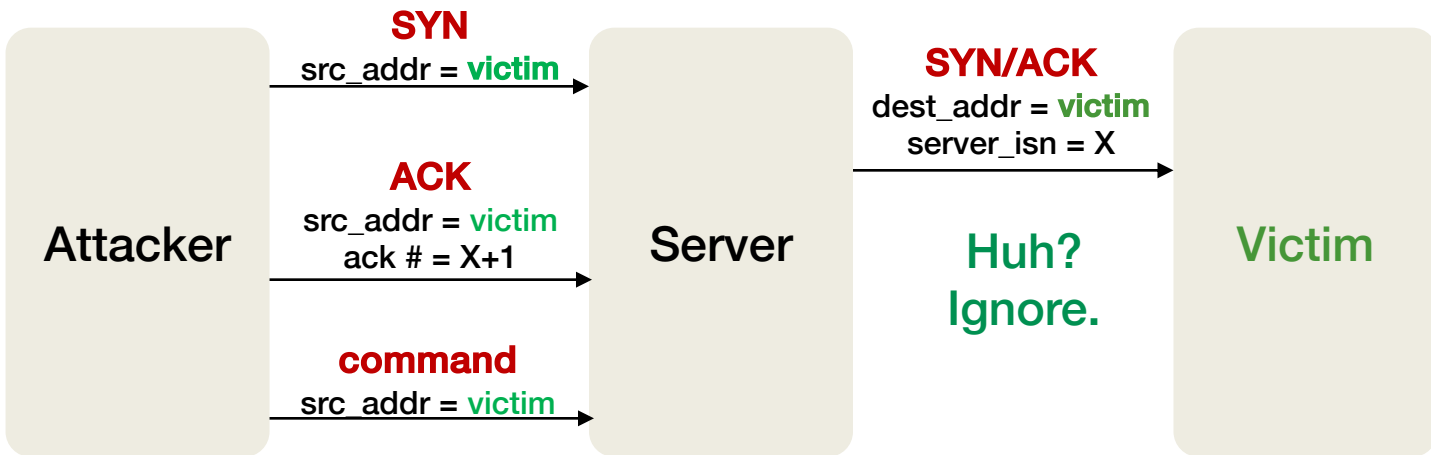
Random numbers make this attack harder – especially if the attacker cannot sniff the network

Denial of service: SYN Flooding

An OS will allocate only a finite # of TCP buffers

- **SYN Flooding attack**
  - Send lots of SYN segments but never complete the handshake
  - The OS will not be able to accept connections until those time out

- **SYN Cookies: Dealing with SYN flooding attacks**
  - Do not allocate buffers & state when a SYN segment is received
  - Create initial sequence # =
    hash(src_addr, dest_addr, src_port, dest_port, SECRET)
  - When an ACK comes back, validate the ACK #
    Compute the hash as before & add 1
  - If valid, then allocate resources necessary for the connection & socket

Denial of service: Reset

- **Attacker can send a RESET (RST) packet to an open socket**

- **If the server sequence number is correct, then the connection will close**

- **Sequence numbers are 32 bits**
  - Chance of success is $1/2^{32} \approx 1$ in 4 billion
  - But many systems allow for a large range of sequence numbers
  - Attacker can send a flood of RST packets until the connection is broken

Network Routing Protocols

IP Routing Protocols

Network operators (autonomous systems) need to know how to route packets within their network and the best connection to use for packets that are routed outside their network

OSPF: Open Shortest Path First

– Interior Gateway Protocol (IGP) within an autonomous system (AS)

– Uses a link state routing algorithm (Dijkstra's shortest path)

BGP: Border Gateway Protocol

– Exterior Gateway Protocol (EGP) between autonomous systems (AS)

– Network operators exchange routing and reachability information

 • Each sends a list of blocks of addresses they can route to and the distance to each block

 • Identifies the owner and AS route to reach the owner

– Distance vector routing protocol
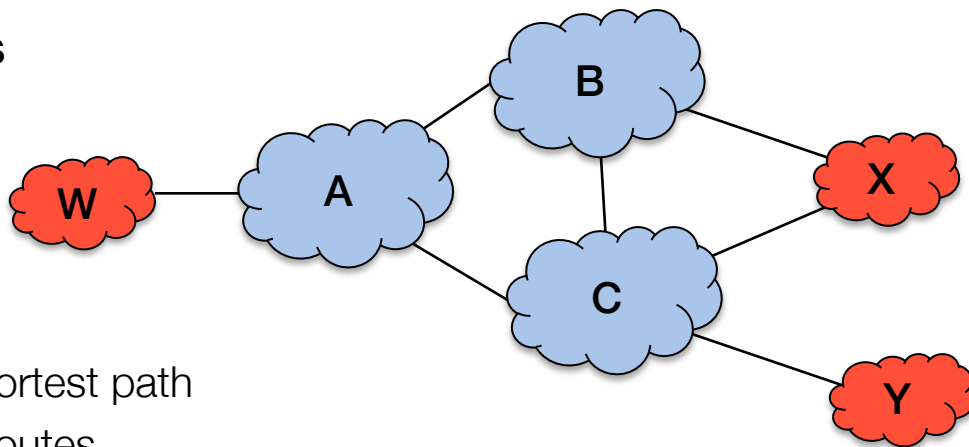
BGP sessions maintained via TCP links

Pairs of routers exchange information via semi-permanent TCP connections

- One connection for each link between gateway routers
  - External BGP (eBGP) session
- Also, BGP TCP connections between routers *inside* an AS
  - Internal BGP (iBGP) session

- **A, B, C: transit ASes – ISPs & backbone**

- **W, X, Y: stub ASes – customers**

BGP route selection

- In general, pick the route with the shortest path
- Policies allow selection of preferred routes
- More specific route definitions get priority:
  - An advertised route for 128.6.48.0/24 gets chosen over 128.6.0.0/16 if the address matches both

This is a **prefix**: it tells the router to match the hig—order (most significant) 24 bits of the IP address. Prefixes represent a range of addresses.

An Autonomous System (AS) is a network or group of networks under one administrative control that uses BGP for routing.

BGP Prefix Hijacking

- **BGP was built based on trust**
  - Each network operator trusts others & believes the information it receives is accurate
  - The trust is a chain: a network operator sends route advertisements that are built from data it received from other network operators

- **Route advertisements are not authenticated**
  - A malicious network operator can inject advertisements for arbitrary routes
  - Information will propagate throughout the Internet
  - Can be used for DoS (dropping packets), eavesdropping, man-in-the-middle attacks, or redirecting traffic to malicious computers

> **A prefix is the # of bits in an IP address to use for routing**
> **More bits = a more specific the network (fewer machine addresses) = higher priority over an address with fewer bits**

Pakistan's attack on YouTube in 2008

- **YouTube service was cut off the global web for over an hour**

- **Pakistan Telecom received a censorship order from the telecommunications ministry to block YouTube**
  - The company sent spoofed BGP messages claiming to be the best route for YouTube's range of IP addresses

Pakistan's attack on YouTube in 2008

- **Pakistan Telecom sent BGP advertisements that it was the correct route for 256 addresses in YouTube's 208.65.153.0 network**
  - Advertise a /24 network

- **That is a more specific destination than YouTube's broadcast, which covered 1024 addresses**
  - YouTube advertised a /22 network
  - Within minutes, all YouTube traffic started to flow to Pakistan

- **YouTube immediately tried countermeasures**
  - Narrowed its broadcast to 256 addresses … but too late
  - Then tried an even more specific group: 64 addresses
    Advertise a /26 network ⇒ priority over /24 routes
    - Routes for more specific addresses overrule more general ones
  - Route updates were finally fixed after 2 hours

2013 – Repeated attacks

- **38 events observed where traffic to 1,500 blocks of IP addresses was redirected to Iceland or Belarus**
  - Redirection ranged from a few minutes to several days
  - Over 60 days of man-in-the-middle attacks observed

- **Data targeted to 150 cities was intercepted**

Traceroute Path 1: from Guadalajara, Mexico to Washington, D.C. via *Belarus*

7. Moscow, Russia
8. Minsk, Belarus
6. London, UK
9. Frankfurt, Germany
4. Ashburn, VA
10. New York, NY
5. Washington, D.C.
11. Washington, D.C. **END**
3. Laredo, TX
2. Monterrey, Mexico
START 1. Guadalajara, Mexico

renesys

Source: *Renesys Path Measurements*

Traceroute Path 2: from Denver, CO to Denver, CO via *Iceland*

6. Reykjavik, Iceland
5. London, UK
7. Montreal, Canada
2. Chicago, IL
8. Chicago, IL
START/END
1. Denver, CO
12. Denver, CO
4. New York, NY
9. New York, NY
11. Kansas City, MO
3. Ashburn, VA
10. Dallas, TX

renesys

Source: *Renesys Path Measurements*

URL: https://arstechnica.com/information-technology/2013/11/repeated-attacks-hijack-huge-chunks-of-internet-traffic-researchers-warn/

- **Russian domestic traffic was repeatedly rerouted to routers operated by China Telecom**

- **Occurred after Russian mobile provider Vimpelcom and China Telecom signed a peering agreement to carry traffic over each other's network at no cost**

- **The rerouting could have been a configuration error**

  – But could also have been espionage or hacking

Traceroutes from **Moscow, RU** to **Yaroslavl, RU**

https://arstechnica.com/information-technology/2014/11/wtf-russias-domestic-internet-traffic-mysteriously-passes-through-china/

2017 – Selected traffic routed to Russian ISP

- **Traffic belonging to Google, Facebook, Apple, Microsoft, Twitch, and Riot Games was routed through a Russian Internet provider**
  - Eight months earlier, traffic for MasterCard, Visa, and more than two dozen other financial services was routed through a Russian government-controlled telecom

- **Considered suspicious & not a configuration error**
  - Targeted very specific companies
  - Advertised IP address blocks were broken into small chunks
    - BGP prioritizes more specific blocks of addresses; this ensures they get selected over broader advertisements for the same groups of addresses

https://arstechnica.com/information-technology/2017/12/suspicious-event-routes-traffic-for-big-name-sites-through-russia/

https://arstechnica.com/information-technology/2017/04/russian-controlled-telecom-hijacks-financial-services-internet-traffic/

- China Telecom redirected large chunks of Internet traffic through their routers

- This took place for 2.5 years

- China Telecom was incorrectly advertising routes to Verizon's Asia-Pacific AS (AS703)
  – Packets would be routed to China

- Could have been a configuration error

China Telecom's Internet Traffic Misdirection
Routing leak sent US domestic traffic through China

Washington, DC
Packets arrive from Asia to their destination in the US

Los Angeles, CA
Packets originate from LA (depicted as ⊙)

Eastern Asia
Packets travel from LA to Shanghai, China and on to Hong Kong before returning to the United States
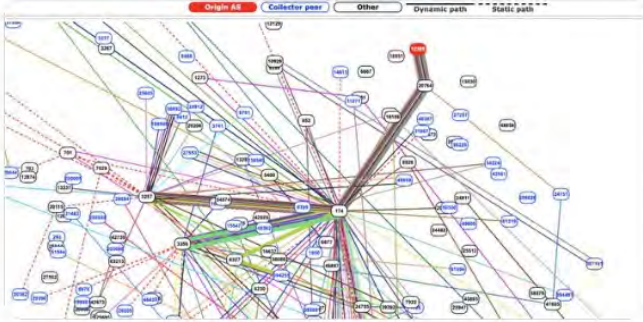
INTERNET INTELLIGENCE | ORACLE Cloud Infrastructure

https://arstechnica.com/information-technology/2018/11/strange-snafu-misroutes-domestic-us-internet-traffic-through-china-telecom/

2020 – Traffic from >200 providers redirected to Russia

- **Traffic for content delivery networks and cloud providers was redirected through Rostelecom, Russia's state-owned telecom provider**

- **Affected over 8,800 routes from over 200 networks**

- **Lasted for an hour**

- **Companies affected included Google, Amazon, Facebook, Akamai, and Cloudflare**

- **Could have been a mistake**

Cisco BGPmon
@bgpmon

Earlier this week there was a large scale BGP hijack incident involving AS12389 (Rostelecom) affecting over 8,000 prefixes.
Many examples were just posted on @bgpstream , see for example this example for @Facebook
bgpstream.com/event/230837

12:51 PM · Apr 5, 2020 · Twitter Web App

226 Retweets   56 Quote Tweets   318 Likes

https://www.zdnet.com/article/russian-telco-hijacks-internet-traffic-for-google-aws-cloudflare-and-others/

2022 – Klayswap crypto attack

- **Hackers stole almost $2M from South Korean cryptocurrency platform KLAYswap**

- **Used a rogue autonomous system**
  - Advertised IP address for `developers.kakao.com`
  - `developers.kakao.com` hosts the Kakao SDK used by third-party developers

- **Attackers hijacked the address and served a malicious version of KakaoTalk's JavaScript SDK file**
  `https://developers.kakao.com/sdk/js/kakao.min.js`

- **Users thought they were downloading it from the official site, but it came from the attacker's servers**

- **Code waited for a transaction and transferred funds to an attacker's wallet**

- **Attack lasted two hours and incurred 407 transactions across 325 customer wallets**

https://therecord.media/klayswap-crypto-users-lose-funds-after-bgp-hijack/

BGP Attacks Continue

13,626 BGP hijacks detected in Q2 2024; 13,438 in Q3 2024

BGP ROUTE LEAKING ASes	2024, Q3	BGP HIJACKING ASes
2 036	JULY	9 711
2 040	AUGUST	9 465
1 883	SEPTEMBER	4 570
Unique BGP Route Leakers: 3 123		Unique BGP Hijackers: 13 438

https://qrator.net/blog/details/q3-2024-ddos-bots-and-bgp-incidents-statistics-and

Defending against BGP Hijacking – RPKI

RPKI (Resource Public Key Infrastructure) framework

RPKI ensures the origin is legitimate

– Provides a way to validate that the AS that is announcing the route is authorized to do so for the addresses (prefix) it is advertising

– Allows network operators to identify & ignore fake announcements & prevents IP prefix hijacking

How it works:

1. IP address holder gets a **digital certificate** from a Regional Internet Registry (RIR). This proves ownership of a prefix.
2. IP address holder creates a **Route Origin Authorization** (ROs)
   – The ROA identifies which network operator is allowed to announce an organization's IP addresses using BGP
   – This ROA is digitally signed with the private key associated with the certificate.
3. Other ISPs download ROAs from the RIRs & cache them.
4. An ISP uses an ROA to validate the route's prefix and origin AS against cached ROA data
   – Route advertisements without a valid, signed ROA are ignored

See: https://labs.ripe.net/author/niklas-vogel/crashing-the-party-vulnerabilities-in-rpki-relying-party-software/

RPKI (Resource Public Key Infrastructure) framework

Does not require any changes to the BGP protocol

– BGP messages are unchanged.

– ISPs contact the RIR (responsible for allocating and managing IP addresses) to download the ROA records to check if the route advertisement came from an approved place.

Challenges

– Only about 57% of advertised routes use ROA as of April 2025

– Doesn't stop all hijacks: a malicious AS can intercept traffic by sending BGP *UPDATE* messages with a legitimate source but claiming to have a more efficient path (lower hop count)

– Risk of misconfiguration

Address owner gets certificate → **Creates ROA** → **ISP downloads ROAs** → **Checks prefix & origin against ROA**

See: https://labs.ripe.net/author/niklas-vogel/crashing-the-party-vulnerabilities-in-rpki-relying-party-software/

Defending against BGP Hijacking – BGPsec

BGPsec

– Security-enhanced version of the BGP protocol

– Protects integrity of BGP update messages – each AS adds its signature to the advertised route

– Downgrade attacks possible if all ASes don't support BGPsec

– Requires all ~71,000 ASes to be able to send *UPDATE* messages to each other

Prefix, P: 128.6.0.0/16

P: {AS46}, {sig46}   **P: {AS46, AS2}, {sig46, sig2}**

P: {AS46, AS2, AS3}, {sig46, sig2, sig3}

Ensures the route's path is valid as it propagates through the network

See: https://labs.ripe.net/author/niklas-vogel/crashing-the-party-vulnerabilities-in-rpki-relying-party-software/

Defending against BGP Hijacking – BGPsec

BGPsec

Challenges

– High overhead: every signature along the path must be validated and routers can get high volumes of BGP advertisements

– Every router in the path must support BGPsec – otherwise it's useless

– It's complex to deploy
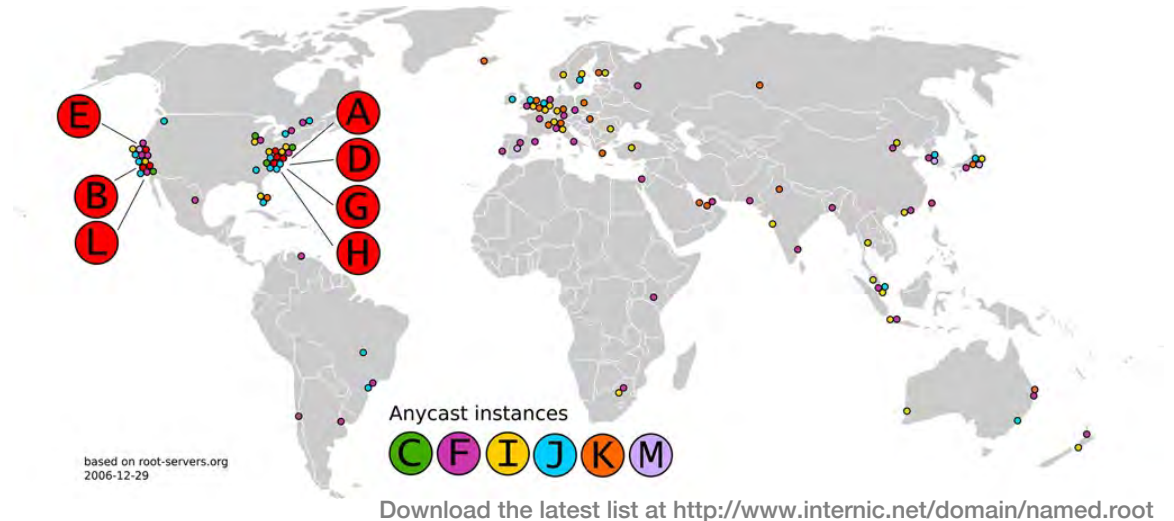
• **Adoption of BGPsec is extremely limited**

See: https://labs.ripe.net/author/niklas-vogel/crashing-the-party-vulnerabilities-in-rpki-relying-party-software/

Domain Name System (DNS) Vulnerabilities
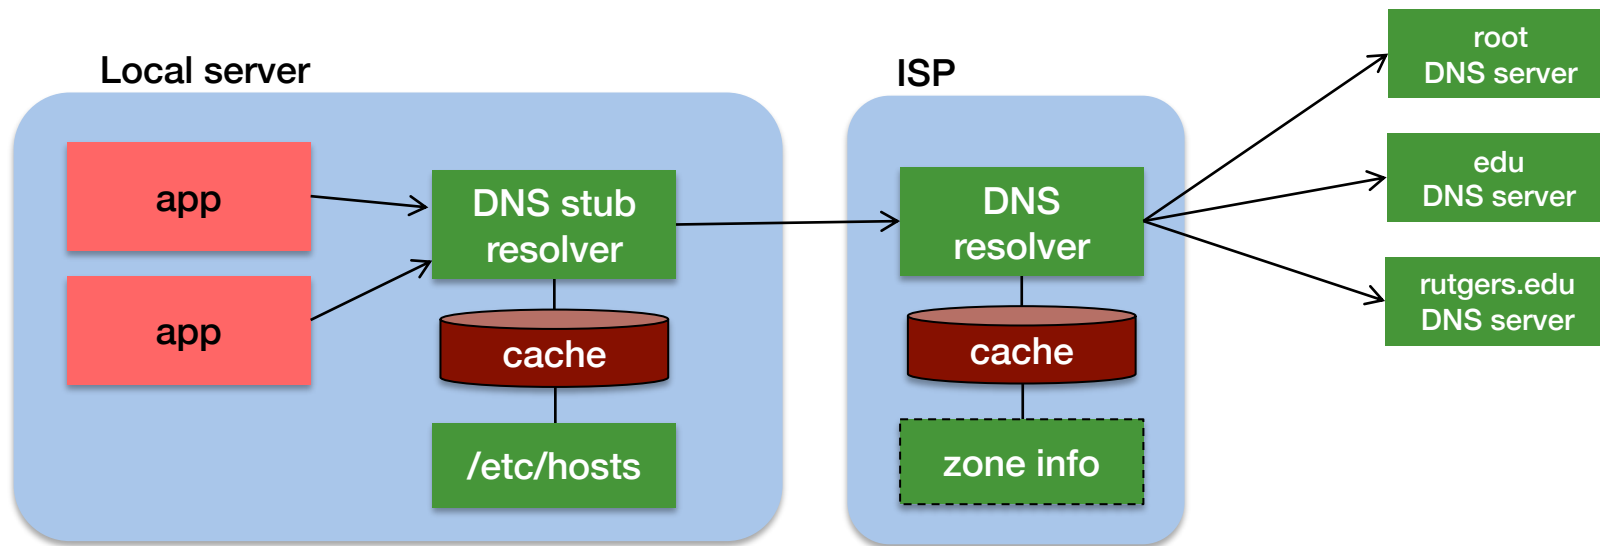
Domain Name System (DNS)

- **Hierarchical service to map domain names to IP addresses**

- **How do you find the DNS Server for rutgers.edu?**
  - That's what the domain registry keeps track of
  - When you register a domain
    - You supply the addresses of at least two **DNS servers** that can answer queries for your zone
    - You give this info to the **domain registrar** (e.g., Namecheap, GoDaddy) who updates the database at the **domain registry** (e.g., Verisign for .com, .net, .edu, .gov, … domains)
      - **Domain registrar:** Sells domain names to the public
      - **Domain registry:** Maintains the top-level domain database

DNS: Root name servers

- ***So how do you find the right DNS server for a domain?***
  - Start at the root

- **The root name servers provide lists of authoritative name servers for top-level domains**

- **13 root name servers:** `A.ROOT-SERVERS.NET, B.ROOT-SERVERS.NET, …`
  - Each server has redundancy (via *anycast* routing or load balancing) and is a set of machines

Anycast instances
C F I J K M

based on root-servers.org
2006-12-29

Download the latest list at http://www.internic.net/domain/named.root

How do DNS queries work?

Local server

ISP

app

app

DNS stub resolver

cache

/etc/hosts

DNS resolver

cache

zone info

root
DNS server

edu
DNS server

rutgers.edu
DNS server

<u>Local stub resolver</u>:
- check local cache
- check local hosts file
- send request to external resolver

E.g., on Linux: resolver is configured via the `/etc/resolv.conf` file

<u>External resolver</u>:
- Running at ISP, Cloudflare, Google Public DNS, OpenDNS, etc.

DNS Vulnerabilities

Programs (and users) trust the host-address mapping

– This is the basis for some security policies

- Browser same-origin policy, URL address bar

- **But DNS responses can be faked**

  – If an attacker gives a DNS response first, the host will use that

  – Malicious responses can direct messages to different hosts

  – A receiver cannot detect a forged response

- **DNS resolvers cache their results (with an expiration)**

  – If it gets a forged response, the forged results will be passed on to any systems that query it

Pharming attack

Pharming attack: the attacker changes DNS results provide addresses that will redirect domains to a malicious site

Forms of attack

1. Use malware or social engineering to modify the victim's *hosts* file

   This file maps *names→IP addresses* and avoids making external DNS queries

2. Attack the router or DHCP server & modify its DNS server setting

   Direct traffic to the attacker's DNS server, which will give the wrong IP address for certain domain names

3. Attack the DNS server to provide a malicious address for a domain

   Exploit vulnerabilities in the system hosting the DNS service

DNS spoofing attack

Redirect traffic to an attacker via DNS **cache poisoning**

An attacker sends a malicious DNS response to the victim

The DNS resolver requesting it will cache it and provide that to anyone else who asks in the near future

- **How do we prevent spoofed responses?**
  - Each DNS query contains a 16-bit Query ID (**QID**)
    - Response from the DNS server must have a matching QID
  - DNS uses UDP and this was created to make it easy for a system to match responses with requests

- **An attacker will have to guess the QID number –** *but there are only 65,536 possible #s*
  - But Query IDs were typically sequential and not hard to guess (snoop on previous queries)
  - Fix by using random Query IDs

DNS Spoofing:	focuses on manipulating DNS responses to mislead users temporarily.
Pharming:	Focuses on achieving persistent redirection of users to malicious sites.

# DNS spoofing via Cache Poisoning

How does it work?

– Suppose an attacker wants to change a victim's address for the domain `bank.com`

– Malicious JavaScript on a web page makes a request for `a.bank.com`, which causes a DNS lookup because the domain really doesn't exist and thus isn't cached.

– At the same time, the attacker sends a stream of DNS "responses" for `a.bank.com` hoping that one will have a matching query ID (QID)

If the attacker is successful, one of the responses matches the request

– But the attacker wants to change the address of `bank.com`, not `a.bank.com`

– However… the DNS response can also define a new DNS server for the domain `bank.com` !

– This overwrites any saved DNS info for `bank.com` that may be cached

– *The attacker can take over any requests for bank.com!*

DNS spoofing via Cache Poisoning

JavaScript on the attacker's web page launches a DNS attacker
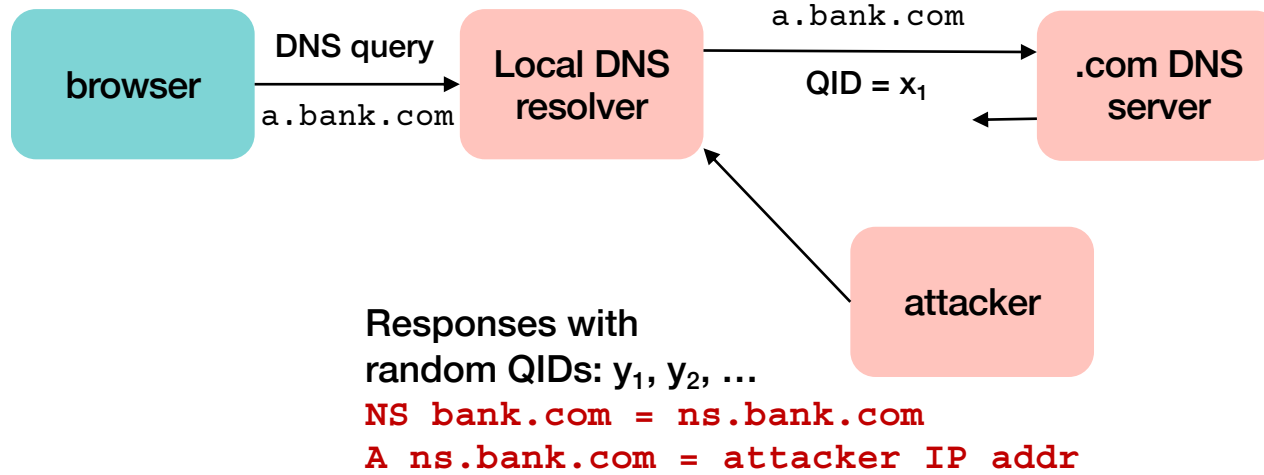
If there is some j such that $x_1 = y_j$ then the response will be cached

All future DNS queries for anything at bank.com will go to attacker_IP_addr

If it doesn't work … try again with `b.bank.com`, `c.bank.com`, etc.

Defenses against DNS cache poisoning

- **Query IDs used to be predictable – use random QIDs**
  - An attacker could have a web page make a DNS query to a domain under the attacker's control & look at the QID sent by the victim
  - The attacker can then guess the next QID

- **Randomize source port # –** *where DNS queries originate*
  - An attack will require several hours instead of a few minutes
  - Will have to send responses to a range of ports
  - But this is tricky in environments that use NAT (network address translation) and may limit the exposed UDP ports

- **Issue double DNS queries**
  - The attacker will have to guess the Query ID twice (32-bit complexity)

- **Use TCP instead of UDP for DNS queries**
  - It's much harder to inject a fake response into a TCP stream
  - But – TCP queries have a much higher latency & much more overhead at the DNS resolver

Defenses against DNS cache poisoning

The better long-term solution: DNSSEC

- **Secure extension to DNS that provide authenticated responses**

- **How it works**
  - Each DNS zone (e.g., bank.com) signs its DNS records with a private key
    - The corresponding public key (in a DNSKEY record) published in DNS, allowing signature verification
  - DNS query responses contain a digital signature (an RRSIG record)
    - A resolver retrieves the associated signature and DNSKEY
    - Verifies the signature
    - Checks the chain of trust to the root

- **But**
  - Adoption has been very slow
  - DNSSEC response size is much bigger than a DNS response, which makes it more powerful for DDoS reflection attacks

Current DNSSEC Deployment – early 2025

SE: 86%

GB: 18%

50%

28%

38%

0.09%

14%

20%

59%

32%

30 day average (14/03/2025 – 12/04/2025

https://stats.labs.apnic.net/dnssec

DNS Rebinding Attack

DNS Rebinding

DNS Rebinding: trick a user's program (usually a browser) into resolving a domain name repeatedly to different IP addresses, allowing the attacker to run scripts that access private network resources.

- **The web's security model relies on comparing domain names**

- **If we can change the underlying address:**
    - We can send messages other systems
        … while the browser thinks it's still going to the same domain
    - This can let us access private machines in the user's local area network
    - Example: access local web services, cameras, thermostats, printers, …

DNS Rebinding: How it Works

- **Attacker**
  - Registers a domain (`attacker.com`)
  - Sets up a DNS server for that domain
  - The DNS server responds with very short TTL values – so the response won't be cached

- **Client (browser)**
  - JavaScript on the attacker's web page causes access to a malicious domain
  - Attacker's DNS server responds with IP address of a server hosting malicious client-side code
  - Malicious client-side code makes additional references to the same domain name
    - This is allowed under the web's same-origin policy
      - Scripts in a page may access data in another page only if both pages have the same origin (protocol, address, port)
    - Because of the short TTL, the script causes the system to issue a new DNS request
    - The attacker's DNS server replies with a new IP address (e.g., a target somewhere in the victim's LAN)
    - The script can continue to access content in the same domain
      - But it really isn't in the same domain!

Defending against DNS rebinding

- **Force minimum time-to-live (TTL) values**
  - This may affect some legitimate dynamic DNS services
  - Many resolvers will only accept a minimum TTL of 30 seconds

- **DNS pinning: refuse to switch the IP address for a domain name**
  - This is similar to forcing minimum TTL values
  - But this can mess up load balanced or other dynamic services

- **Have the local DNS resolver make sure DNS responses don't contain private IP addresses**

- **Server-side defense within the local area network**
  - Reject HTTP requests with unrecognized `Host` headers
  - Authenticate users

Human Factors: Sitting Ducks Attacks

- **Attackers can exploit misconfigured or vulnerable DNS servers and reconfigure them**
  - DNS providers fail to properly verify domain ownership
  - DNS services delegated to another provider that is exploitable
  - Incorrect configurations at the domain registrar

- **35,000 domains have been hijacked between 2018 & 2024**

- **800,000+ domains could be hijacked and 70,00 of those been hijacked since the Infoblox team's investigation**

SECURITY WEEK

CYBERCRIME

Known Brand, Government Domains Hijacked via Sitting Ducks Attacks

Threat actors have hijacked over 70,000 domains, including known brands and government entities, because of failed domain ownership verification.

by Ionut Arghire
November 15, 2024

Tens of thousands of domains, including those of well-known brands, non-profits, and government entities, have been hijacked over the past five years because DNS providers failed to properly verify domain ownership, cybersecurity firm Infoblox reports.

The issue was initially disclosed in late July, when Eclypsium and Infoblox said that roughly 35,000 domains had been hijacked since 2018 by abusing the weakness as part of so-called Sitting Ducks attacks.

However, that was just the tip of the iceberg, Infoblox says in a new report. Further investigation into this configuration-oriented attack vector has revealed that at least 800,000 domains could be hijacked, and that 70,000 have already fallen victim to attackers.

"We know these numbers do not accurately reflect the attack surface: they are derived from a limited monitoring system. The challenge with a Sitting Ducks attack is that it is easy to perform and very hard to detect," Infoblox warns in a new report (PDF).

https://www.securityweek.com/known-brand-government-domains-hijacked-via-sitting-ducks-attacks/

Deliberate DNS Poisoning

March 2025:

- Italy uses its *Privacy Shield* law to demand Google to poison its DNS servers to prevent people from reaching pirate streams of football games

- A similar case was made against Cloudflare's DNS server and WARP VPN.

https://arstechnica.com/gadgets/2025/03/italian-court-orders-google-to-block-iptv-pirate-sites-at-dns-level/

The End