DDOS MITIGATION FUNDAMENTALS

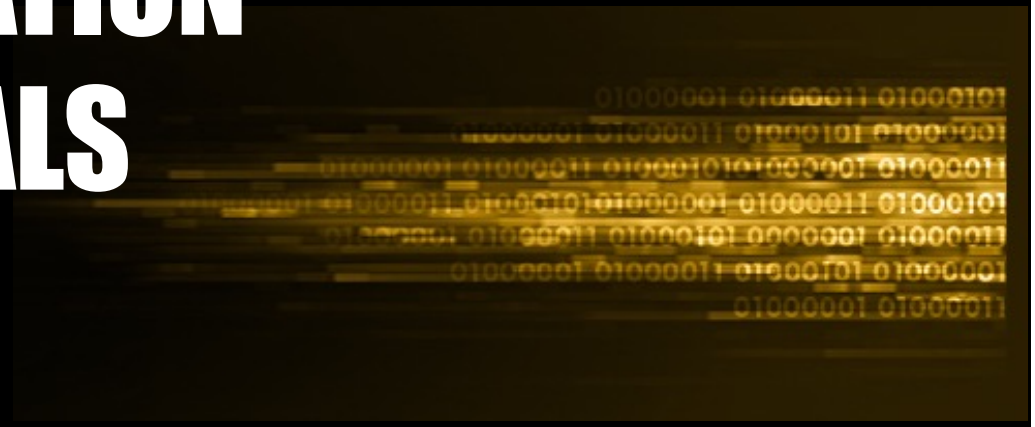

PURDUE UNIVERSITY

Introduction

- Who am I?
  - Graduate student at Purdue University
  - Director of Security Engineering for Hydrolix
  - krassi@purdue.edu // krassi@hydrolix.io
  - https://www.linkedin.com/in/krassi/

PURDUE UNIVERSITY

THE CYBER TAP

hydrolix

- Who are you?
- Let's make it interactive!

- No pictures, please!

WHAT IS DOS/DDOS?

Terminology and general concepts

What is a Denial of Service attack?

- Discussion: so what is a DDoS attack?
- Resource exhaustion… which leads to lack of availability
- Consider:
  - How is it different from a major website highlighting a small one and the resulting traffic?
  - How is that different from company's primary Internet connection going down?

What is a Denial of Service attack?

- From security point of view?
  Decreased availability
- From operations point of view?
  An outage
- From business point of view?
  Financial losses

DoS vs DDoS

- What is the difference?
  - The traffic originator – one system vs many systems
  - Consider reflected attacks
- How does that change the attacks volume?
  - More systems – more capacity

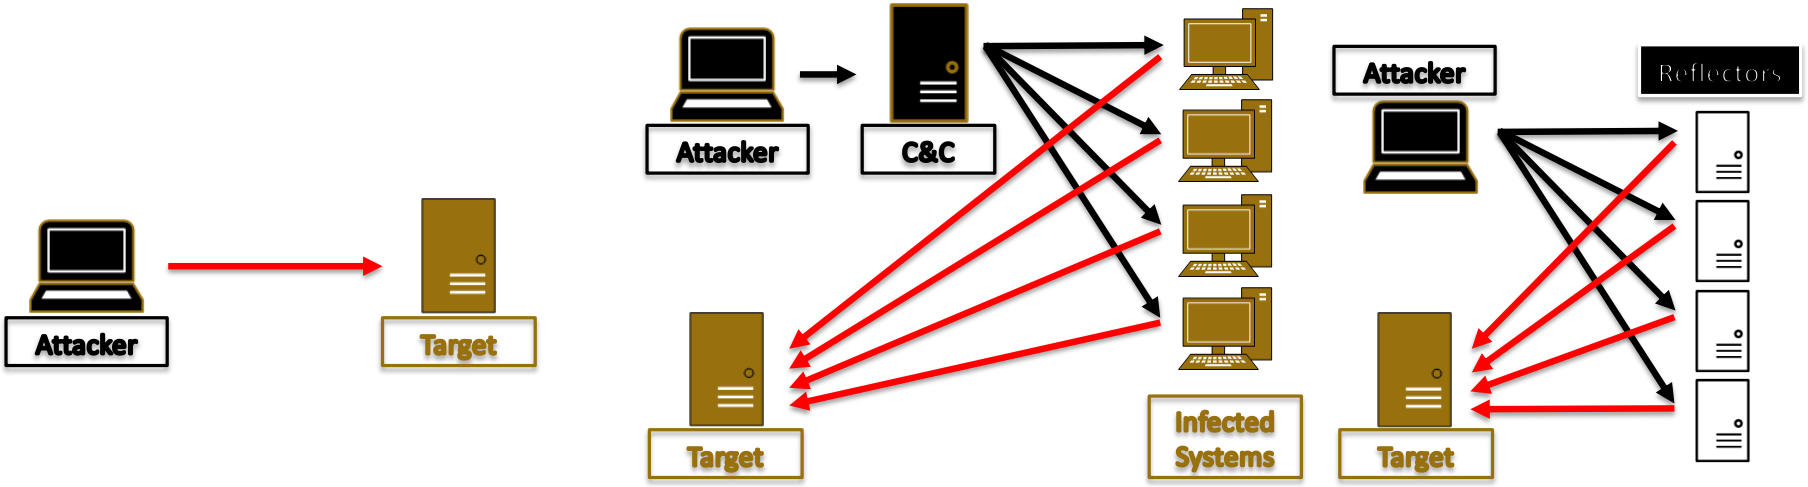

DDOS VOLUME FACTORS

Trends in DDoS growth?

Additional factors supporting and accelerating DDoS

- Overall bandwidth
- **IOT/Embedded home and SOHO devices**
- Booters/Stressors (lowers threshold)
- Reflectors (and ability to spoof source IP)
- Content management systems
- Accessible information

Home routers

- Embedded/IoT devices
  - Default username/password
  - Open DNS recursive resolvers
  - Software bugs (NetUSB)
  - Network diagnostic tools
  - Some do not allow the user to turn off DNS
- XBOX and Sony attacks over Christmas (2014)
  - Lizard Stresses, 2015
  - Mirai, 2016

Compromised Content Management Systems (CMS)

- Most targeted Content Management Systems:
  - WordPress
  - Joomla
- Started in early 2013 - notably around the attacks against US financial institutions
- Now it is an easy way to build a botnet and other groups abuse it as well

Booters/Stressors

- Inexpensive
- Popular among gamers
- Tools are sold for cheap on the black market (forums)
- Range 5-10 Gbps and up to 40GBps
- Usually short duration

Low cost resulting from reflection

- Comparing Botnet driven vs Reflections DDoS

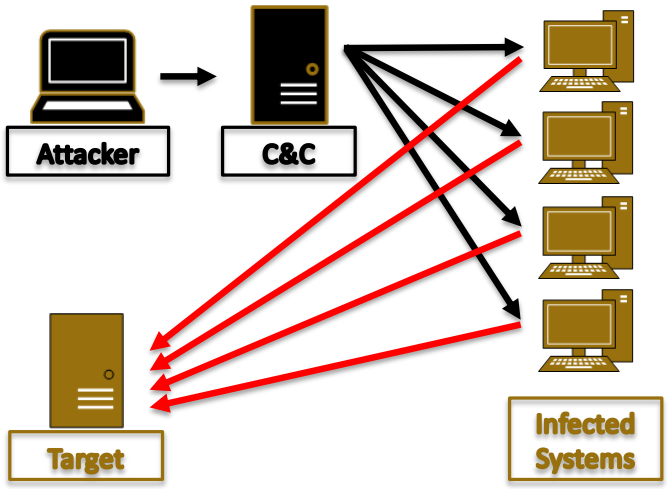

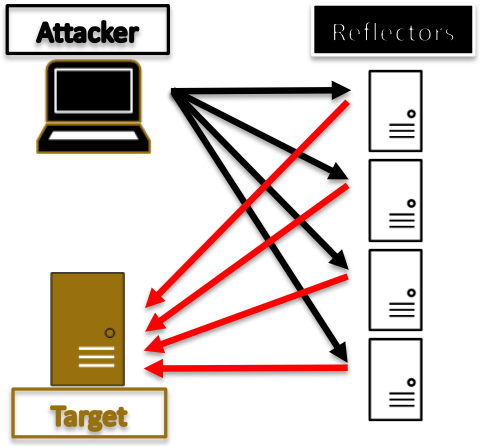

QUESTIONS?

PURDUE UNIVERSITY

THE ADVERSARY

PURDUE UNIVERSITY

The People

- Wide range of attackers
  - Professional DDoS operators
  - Booters/stressors
  - Some of the attacks have been attributed to nation states
  - Hacktivists – though not recently
  - …and more.

Motivation

- Wide range of motivating factors as well
  - Financial gain

    extortion (DD4BC/Armada Collective/copy cats)

    taking the competition offline during high-gain events (online betting, superbowl, etc).
  - Political statement
  - Divert attention (seen in cases with data exfiltration* or financial fraud)
  - Disable firewalls (WAF)
  - Immature behavior

BOOTERS: MO AND TTPS

Booter Services

- Gained popularity since around 2015
- Mostly reflected attack (no need for additional infrastructure)
- Mostly computer gaming industry related
  - Short, bursty attacks
  - Rudimentary scripts
- Fairly inexpensive

Variety of service packages

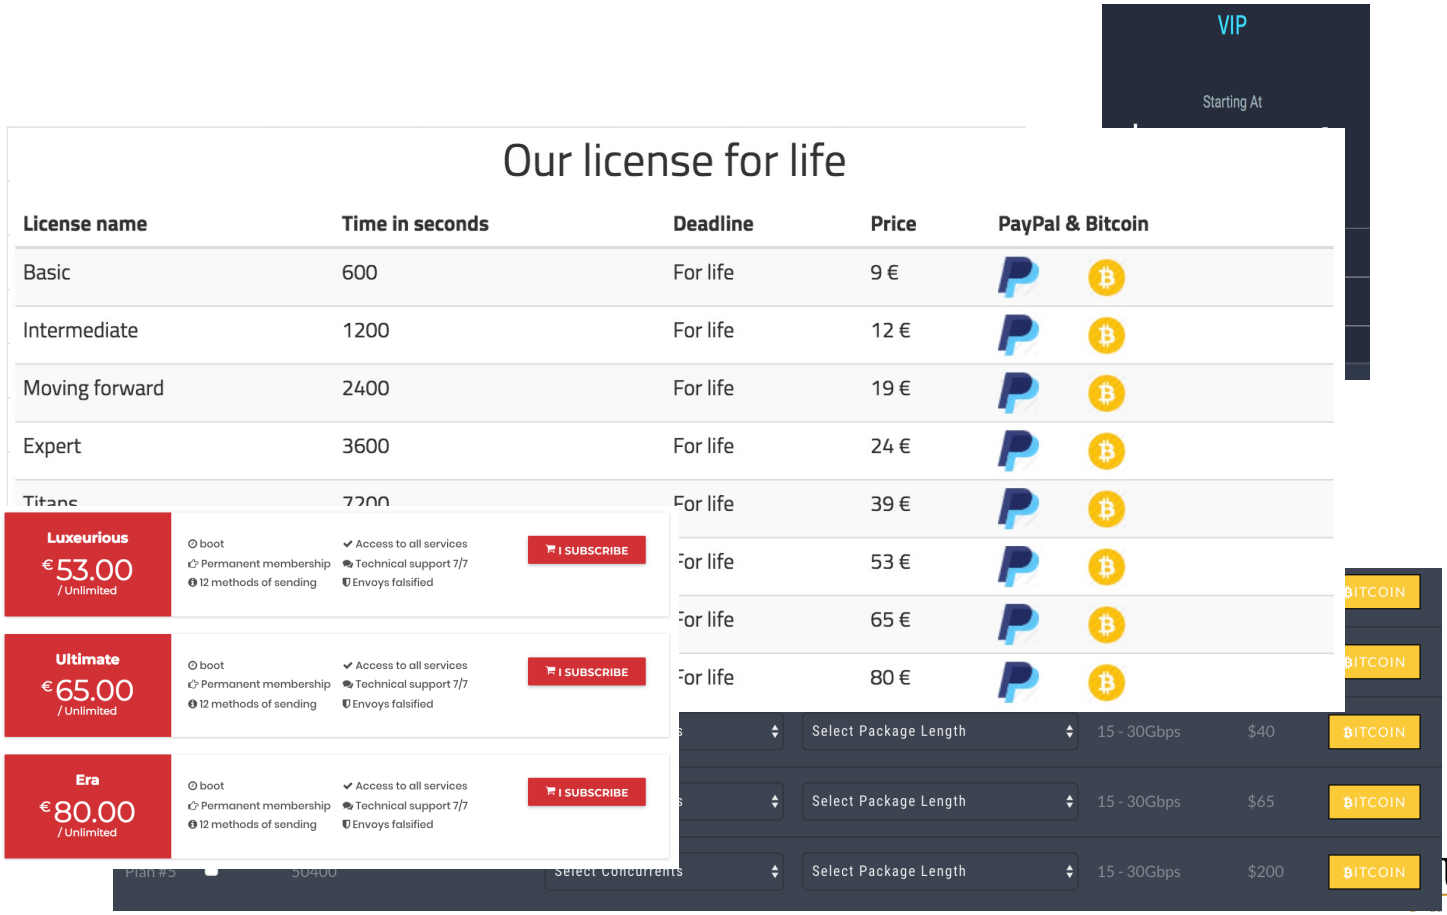

Functionality

- Fancy dashboard
- Different attack types
- Network tools, etc.

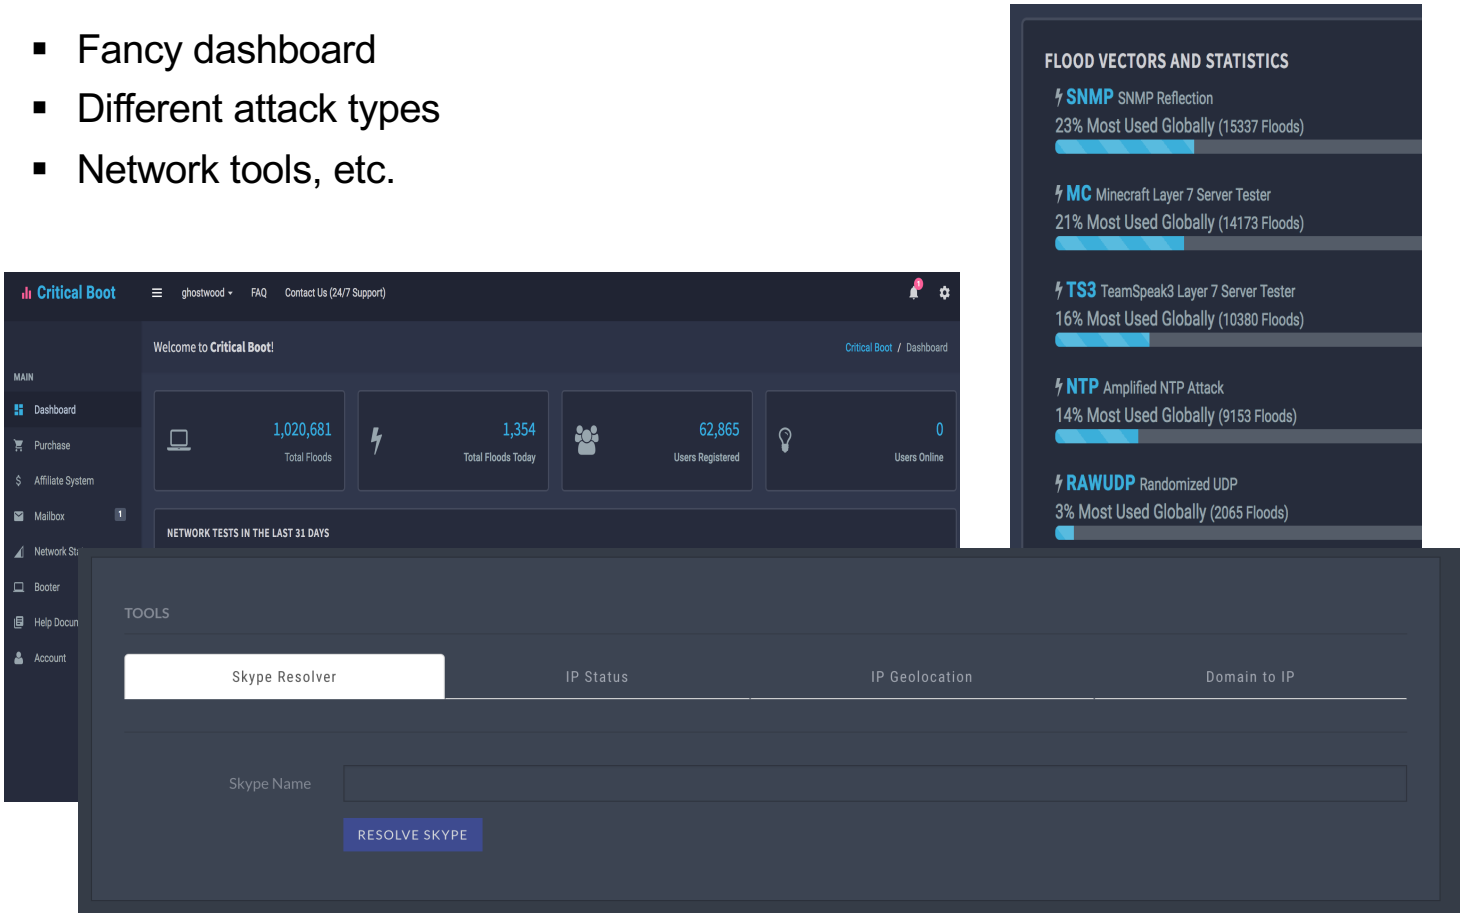

PURDUE UNIVERSITY

Code reuse

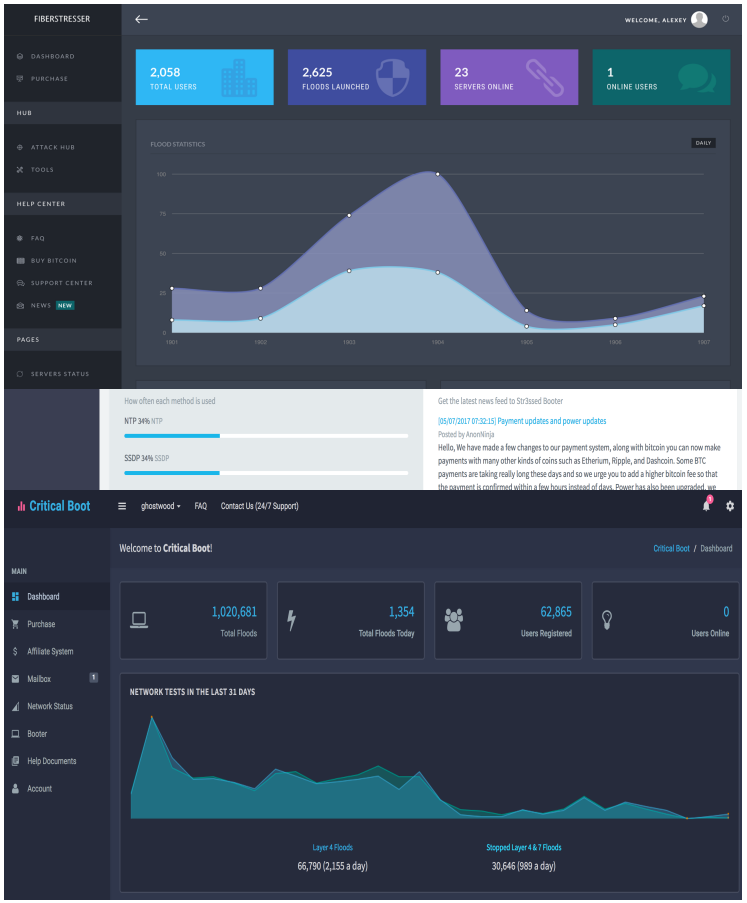

- Individual attack scripts reused widely
- Also some operators set multiple front end sites

PURDUE UNIVERSITY

QUESTIONS?

PURDUE UNIVERSITY

SYN FLOOD

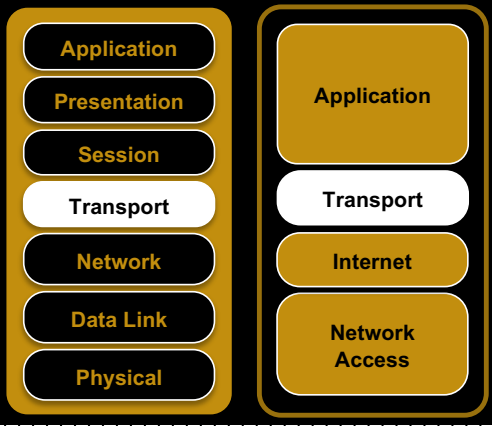

PURDUE UNIVERSITY

Three way handshake?

- What is a 3-way handshake?

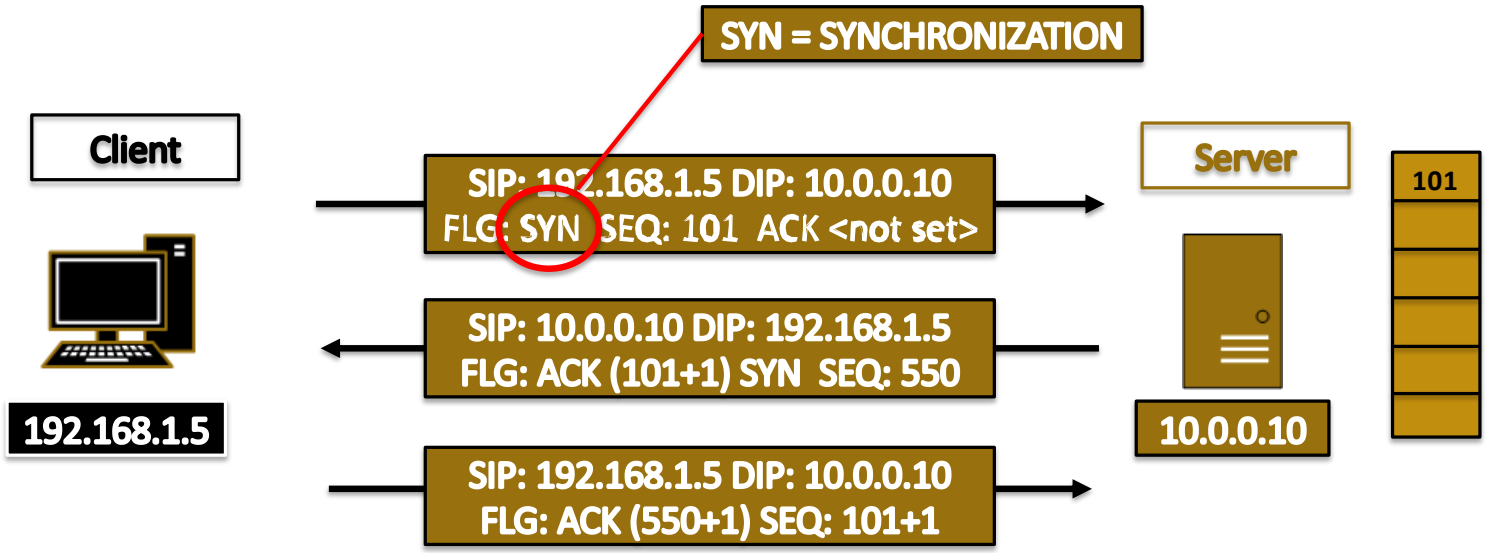

What is a SYN flood?

- The attacker exploits the pending connections queue size (backlog)

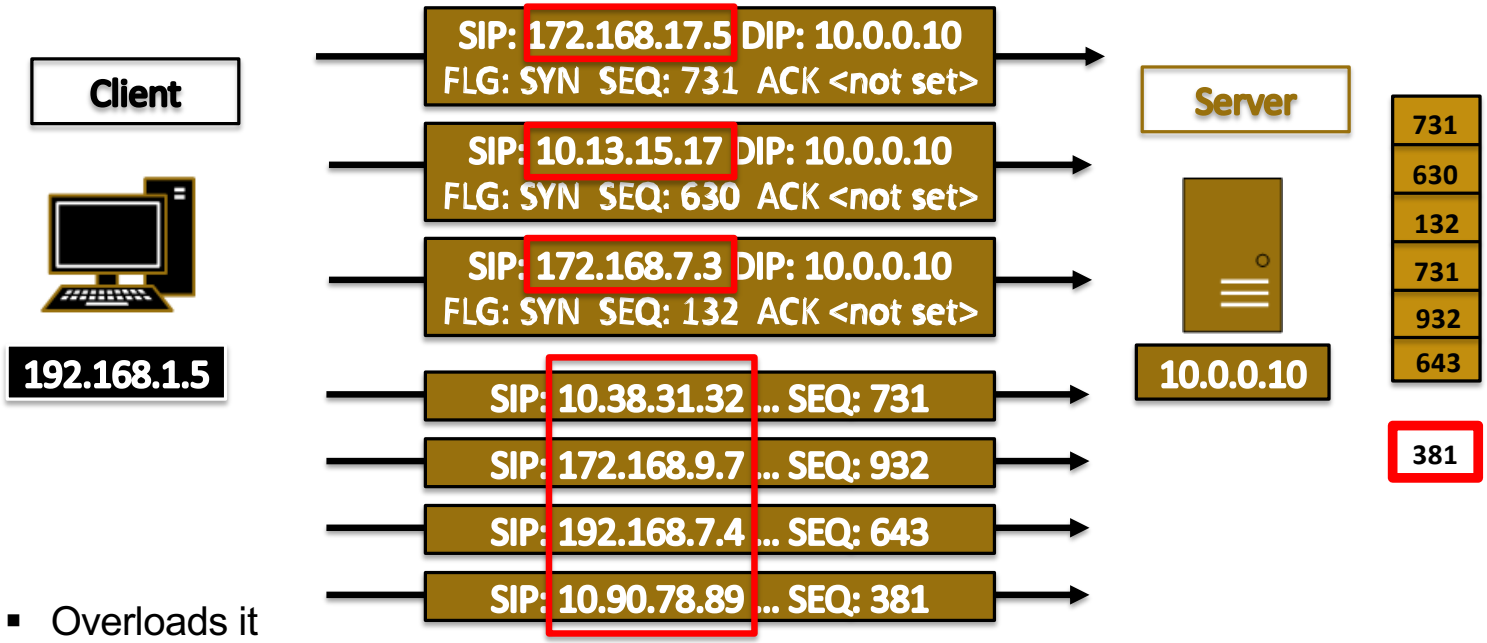

- Overloads it
- To combat easy mitigation the sender randomizes the source IP address, usually using PRNG
- If tcp_abort_on_overflow is set, it will return RST, instead of ACK

PURDUE UNIVERSITY

Listen backlog queue

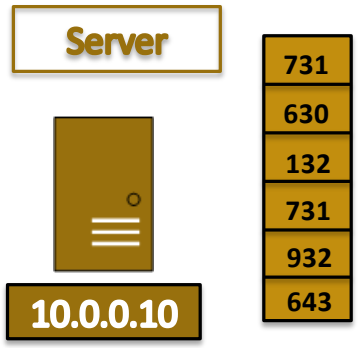

- Connection queue semantics
  - BSD: behaves as if there is only one queue
  - Linux: two queues. In kernel 2.2 the backlog queue also holds ESTABLISED connections which have not been “accepted” by the application
- Size
  - /proc/sys/net/ipv4/tcp_max_syn_backlog – limits the kernel size of the table per socket (4.18.0 defaults to 128)
  - /proc/sys/net/core/somaxconn – limits the backlog argument in the listen() syscall (default 128)
- Tuning up helps with busy servers

Math is hard, let's go shopping

- How much bandwidth does one need to send to enough packets to saturate the queue?
  - Backlog queue size?
    for this example, assume 1000
  - Backlog SYNRECV timeout?
    60 seconds
  - SYN packet size?
    84 bytes (64 bytes + IPG)
- If you are still here (and didn't go shopping):
  - 1000 pkts per minute (~16 pps)
  - 1.4kbps
- What's the effect on lowering the timeout?
- What's the effect of increasing the backlog?

SYN flood through the eyes of netstat

netstat -nap

Active Internet connections (servers and established)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State	PID/Program name
tcp	0	0	0.0.0.0:111	0.0.0.0:*	LISTEN	1339/rpcbind
tcp	0	0	0.0.0.0:33586	0.0.0.0:*	LISTEN	1395/rpc.statd
tcp	0	0	192.168.122.1:53	0.0.0.0:*	LISTEN	1962/dnsmasq
tcp	0	0	127.0.0.1:631	0.0.0.0:*	LISTEN	1586/cupsd
tcp	0	0	127.0.0.1:25	0.0.0.0:*	LISTEN	2703/sendmail: acce
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49718**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49717**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49722**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49720**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49719**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49721**	**SYN_RECV**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:49716**	**SYN_RECV**	**-**

Mitigation: What is a SYN cookie?

- Preserves information in ISN (initial sequence number)
- SYN Cookie:
  - Timestamp % 32 (5 bits)+ MSS (3 bits) + 24-bit hash
  - Components of 24-bit hash:
    - server IP address
    - server port number
    - client IP address
    - client port
    - timestamp >> 6 (64 sec resolution)
  - Maximum Segment Size (MSS) – 8 common values
- All TCP options (but MSS) are lost
- Generated, only when the backlog is full

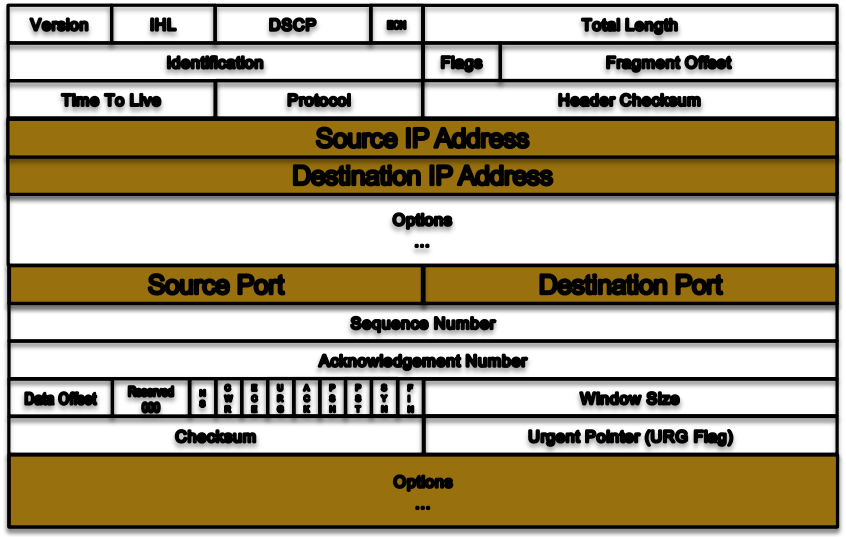

SYN cookie flow

- First SYN packet:
  - Source IP/port
  - Destination IP/port
  - TCP options
  - Time of arrival
  - Initial sequence number: 200
  - Calculate SYN cookie: **345**
- Second packet SYN/ACK:
  - ACK = ISN+1 (200+1)
  - ISN reverse: **345** (== SYN cookie)
- Third packet:
  - ACK = ISN + 1 (501) (== SYN Cookie + 1)
  - SEQ: 101+1 (102 -1 => ISN)

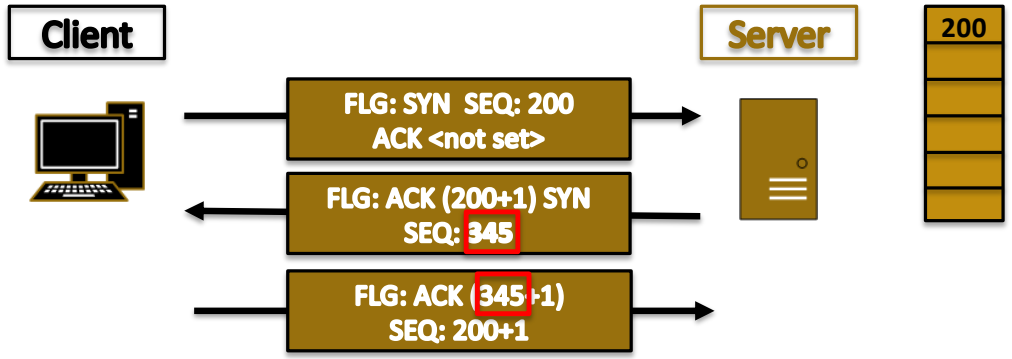

QUESTIONS?

PURDUE UNIVERSITY

REFLECTION AND AMPLIFICATION

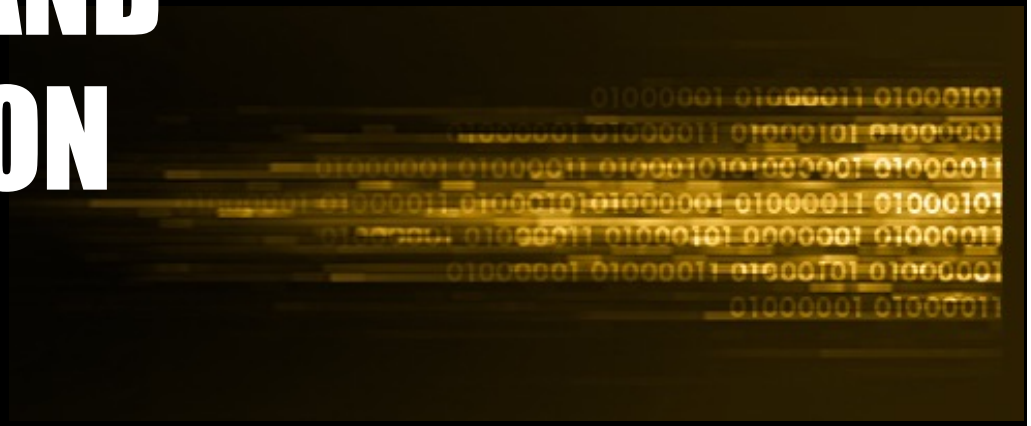

PURDUE UNIVERSITY

Two different terms

- Amplification
  ability to deliver larger response than the query traffic
- Reflection
  using an intermediary to deliver the attack traffic

REFLECTION

Reflection attacks

- A class of attacks, where an unwilling intermediary is used to deliver the attack traffic.
- The attacker would normally send a packet with a forged source IP address to an intermediary. The forged source address is going to be the one of the target. The intermediary will respond and this packet will go to the target instead of the attacker
- Usually those would use the UDP transport level protocol since it does not have the notion of a session. However, there are exceptions like the reflected SYN attacks.

Reflected attack

- Unwilling intermediary, called Reflector, is used to deliver the attack traffic
- Attacker sends a packet with a spoofed source IP set to the victim's IP
- Reflectors responds to the request and send the response to the victim

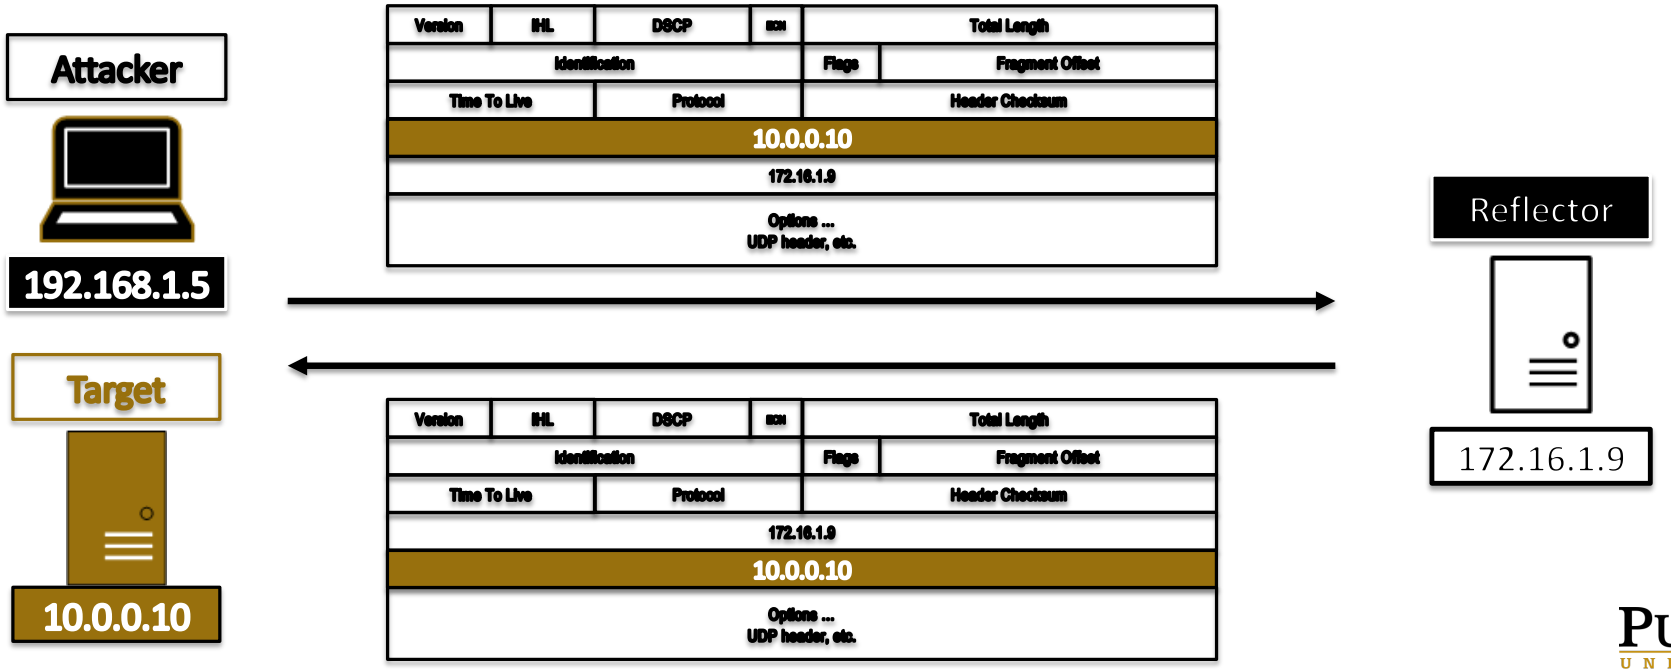

Protocols predisposed to reflection

- Discussion: What protocols do you think can be used for reflected attacks?
- Currently abused:
  - DNS
  - NTP
  - SSDP

AMPLIFICATION

PURDUE UNIVERSITY

Amplification

- Amplification occurs when the query size is disproportionately smaller than the answer size. In such a case the attacker can send a query in a small packet and elicit a large number of packets as response.
- Consider the following DNS query:

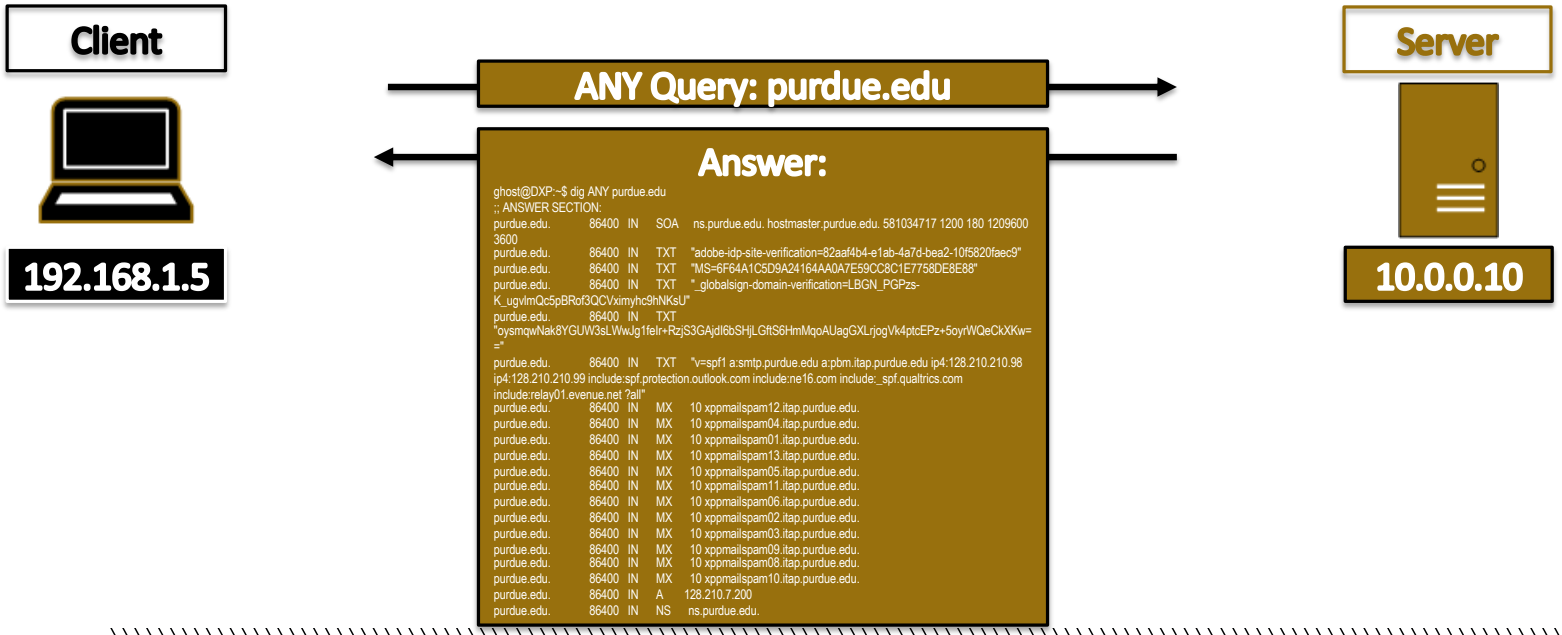

Amplification: Let's do the math

- What's the smallest DNS query possible?
  - Let's build the packet
    - IP Header
    - UDP Header
    - Query
  - What is the minimum legal frame size?
- Let's see it in action:

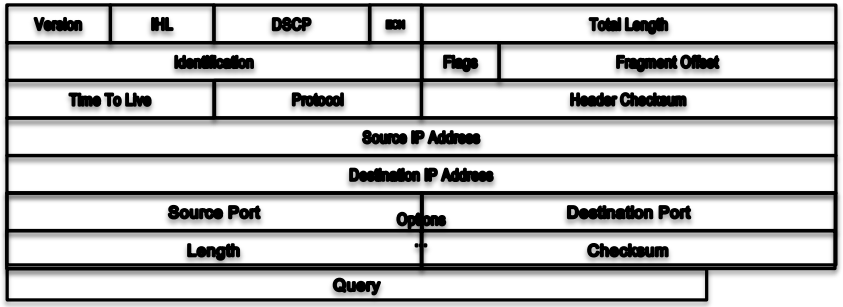

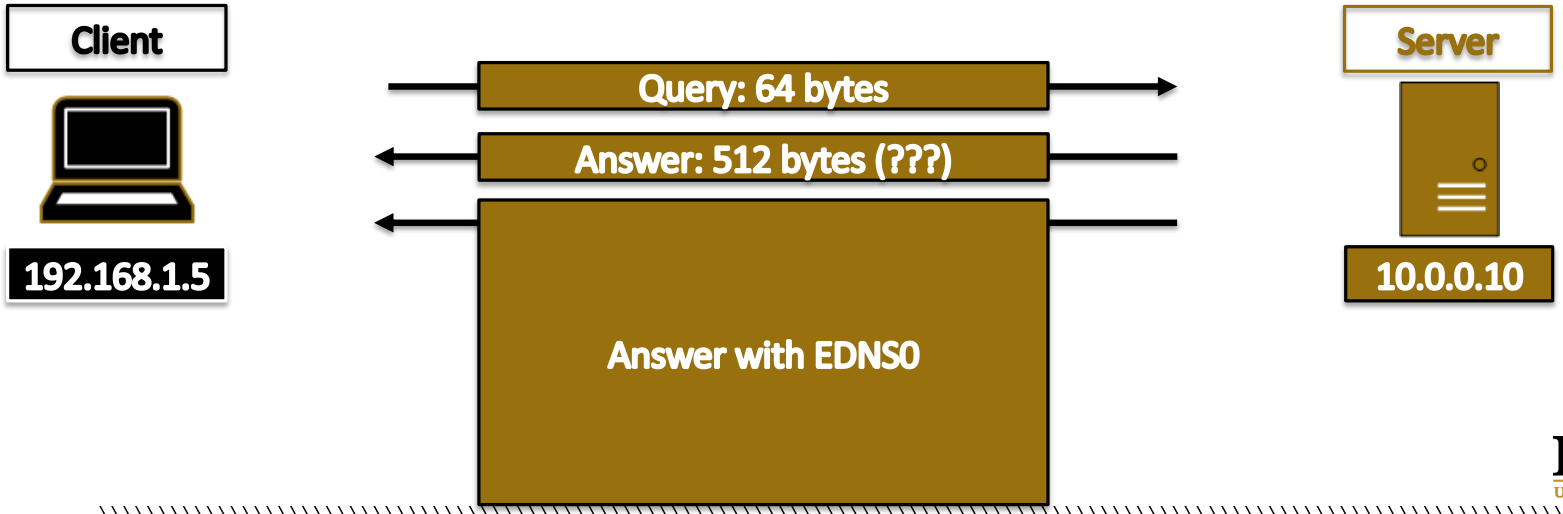

Other amplifier types

- The ones that are of interest and provide amplifications are:
  - DNS
  - SSDP
  - NTP

- Amplification factors:

  https://www.us-cert.gov/ncas/alerts/TA14-017A

Amplification quotients

Protocol	Bandwidth Amplification Factor	Vulnerable Command
DNS	28 to 54	Multiple
NTP	556.9	Multiple
SNMPv2	6.3	GetBulk request
NetBIOS	3.8	Name resolution
SSDP	30.8	SEARCH request
CharGEN	358.8	Character generation request
QOTD	140.3	Quote request
BitTorrent	3.8	File search
Kad	16.3	Peer list exchange
Quake Network Protocol	63.9	Server info exchange
Steam Protocol	5.5	Server info exchange
RIPv1	131.24	Malformed request
CLDAP	56 to 70	
Memcached	10,000 to 51,000	

Source: US-CERT: https://www.us-cert.gov/ncas/alerts/TA14-017A

DNS REFLECTION

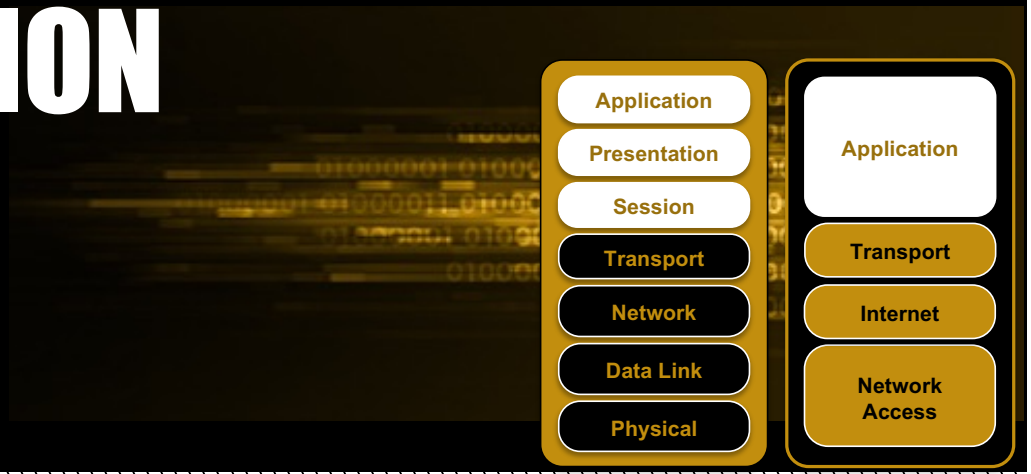

PURDUE UNIVERSITY

QTYPE=ANY

- Consider this query:
  - dig ANY purdue.edu
- What is an ANY query?
  - Query type “*”.
  - All Resource Records (RR) pertaining to a Fully Qualified Domain Name (FQDN).
  - May return any of the RR types (SOA, NS, A, AAAA, MX, TXT, RSIG, etc).
- Different from a zone transfer (AXFR/IFXR).

Consider this (cont'd)


```
ghost@DXP:~$ dig ANY purdue.edu
;; ANSWER SECTION:
purdue.edu.          86400   IN     SOA    ns.purdue.edu. hostmaster.purdue.edu. 581034717 1200 180 1209600 3600
purdue.edu.          86400   IN     TXT    "adobe-idp-site-verification=82aaf4b4-e1ab-4a7d-bea2-10f5820faec9"
purdue.edu.          86400   IN     TXT    "MS=6F64A1C5D9A24164AA0A7E59CC8C1E7758DE8E88"
purdue.edu.          86400   IN     TXT    "_globalsign-domain-verification=LBGN_PGPzs-K_ugvlmQc5pBRof3QCVximyhc9hNKsU"
purdue.edu.          86400   IN     TXT
"oysmqwNak8YGUW3sLWwJg1feIr+RzjS3GAjdI6bSHjLGftS6HmMqoAUagGXLrjogVk4ptcEPz+5oyrWQeCkXKw=="
purdue.edu.          86400   IN     TXT    "v=spf1 a:smtp.purdue.edu a:pbm.itap.purdue.edu ip4:128.210.210.98 ip4:128.210.210.99
include:spf.protection.outlook.com include:ne16.com include:_spf.qualtrics.com include:relay01.evenue.net ?all"
purdue.edu.          86400   IN     MX     10 xppmailspam12.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam04.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam01.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam13.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam05.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam11.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam06.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam02.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam03.itap.purdue.edu.
purdue.edu.          86400   IN     MX     10 xppmailspam09.itap.purdue.edu.
Purdu
;; ADDITIONAL SECTION:
ns.purdue.edu.         86400   IN     A      128.210.11.5
ns2.purdue.edu.        86400   IN     A      128.210.11.57
harbor.ecn.purdue.edu.  86400   IN     A      128.46.154.76
pendragon.cs.purdue.edu. 86400  IN     A      128.10.2.5
ns.purdue.edu.         86400   IN     AAAA   2001:18e8:800:202::a
ns2.purdue.edu.        86400   IN     AAAA   2001:18e8:800:202::b
e.edu.           86400   IN     MX     10 xppmailspam08.itap.purdue.edu.
```

PURDUE UNIVERSITY

Reflection and Amplification

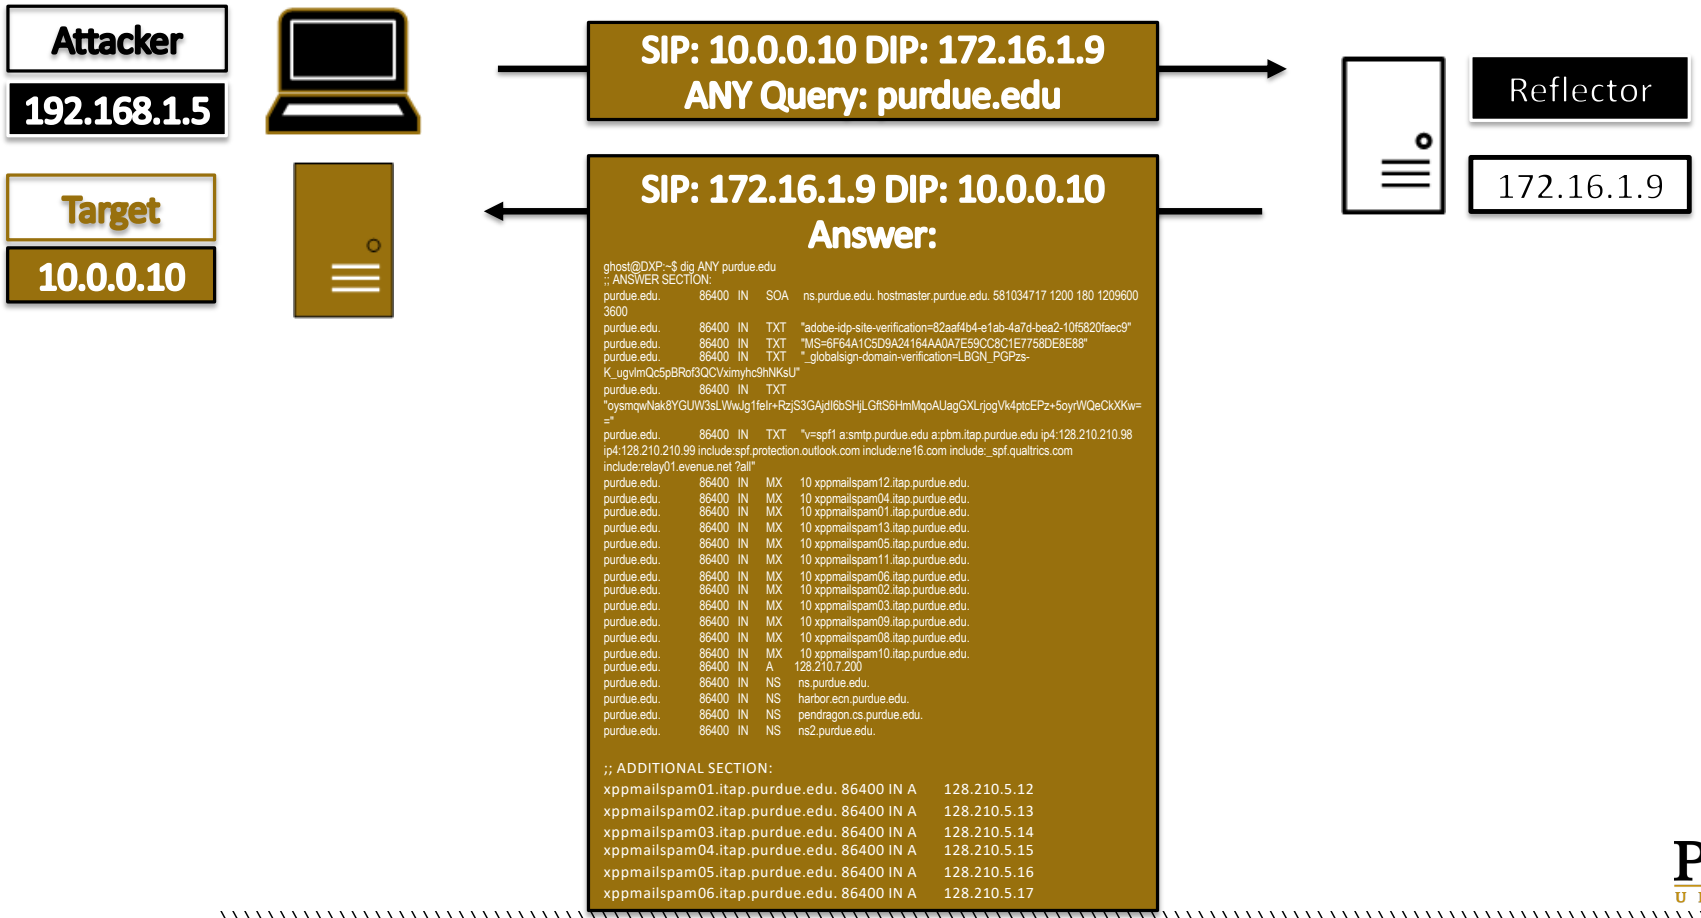

What is DNS reflection attack?

- What happens if an attacker forges the victim address as its source?

  …the reflected traffic goes to the target server.

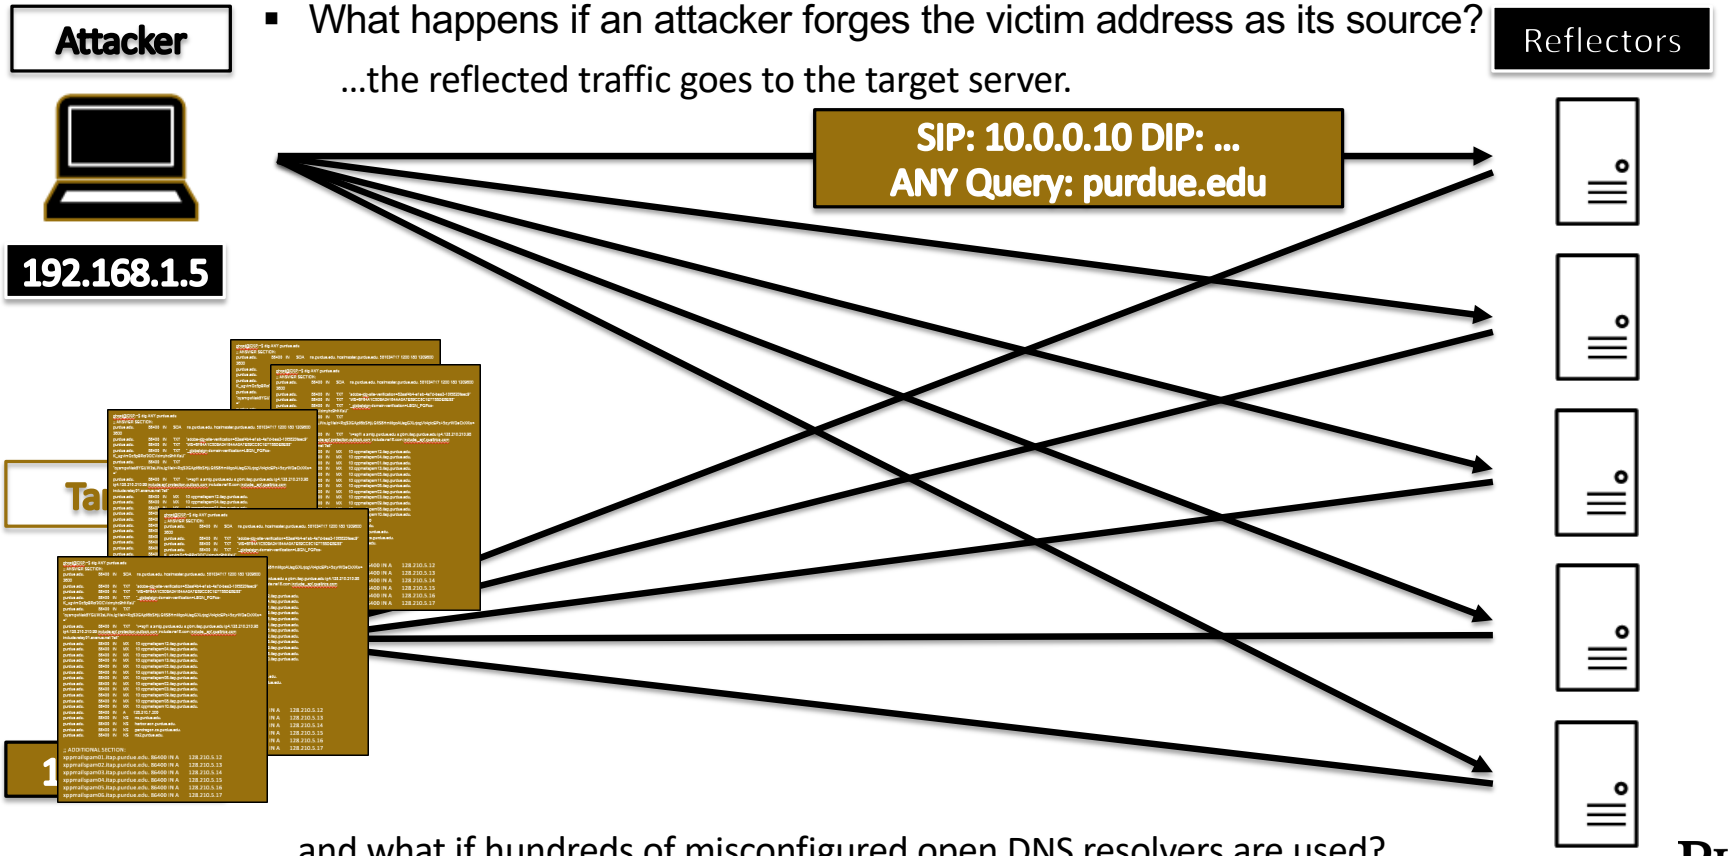

… and what if hundreds of misconfigured open DNS resolvers are used?

PURDUE UNIVERSITY

NTP REFLECTION

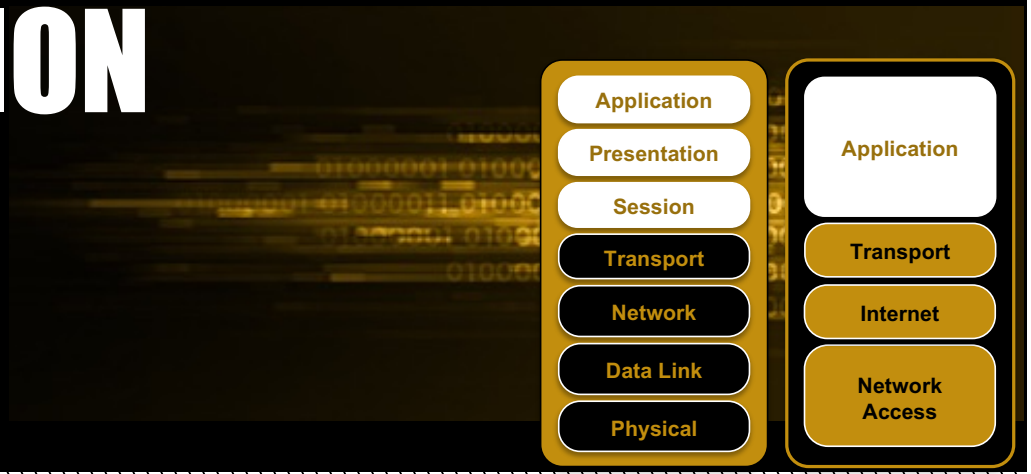

PURDUE UNIVERSITY

What is NTP

- NTP = Network Time Protocol
- RFC 778 (historical), 1981, Internet Clock Service
- Designed to allow clock synchronization
- ICMP, Echo
- UDP based
- Stratum Servers
  - 0 – high precision – atomic or radio clock, GPS synchronized;
  - 1 – within few milliseconds from their Stratum 0 source;
  - 2 – connected to multiple Stratum 1 over the network; may peer with other Stratum 2 servers;
  - 3 – computers sourcing time from Stratum 2 systems.

NTP Operation

- Lower stratum servers talk to higher stratum ones.
- The server maintains a list of all clients that have talked to them.
- A system administrator (or any other user) can query that list.

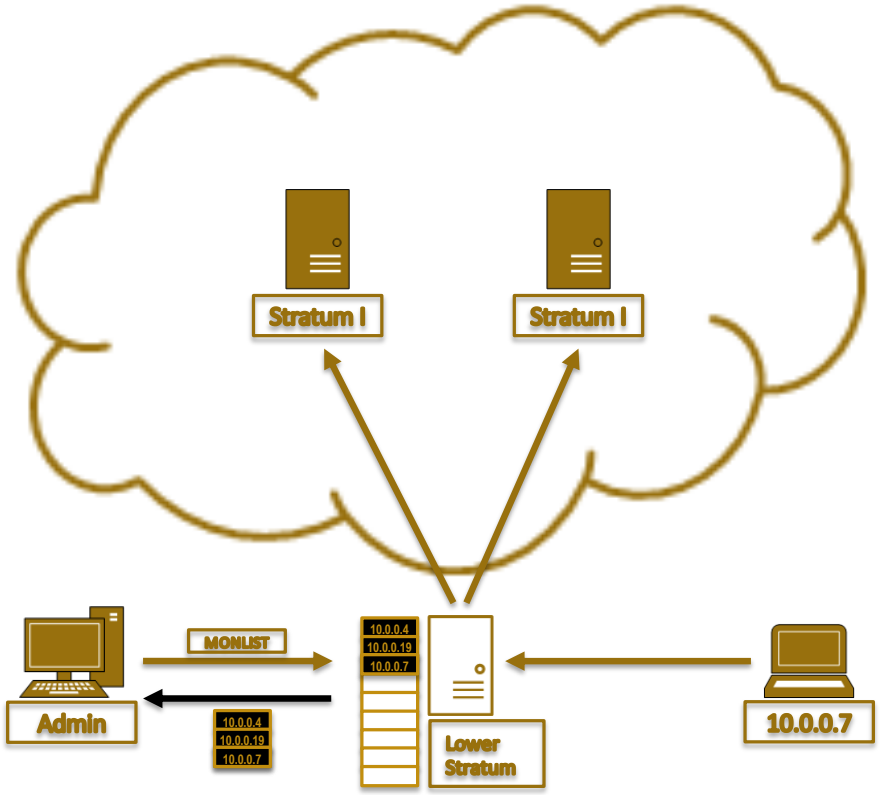

NTP Reflection

- What if an attacker wants to abuse this?
- The attacker sends a MONLIST request spoofing the source IP with the one of the victim.
- The NTP server, as designed, sends the list of IPs to the requestor.
- Now let's add a number of other misconfigured servers on the Internet.

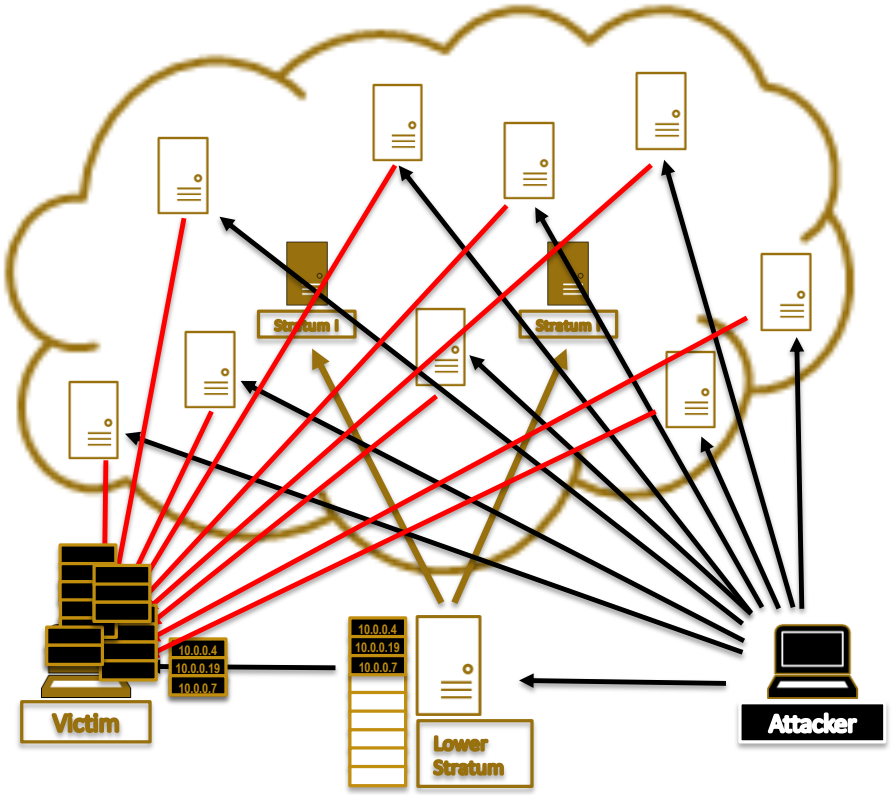

QUESTIONS?

PURDUE UNIVERSITY

MITIGATION STRATEGIES

PURDUE UNIVERSITY

The cost of a minute?

- How much does a minute of outage cost to your business?
- Are there other costs associated with it? Reputation?
- Are you in a risk category?
- How much is executive management willing to spend to stay up?
- Are there reasons you need to mitigate on-site vs offsite? Latency?

The Risk Pyramid

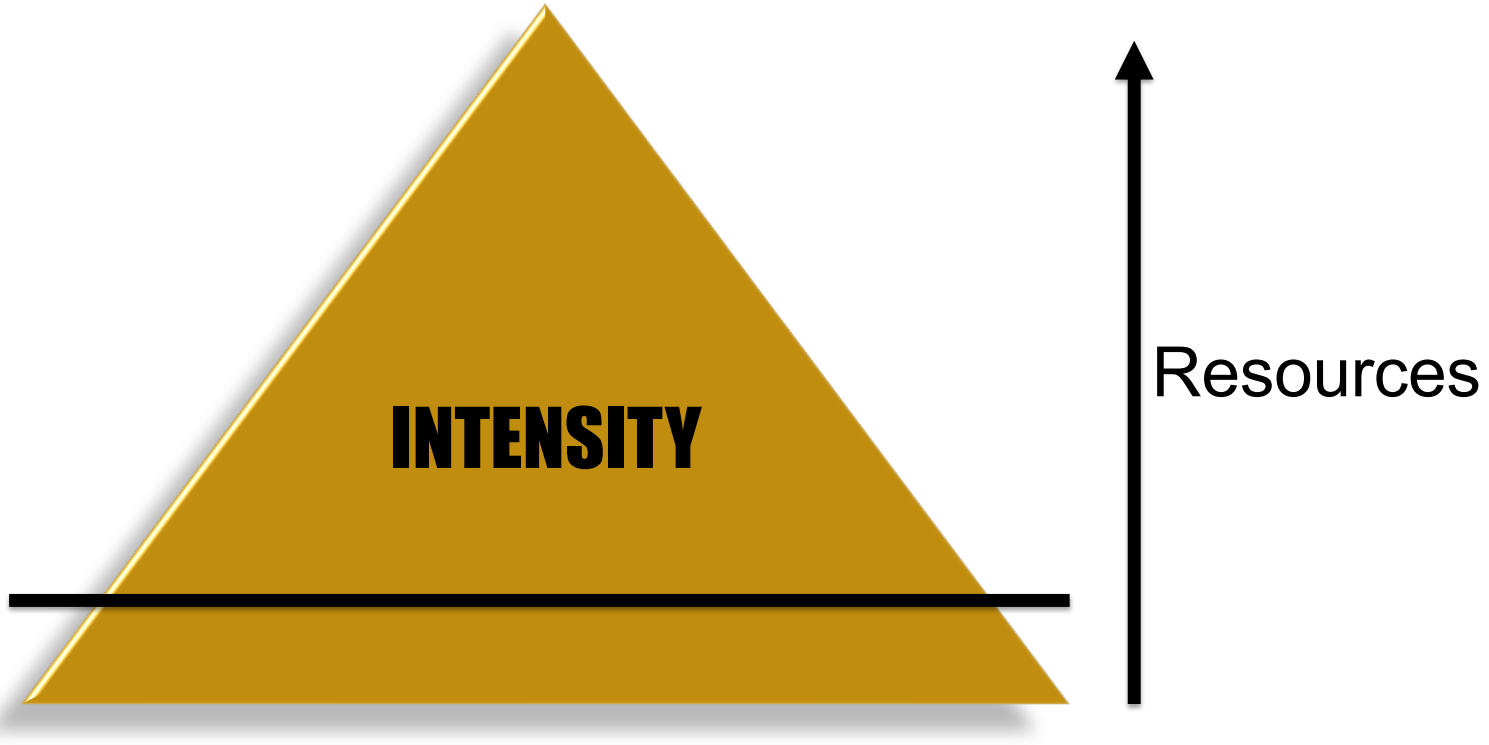

Mitigation

Different approaches:

- Do it yourself (DIY)
- Outsource/service
  - On demand
  - Always on
- Hybrid

Do it Yourself (On Premise)

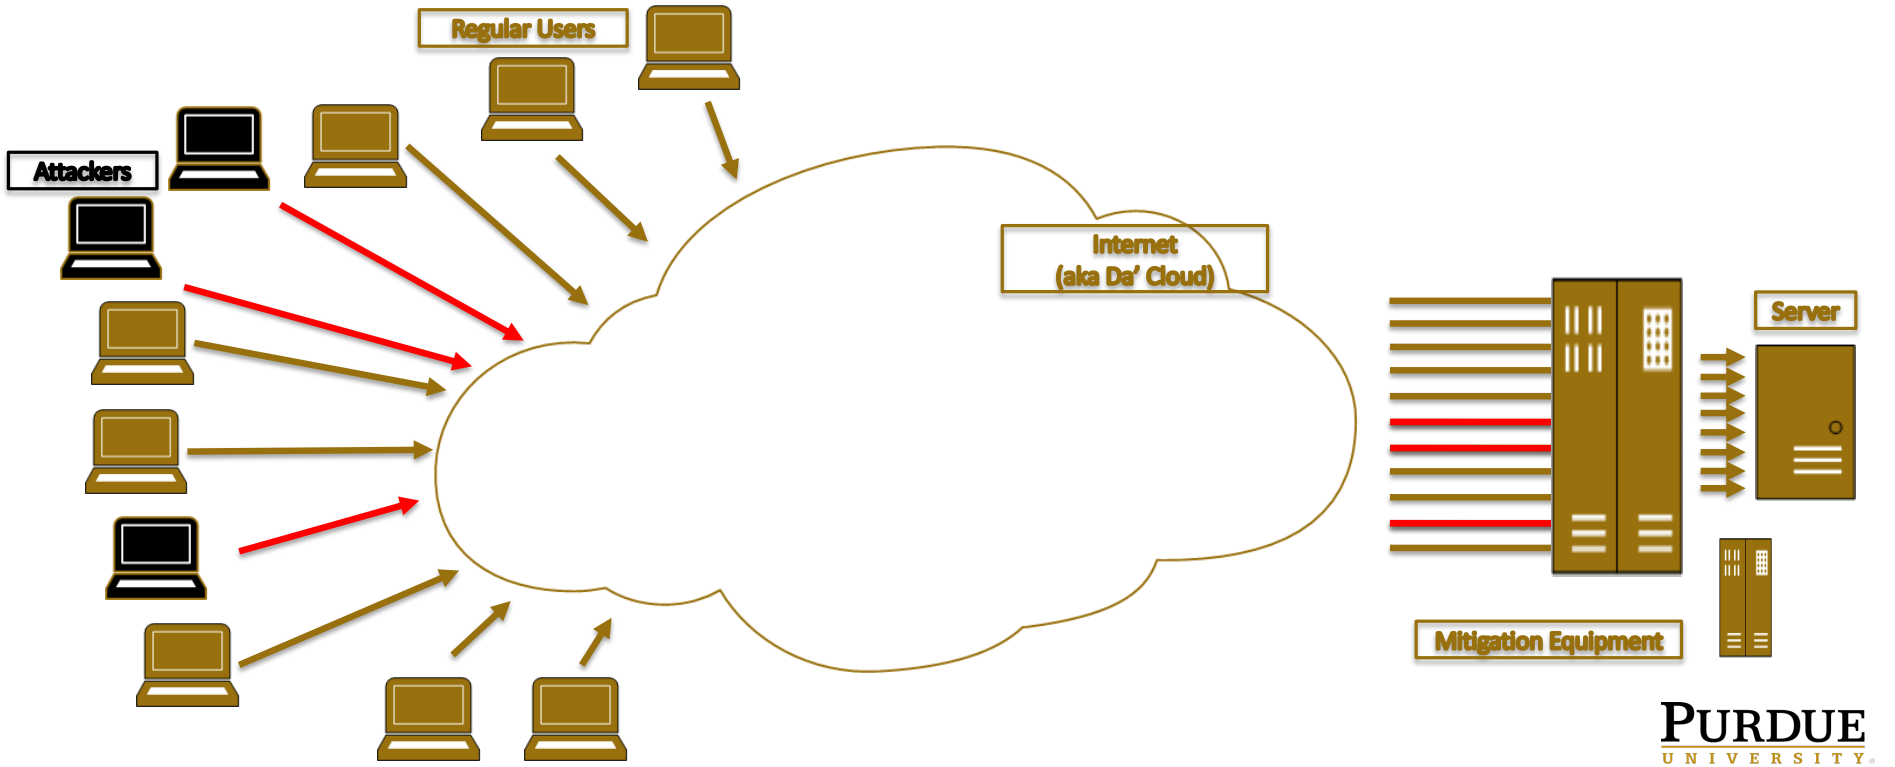

DIY: Considerationss

- Network capacity: bandwidth
- Hardware capacity: packet rates, inspecting headers and content?
- One time cost (refresh every 3-4 years)
- Depending on attacks size can be in $100,000s

DIY: Benefits

- Very low latency
- Can be application specific (non-http, gaming industry)
- Better control of the mitigation
- If inspecting TLS traffic keeps the keys in the company

DIY: Drawbacks

- Network capacity:
  - Fluctuates
  - How much do you over provision? Double, triple, ten times?
- Need to procure
  - bandwidth - monthly recurring - expensive, adds up
  - compute and network hardware
  - qualified personnel – hard to find; expensive; hard to retain

DIY: Conclusions

- At present DDoS attacks are at a very large scale but DIY is not easy to scale for small and medium networks
- Leverages economy of scale – requires a large infrastructure
- Infrastructure is very expensive to build and maintain
- Requires significant amount of know-how
- Unless hosting a very large site it's better left to the professionals

External service

- DDoS mitigation service providers and CDNs
- Pricing:
  - based on size of attack
  - based on clean traffic
- Operating model:
  - on demand
  - always on

On Demand DDoS Protection

- Target: detect and signal the mitigation provider
- Mitigation provider: Inject BGP routes
- Traffic is redirected to the mitigation provider
- Clean traffic is delivered to the origin server (usually over GRE tunnel)

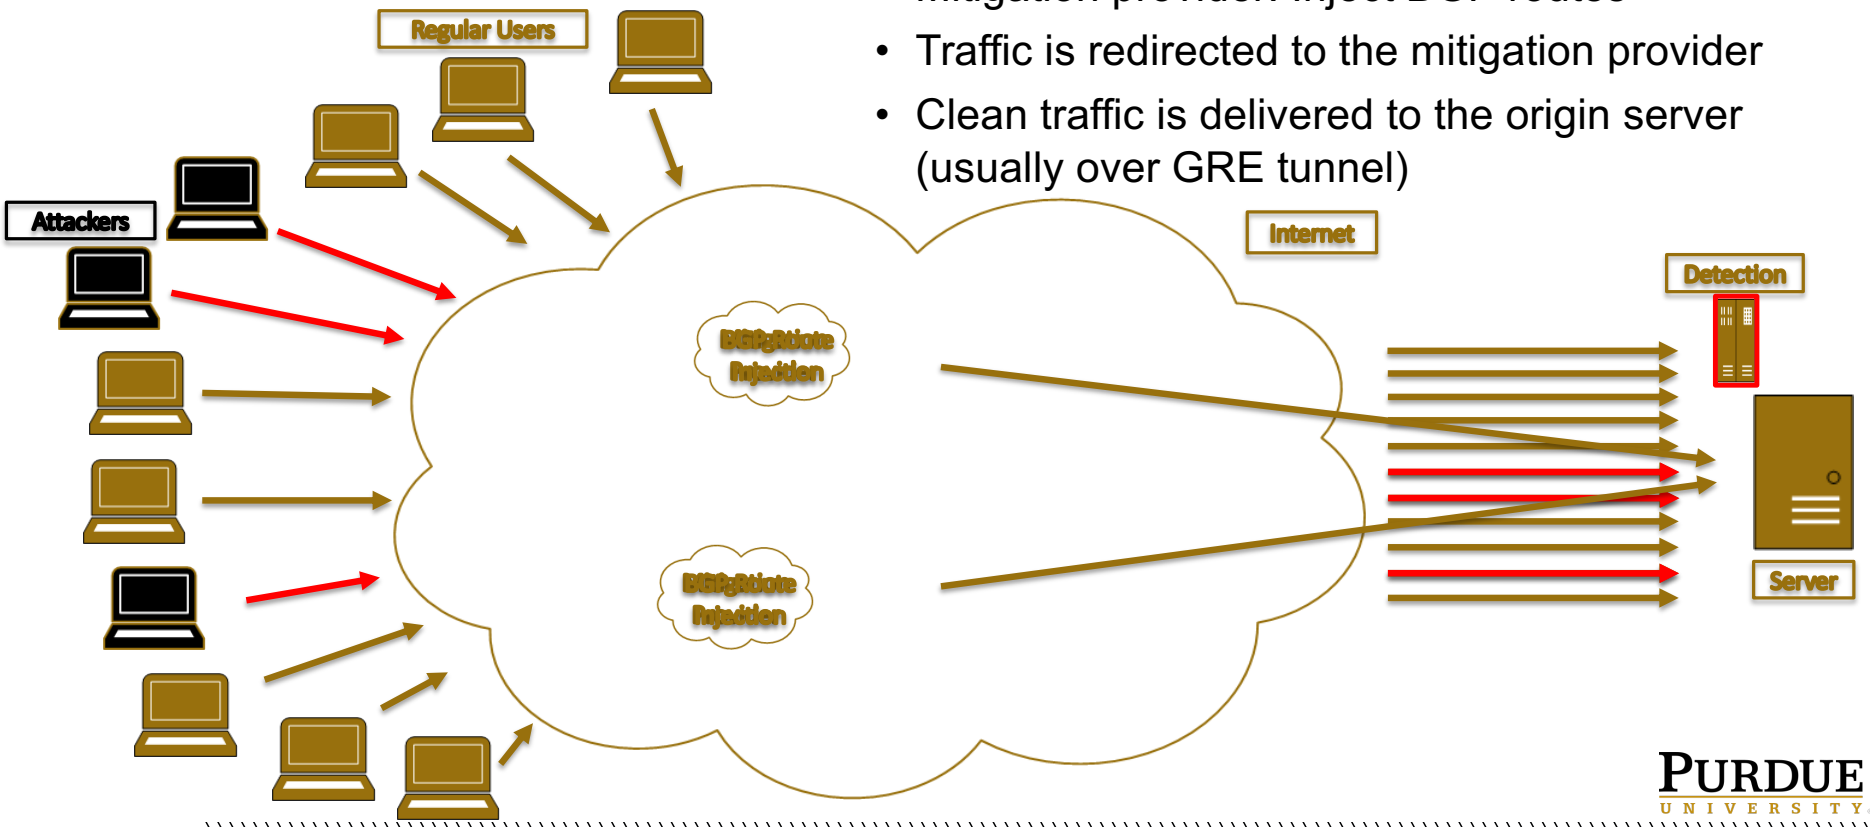

On Demand Mitigation - benefits

- Scales up very easily
- Since most applications are HTTP/S based, it is compatible with them
- Easier to deploy
- May leave the target vulnerable to bypass

On Demand Mitigation - drawbacks

- Takes time between the site being attacked until it switches to the service provider
- Potential outages
- Difficult to establish TLS
- May have increased latency
- **Target may still be exposed**
- Detection is not Application Aware
- GRE Tunnels create complexity

Always On Mitigation

- Permanently serve the customer space
  - Advertise IP address space
  - Use shared delivery infrastructure (CDN)
- Traffic is always flowing through the mitigation systems
- Usually combined with services like CDN, which further increases website performance (even during peace time)

Always On DDoS Mitigation (advertise IP space)

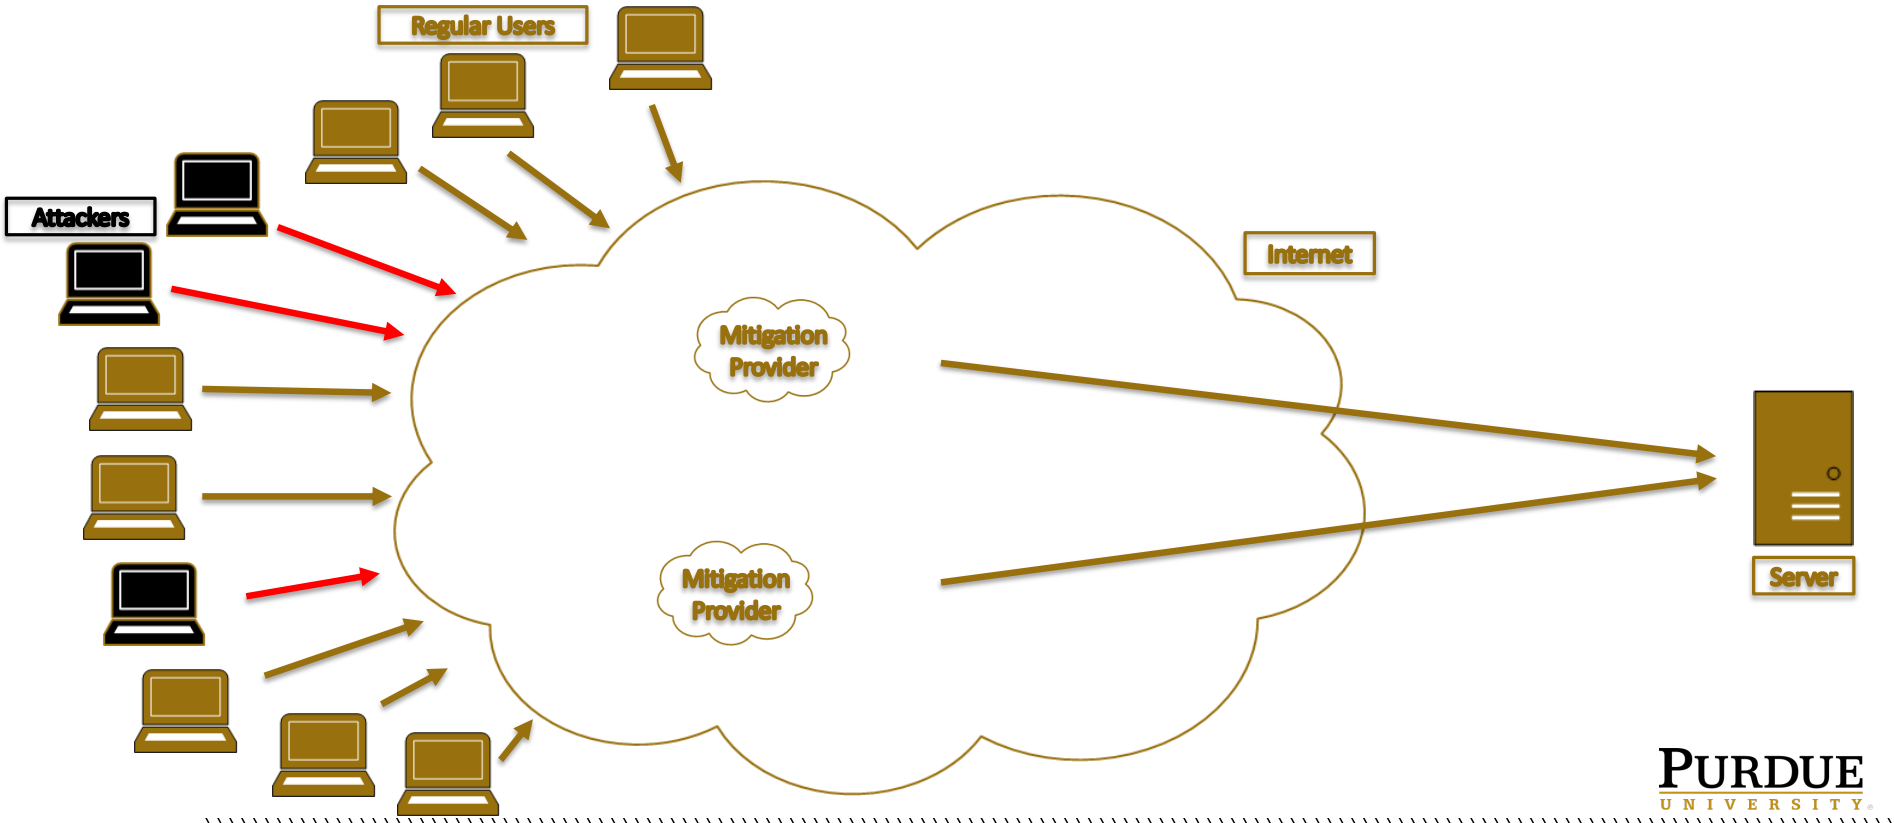

Always On DDoS Mitigation (CDN)

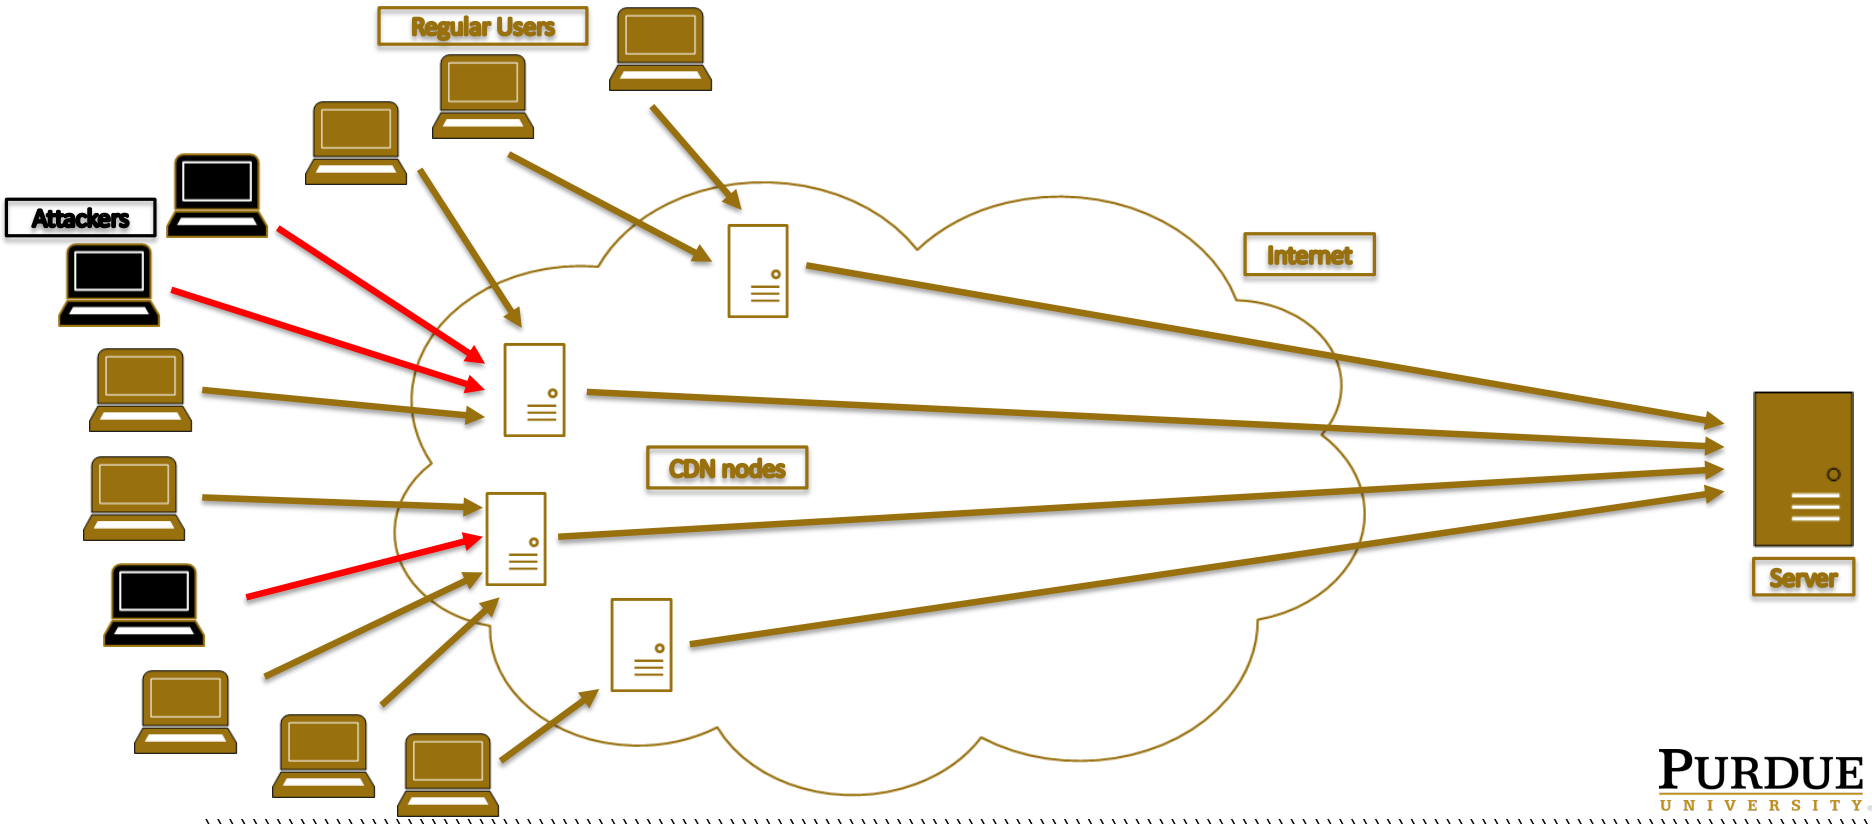

Always On Mitigation - benefits

- Scales up very well during volumetric attacks
- Mitigation can be virtually instantaneous
  - No moving parts during the attack
- Can protect most applications
- Once it's on there are no moving parts
- Very hard to bypass
- (proxy/caching) If deployed properly, it may improve website performance
- Cost depends on the website traffic (not the attack)

Always On Mitigation - drawbacks

- Can increase latency
- Challenges around TLS
- Stale caches
- May be much more expensive

Hybrid

- Combination of DIY and service providers
- Helps customers manage their risk profile in a more flexible way

Benefits:

- Provides protection against large scale events without the added service cost
- Allows for escalating response postures and risk/finance management
- Overall most of the benefits of On Demand

Drawbacks:

- Increased complexity
- Requires skilled personnel
- May have interoperability issues

DDoS mitigation service providers

Pros

- Hides the complexities of managing the problem
- May accelerate content delivery
- May be much cheaper, especially as attack sizes grow but are not common
- Cost: much, much lower than DIY

Cons

- May not be applicable to all applications, such as gaming
- May increase latency
- May end up expensive
- Third party sees the users (and maybe the content) - privacy, security
- Issues with stale cache

QUESTIONS?

PURDUE UNIVERSITY

THANK YOU!

PURDUE UNIVERSITY