CSE331 Computer Security Fundamentals

10/17/2017

Network Scanning and DoS Attacks

Michalis Polychronakis

Stony Brook University

Information Gathering

First step of an attacker: learn as much about a particular target as possible

- human, system, organization, …

Dependencies and third-party interactions are also important

- Example: the Target 2013 breach was achieved through the compromise of a third-party HVAC vendor who had access to the internal network

Peripheral or "forgotten" systems are often less secure than publicized web, application server, and mail endpoints

Every piece of information counts!

Passive reconnaissance: no direct interaction with the target system

- Information gathering from public sources
- Passive network eavesdropping
- Dumpster diving (e.g., recover data from discarded hard disks)
- Information leakage (e.g., through social engineering)

Active reconnaissance: attacker's activities can be directly detected and logged

- Network scanning
- Service enumeration
- OS and service fingerprinting/probing

OSINT (Open-source Intelligence Gathering)

Intelligence collected from publicly available sources

- As opposed to covert or clandestine sources

Wide variety of information and sources

- *Search engines:* public documents, forgotten web pages, exposed login interfaces, dashboards, historical data, …
- *Public data:* courthouse documents, tax forms, budgets, …
- *Media:* articles, interviews, blog posts, …
- *Social media:* LinkedIn/Facebook/Twitter/etc., mailing lists, …
- *Professional/academic sources:* reports, presentations, …
- *Metadata:* documents, EXIF, executables, email headers, …
- …

WHOIS

Protocol for querying databases of registration information about assignees of internet resources

- IP address blocks, domain names, and autonomous systems
- Top registries: AFRINIC, APNIC, ARIN, IANA, ICANN, LACNIC, NRO, RIPE, InterNic

`whois` command-line utility

```
# whois stonybrook.edu
# whois 130.245.27.2
```

Registrars and third-party services provide web interfaces

Useful information

- Registrar information, domain creation/expiration dates, primary DNS name servers associated with the domain
- Registrant information such as First Name, Last Name, Organization, physical address, phone number, and e-mail address
- Assigned domain administrator, billing contact, technical contact

Network Scanning

Identify accessible hosts, running services, service and OS versions, …

Active: target network can observe probe requests

- As opposed to passive reconnaissance or querying of public sources
- Stealthiness matters! IDSs can easily detect noisy scans

Two main dimensions

- **Horizontal scanning:** scan a subnet (or the whole internet) on a particular port number
  - E.g., find all hosts running a vulnerable service (internet worms)
- **Vertical scanning:** scan all (or a subset of) ports on a given host
  - Scan common ports first

Manual scanning using `ping` and `netcat` can be used for quick assessments

Nmap

De facto tool for network scanning

Support for many port scan types

-sS	TCP SYN scan: just wait for the ACK
-sT	TCP connect scan: full connection (useful for non-root)
-sU	UDP scan: protocol-specific payload for known ports
-sA	ACK scan: determine if a firewall is stateful
-sO	IP protocol scan: determine IP protocols (TCP, ICMP, IGMP) used
-p	Specify port range (default: 1000 most common ports)

Beyond simple port scanning: extensible framework with support for third-party scripts

```
auth, broadcast, brute, default, discovery, dos,
exploit, external, fuzzer, intrusive, malware, safe,
version, vuln
```

Service Fingerprinting

After identifying that a port is open, try to gather more information about the service

```
# nmap –sV 192.168.0.1 –p 22
```

Complete the connection and attempt to determine the software type and version

- Version detection "interrogates" those ports to determine more about what is actually running

Server-initiated dialog: banner grabbing

- Upon connection, the server transmits a banner string that often includes version information (e.g., SSH)

Client-initiated dialog: send probe application requests

Nmap has about 6,500 dialogue patterns for more than 650 protocols such as SMTP, FTP, HTTP, etc.

Shodan: let others do the scanning for you

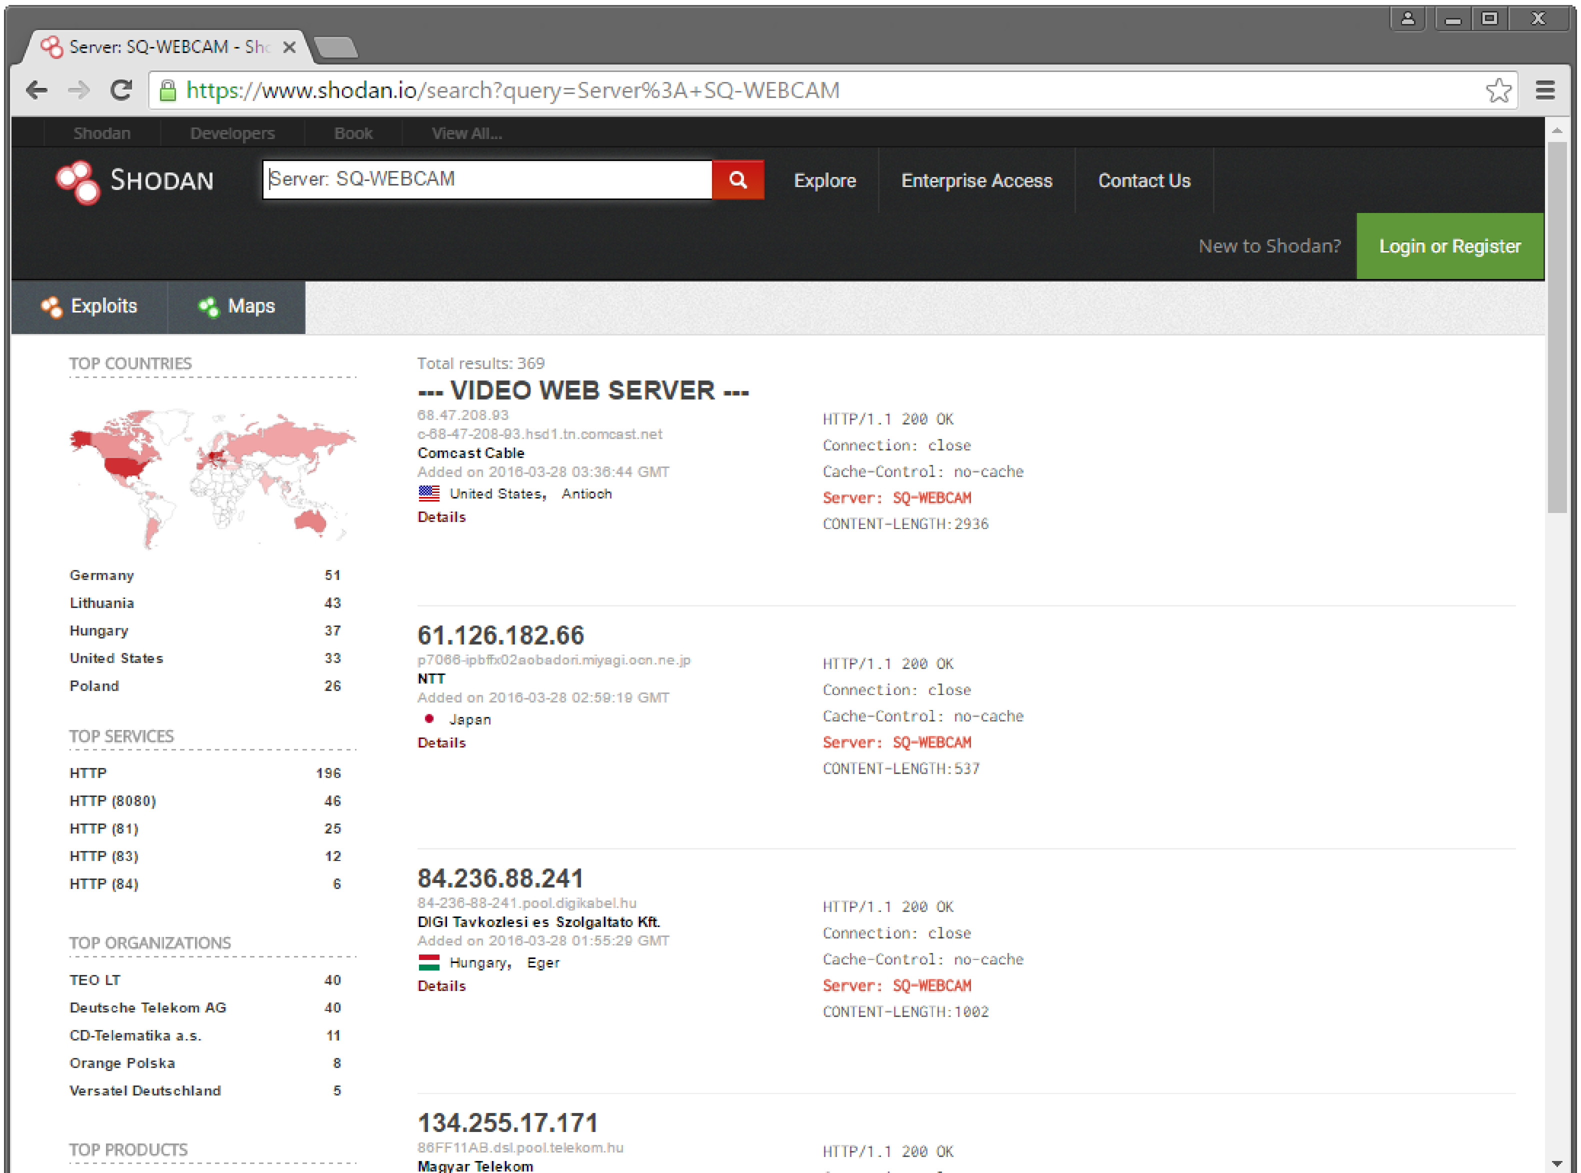

Idle Scan

Hide scan attempts by blaming another "zombie" host

- Zombie must be mostly idle (e.g., network printer)
- Zombie should have sequential/predictable IPID behavior

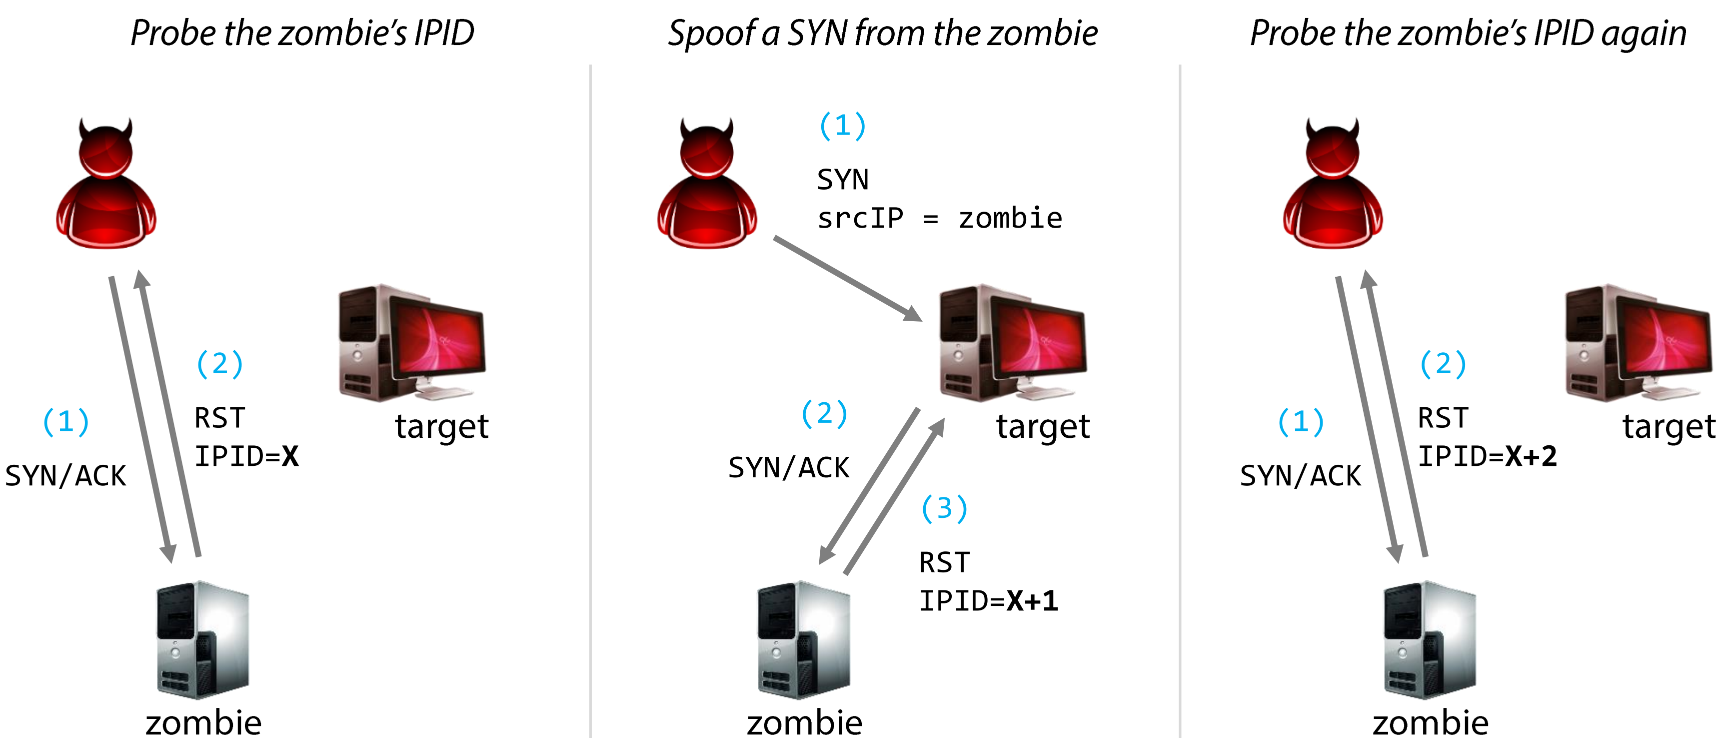

ARP Scan

- Useful technique for host enumeration in a LAN
  - Find every active IPv4 device in the same subnet
- Send a “who has” broadcast packet for each IP address of interest
  - Example: try all 254 host IP addresses for a /24 subnet
  - Retry a couple of times if no response is received
- Linux command-line tool: `arp-scan`

```
# arp-scan 192.168.0.0/24
```

Vulnerability Scanning

Identify vulnerabilities in exposed services

- Typical next step after network scanning
- Exploitable bugs, misconfigurations, default passwords, …

OpenVAS (open-source), Nessus (free/commercial, proprietary), Qualys (commercial), Nexpose (commercial), …

New "vulnerability tests" released every day

- 45,000 in total for OpenVAS as of Feb. 2016

Usually come with user-friendly GUI for configuration, policy management, and report generation

Denial of Service

Goal: harm availability

- Strain software, hardware, or network links beyond their capacity
- Shut down or degrade the quality of a service

Not always the result of an attack

- Flash crowds, “Slashdot effect”

Motives

- Protest/attention
- Financial gain/damage
- Revenge
- Blackmail
- Evasion/diversion

DoS Attack Characteristics

Attack source: single vs. many

- More than a single source: Distributed DoS (DDoS)

Overload vs. complete shutdown

- Degradation vs. completely disabling software or equipment
- Crash, restart, bricking, website defacement, …

Consumed resource

- Network bandwidth, CPU, memory, sockets, disk storage, …

Amplification factor

- Symmetric vs. asymmetric attacks
- Broadcast addresses, large protocol responses, propagation, …

Algorithmic complexity attacks

- Induce worst-case behavior by triggering corner cases

Spoofing

- Hide the true source(s) of the attack

Lower Layer DoS

Physical layer

- Wirecutting, equipment manipulation, physical destruction
- RF jamming, interference

Link Layer

- MAC flooding: overload switch/network
- ARP poisoning: send fake ARP replies to insert erroneous MAC-IP mappings in existing systems' caches
- DHCP starvation
- WiFi Deauthentication

Spectrum Blames Vandal ×

gothamist.com/2017/09/15/hackers_hit_spectrum_mainframe.php

gothamist

TIPS

Spectrum Blames Vandals For Internet Outages In Brooklyn, Queens

BY SCOTT HEINS IN NEWS ON SEP 15, 2017 11:11 AM

137 Like Share Tweet

Damage to Spectrum's fiberoptic cables led to widespread internet loss throughout Queens and Brooklyn this morning. (Getty Images)

GOTHAMIST FILMS

BEST OF GOTHAMIST

The Best Bars In NYC Where You Can Read In Peace

These Are The Best Restaurants In Williamsburg Right Now

The Best Breakfast Eggs

Deauth Attacks

Send a spoofed deauth frame to AP with victims' address (no authentication!)

- Client is disassociated from access point
- Can also use the broadcast address to disassociate all clients
- *They may then connect to an "evil twin" access point…*

Deauthentication is also sometimes used as a protection mechanism

- Prevent the operation of rogue access points

Tools: `aireplay-ng` (aircrack-ng), `deauth` (metasploit)

Also possible: auth attacks

- Flood with spoofed random addresses to authenticate and associate to a target access point → exhaust AP resources

Federal Communications Commission
445 12th Street, S.W.
Washington, D.C. 20554

News Media Information 202 / 418-0500
Internet: http://www.fcc.gov

This is an unofficial announcement of Commission action. Release of the full text of a Commission order constitutes official action. See MCI v. FCC. 515 F 2d 385 (D.C. Cir. 1974).

FOR IMMEDIATE RELEASE:
October 3, 2014

NEWS MEDIA CONTACT:
Neil Grace, 202-418-0506
E-mail: Neil.Grace@fcc.gov

MARRIOTT TO PAY $600,000 TO RESOLVE WIFI-BLOCKING INVESTIGATION

Hotel Operator Admits Employees Improperly Used Wi-Fi Monitoring System to Block Mobile Hotspots; Agrees to Three-Year Compliance Plan

Washington, D.C. –Marriott International, Inc. and its subsidiary, Marriott Hotel Services, Inc., will pay $600,000 to resolve a Federal Communications Commission investigation into whether Marriott intentionally interfered with and disabled Wi-Fi networks established by consumers in the conference facilities of the Gaylord Opryland Hotel and Convention Center in Nashville, Tennessee, in violation of Section 333 of the Communications Act. The FCC Enforcement Bureau's investigation revealed that Marriott employees had used containment features of a Wi-Fi monitoring system at the Gaylord Opryland to prevent individuals from connecting to the Internet via their own personal Wi-Fi networks, while at the same time charging consumers,

Network Layer DoS

Flooding: bombard target with network packets

- Saturate the available network bandwidth (aka "volumetric" attacks)
- Long ICMP packets, UDP/TCP packets with garbage data, …

IP spoofing: conceal the attack source

- Makes it more difficult to block the attack
- Ingress and egress filtering limit its applicability, but not universally deployed
- Applicable only when connection establishment is not needed: ICMP, UDP, TCP SYN, …

Broadcast Amplification

- One packet generates many more packets
- ICMP Smurf Attack (spoofed broadcast Echo request)

IP hijacking

- False BGP route advertisements to attract and drop traffic or cause connectivity instability

Smurf Attack (90's)

Attacker sends spoofed ICMP Echo requests to the victim's network broadcast address

- Src IP == victim's IP
- Victim machine is flooded with responses from all internal hosts
- Initial form of ***amplification***

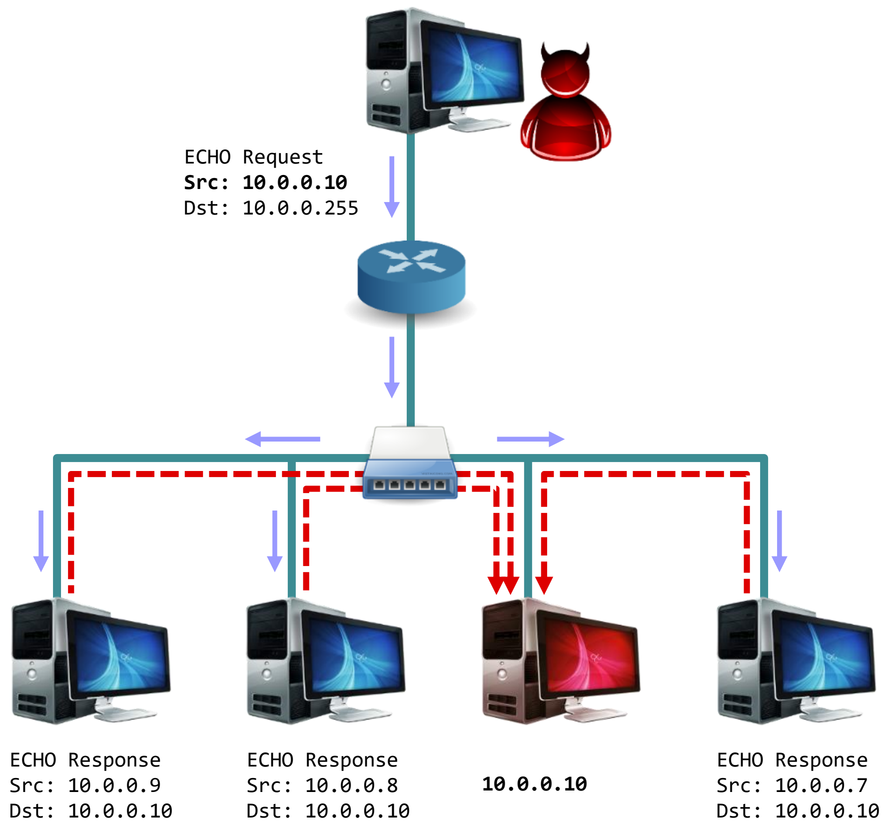

Mitigation

- Configure hosts to not respond to broadcast ICMP requests
- Configure routers to not forward packets destined to broadcast addresses

Transport Layer DoS

SYN flooding

- Server-side resource exhaustion
- Source IP address can be spoofed
- Can be combined with normal flooding to also saturate link

Connection termination

- RST injection
- Mostly used for blocking specific unwanted traffic

SYN Flooding

Flood server with spoofed connection initiation requests (SYN packets)

- Saturate server's max number of concurrent open sockets: no more connections can be accepted
- Each half-open connection consumes memory resources
- Server sends SYN/ACKs back, but ACKs never return…

Mitigation

- Drop old half-open connections after reaching a certain threshold (in FIFO order or randomly)
- SYN cookies: eliminate the need to store state per half-open connection

TCP Connection Termination

FIN: this side is done sending, but can still receive

- "Half-closed" state
- Should be sent by each side and acknowledged by the other

RST: this side is *done sending and receiving*

- No more data will be sent from this source on this connection
- Program closed, abort established connection, …

A MotS attacker can easily send spoofed RST packets

- 5-tuple (src/dst IP/port and protocol) must match
- Sequence number should be *in window*
- More strict stacks will only accept RSTs *in sequence* ➔ Prevent blind TCP RST injection

Legitimate and not so legitimate uses

- Censorship, blocking of non-standard port traffic (e.g., P2P protocols), termination of malicious connections, …

Comcast settles P2P throttling class-action for $16 million

Comcast got itself in hot water when it decided to use reset packets to slow …

by Jacqui Cheng - Dec 22, 2009 4:22pm EST

Comcast has agreed to settle a class-action lawsuit over the throttling of P2P connections that had users up in arms in late 2007 and 2008. The company still stands behind its controversial methods for "managing" network traffic, but claims that it wants to "avoid a potentially lengthy and distracting legal dispute that would serve no useful purpose."

It was more than two years ago when Comcast subscribers began finding evidence that the broadband provider was blocking packets—particularly those being sent through BitTorrent. When the complaints mounted, the Associated Press went ahead with its own investigation and came to the same conclusion: downloads through BitTorrent were either being blocked altogether or being slowed down significantly.

At that time, Comcast vehemently denied that it had anything to do with these mysterious slowdowns. This was despite the fact that numerous customers reported that their Comcast connections were sending reset packets out to the rest of the Internet—the AP discovered that nearly half of the reset packets being received by cable competitor Time Warner were coming from Comcast. Eventually, Comcast acknowledged that it had engaged in "traffic management" techniques in order to keep its network speedy, which eventually resulted in an FCC investigation and a subsequent abandoning of its P2P batin' ways in favor of a more neutral congestion management system.

Application Layer DoS

Connection flooding

Reflection

Software vulnerabilities

- Just crash the client/server if exploitation fails/is not possible

Algorithmic complexity attacks

- Trigger worst-case processing (e.g., hashtable collisions, regular expression backtracking)

Exhaustion of server resources

- Example: fill up FTP server with junk files

Spam can be considered as a DoS attack on our time…

- And server resources

Connection Flooding

Saturate the server with many established connections

- Can't use spoofing: just use bots…

For forking servers, the whole system might freeze (process exhaustion)

Slowloris attack: slowly send a few bytes at a time to keep many concurrent connections open

- Keep the server busy with "infinite-size" HTTP requests by periodically sending more and more bogus HTTP headers
- Alternatives: read response slowly, POST data slowly, …
- Requires minimal bandwidth

Amplification/Reflection Attacks

Like the ICMP Smurf attack

Abuse network services that reply to certain requests with *much larger* responses

- Attacker sends a *small* packet with a forged source IP address
- Server sends a *large* response to the victim (forged IP address)

UDP: connectionless protocol ➔ easy to spoof

Used by many services:

- NTP, DNS, SSDP, SNMP, NetBIOS, QOTD, CharGen, …

Technical Details Behind a

https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/

Technical Details Behind a 400Gbps NTP Amplification DDoS Attack

13 Feb 2014 by Matthew Prince.

On Monday we mitigated a large DDoS that targeted one of our customers. The attack peaked just shy of 400Gbps. We've seen a handful of other attacks at this scale, but this is the largest attack we've seen that uses NTP amplification. This style of attacks has grown dramatically over the last six months and poses a significant new threat to the web.

CloudFlare features

- Overview
- CDN
- Optimizer
- Security

Amplification Factor

Protocol	BAF all	BAF 50%	BAF 10%	PAF all	Scenario
SNMP v2	6.3	8.6	11.3	1.00	*GetBulk* request
NTP	556.9	1083.2	4670.0	3.84	Request client statistics
DNS_{NS}	54.6	76.7	98.3	2.08	ANY lookup at author. NS
DNS_{OR}	28.7	41.2	64.1	1.32	ANY lookup at open resolv.
NetBios	3.8	4.5	4.9	1.00	Name resolution
SSDP	30.8	40.4	75.9	9.92	*SEARCH* request
CharGen	358.8	n/a	n/a	1.00	Character generation request
QOTD	140.3	n/a	n/a	1.00	Quote request
BitTorrent	3.8	5.3	10.3	1.58	File search
Kad	16.3	21.5	22.7	1.00	Peer list exchange
Quake 3	63.9	74.9	82.8	1.01	Server info exchange
Steam	5.5	6.9	14.7	1.12	Server info exchange
ZAv2	36.0	36.6	41.1	1.02	Peer list and cmd exchange
Sality	37.3	37.9	38.4	1.00	URL list exchange
Gameover	45.4	45.9	46.2	5.39	Peer and proxy exchange

TABLE III: Bandwidth amplifier factors per protocols. *all* shows the average BAF of all amplifiers, 50% and 10% show the average BAF when using the worst 50% or 10% of the amplifiers, respectively.

Christian Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse – NDSS'14

Distributed Denial of Service (DDoS)

Any DoS attack that originates from multiple sources

- Early internet worms were the first instances of DDoS

These days usually launched by botnets

- Networks of compromised systems ("bots") awaiting commands by an attacker ("botmaster")
- Not only PCs/servers: mobile and IoT devices equally useful
- Can be rented through online marketplaces ("booter" or "stresser" services)

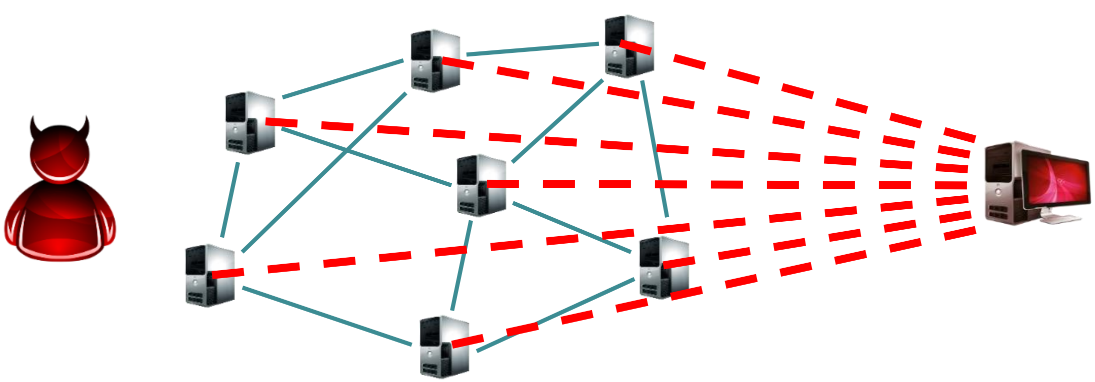

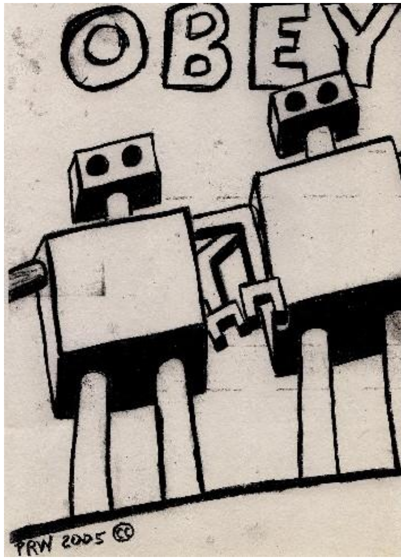

DDoS attack on Dyn cam ×

Secure | https://www.computerworld.com/article/3135434/security/ddos-attack-on-dyn-came-from-100000-infected-devices.html

NEWS

DDoS attack on Dyn came from 100,000 infected devices

DNS service provider Dyn says Mirai-powered botnets were the primary source for Friday's disruption

By Michael Kan

U.S. Correspondent, IDG News Service | OCT 26, 2016 2:21 PM PT

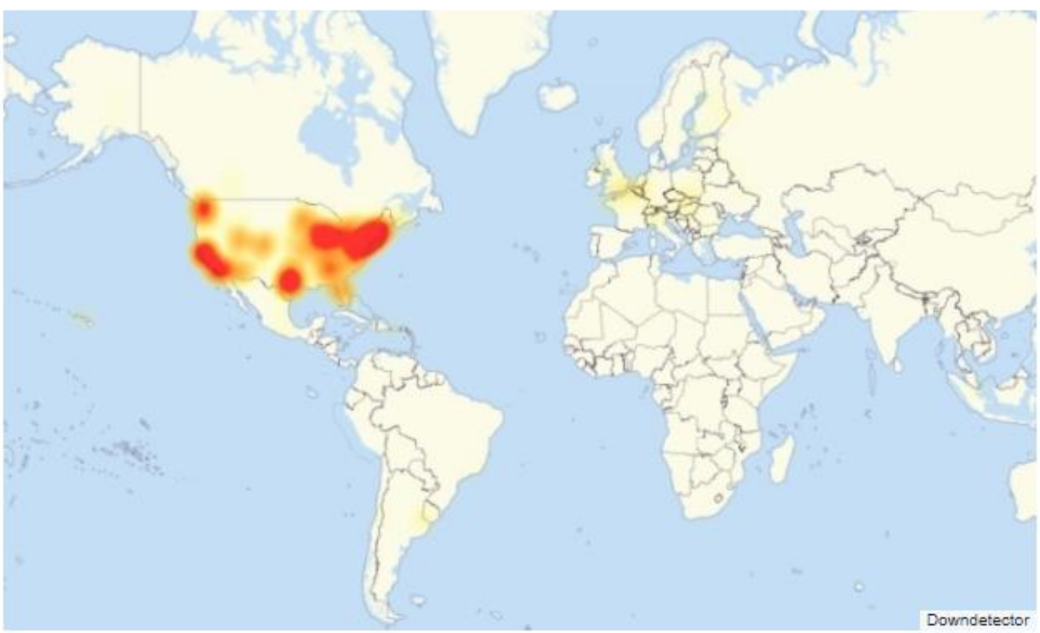

MORE LIKE THIS

Chinese firm admits its hacked products were behind Friday's DDOS attack

An IoT botnet was partly behind Friday's massive DDOS attack

DDoS attack with Mirai malware 'killing business' in Liberia

VIDEO

Tech Talk: Pricey iPhones, intent-based networks, GPS spoofing and smartwatches

Puppetnets: Browser-based Bots

Browsers can be indirectly misused to attack others

- JS code running in the browsers of unsuspecting visitors

Continuously fetch images or other large files from the victim's server

- Can masquerade as "good" bots (e.g., Googlebot, Baiduspider, other legitimate spiders) using a spoofed User-Agent

Many injection ways

- Compromised websites
- Ad networks
- MitM/MotS attacks

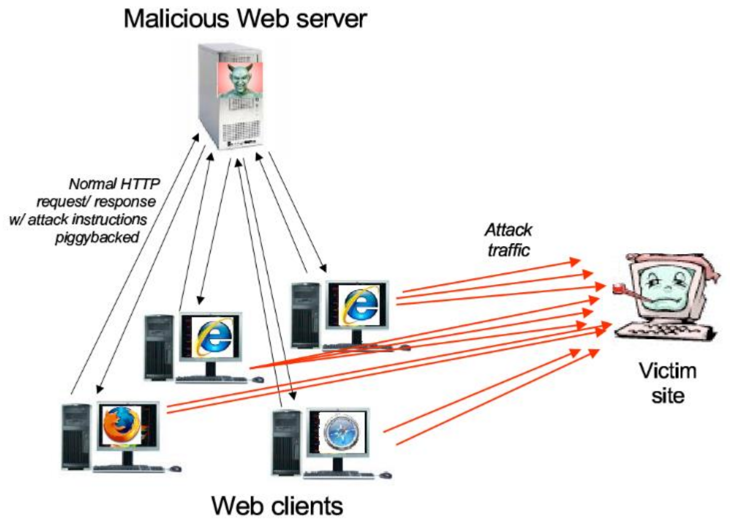

The Pirate Bay is Pirating

Secure | https://cointelegraph.com/news/the-pirate-bay-is-pirating-your-processor-for-bitcoin-mining

THE COINTELEGRAPH
future of money

BTC $ 3,960.18 0.82 %
BCH $ 486.97 -8.01 %
ETH $ 288.34 1.8 %

Submit a post
Submit press release

News | Guides & Analytics | Events | Explained | ICO Calendar

By Darryn Pollock

SEP 18, 2017

The Pirate Bay is Pirating Your Processor for Bitcoin Mining

22977 Total views 332 Total shares

GITHUB ATTACK PERPETRATED BY CHINA'S GREAT CANNON TRAFFIC INJECTION TOOL

by Brian Donohue Follow @TheBrianDonohue April 10, 2015 , 1:06 pm

Chinese attackers used the Great Firewall's offensive sister-system, named the Great Cannon, to launch a recent series of distributed denial of service attacks targeting the anti-censorship site, GreatFire.org, and the code repository, Github, which was hosting content from the former.

The first set of DDoS attacks hit GreatFire.org on March 16. On March 26, Github

Top Stories

Critical Yahoo Mail Flaw Patched, $10K Bounty Paid
January 19, 2016 , 10:02 am

BlackEnergy APT Group Spreading Malware via Tainted Word Docs
January 28, 2016 , 7:00 am

Curious Tale of a Microsoft Silverlight Zero Day
January 13, 2016 , 9:01 am

Oracle to Kill Java Browser Plugin
January 28, 2016 , 12:43 pm

Apple's 'Targeted' Gatekeeper Bypass Patch Leaves OS X Users Exposed
January 15, 2016 , 8:00 am

Data Theft Hole Identified in LG G3 Smartphones

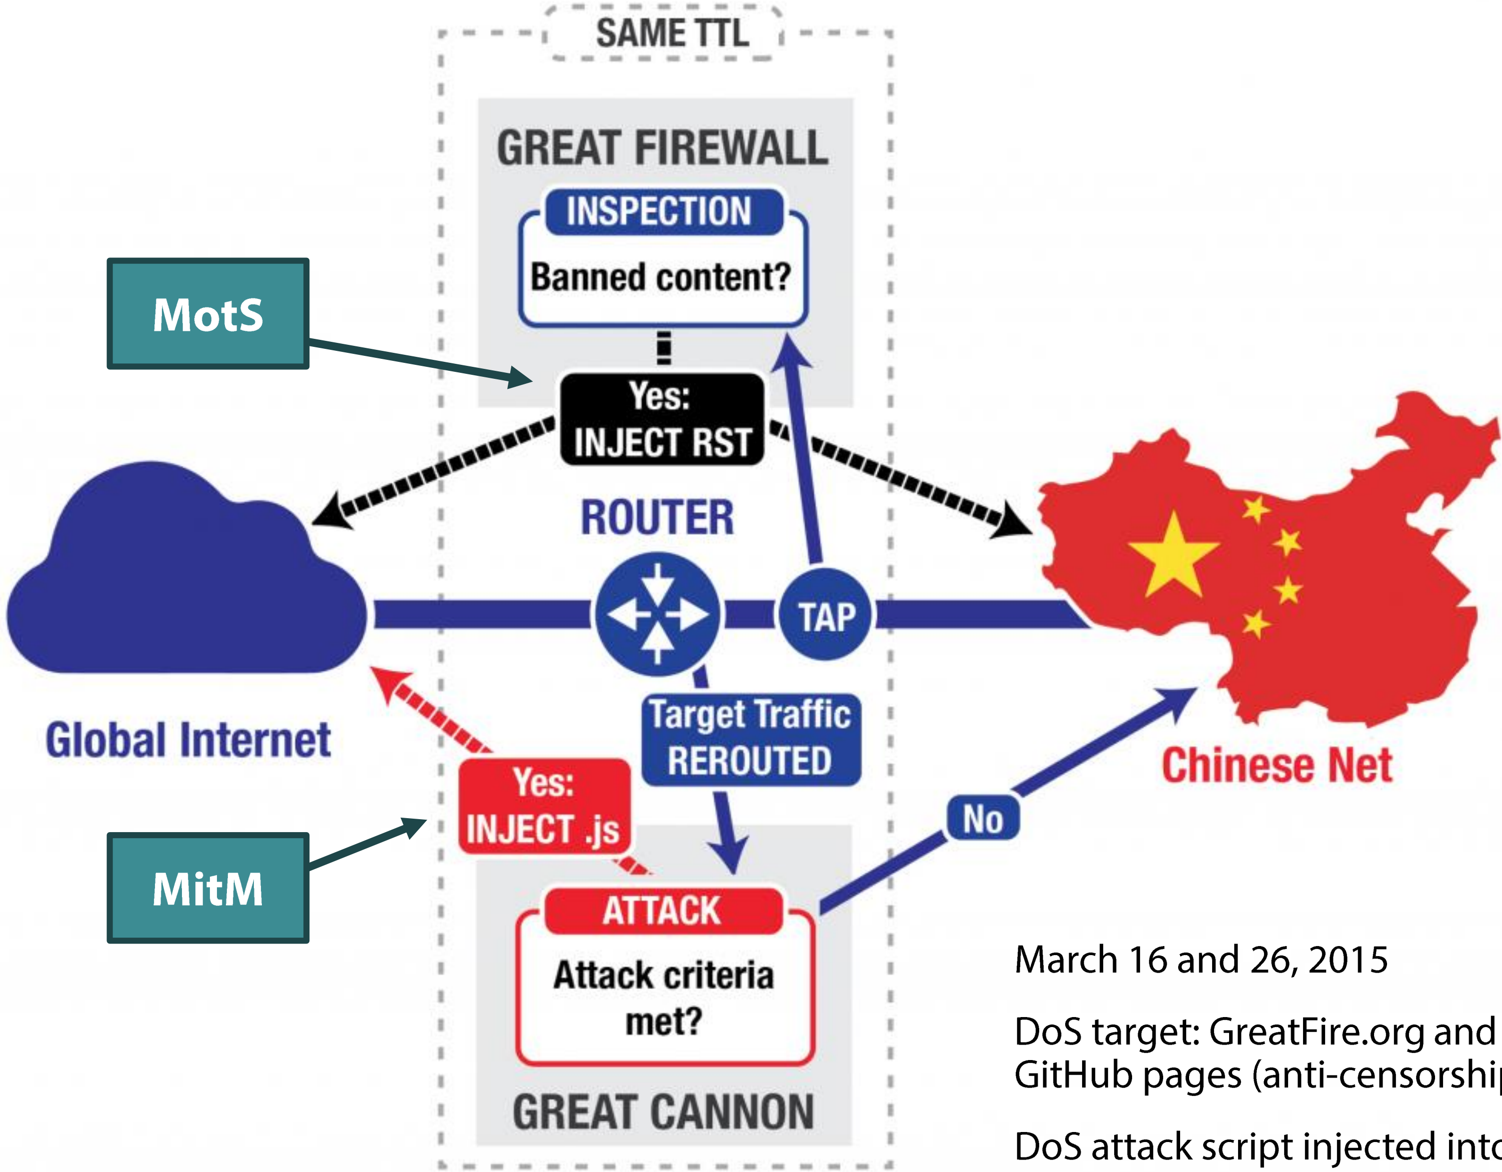

March 16 and 26, 2015

DoS target: GreatFire.org and two related GitHub pages (anti-censorship project)

DoS attack script injected into 1.75% of the requests to Baidu's analytics/ad scripts (probabilistic injection)

China's Great Cannon - https://citizenlab.org/2015/04/chinas-great-cannon/

DoS Defenses

No absolute solution

- Asymmetry: little effort for the attacker, big impact for the victim
- Any public service can be abused by the public
- Prank phone calls, road blockades, …

General strategies

- Filter out bad packets
- Improve processing of incoming data
- Hunt down and shut down attacking hosts
- Increase hardware and network capacity

DoS Defenses

Ingress/egress filtering

Ensure that incoming/outgoing packets actually come from the networks they claim to originate from ➔ drop spoofed packets

Content delivery networks (CDNs) and replication

Distribute load across many servers

Client challenges

Present a CAPTCHA whenever the system is under stress

Other (mostly academic) approaches

IP Traceback: each router "marks" with its own IP the forwarded packets to facilitate determining the actual origin of packets

Pushback filtering: iteratively block attacking network segments by notifying upstream routers

Overlay-based systems: proactive defense based on secure overlay tunneling, hash-based routing, and filtering

To continue, please type the characters below:

Submit

About this page

Our systems have detected unusual traffic from your computer network. This page checks to see if it's really you sending the requests, and not a robot. Why did this happen?