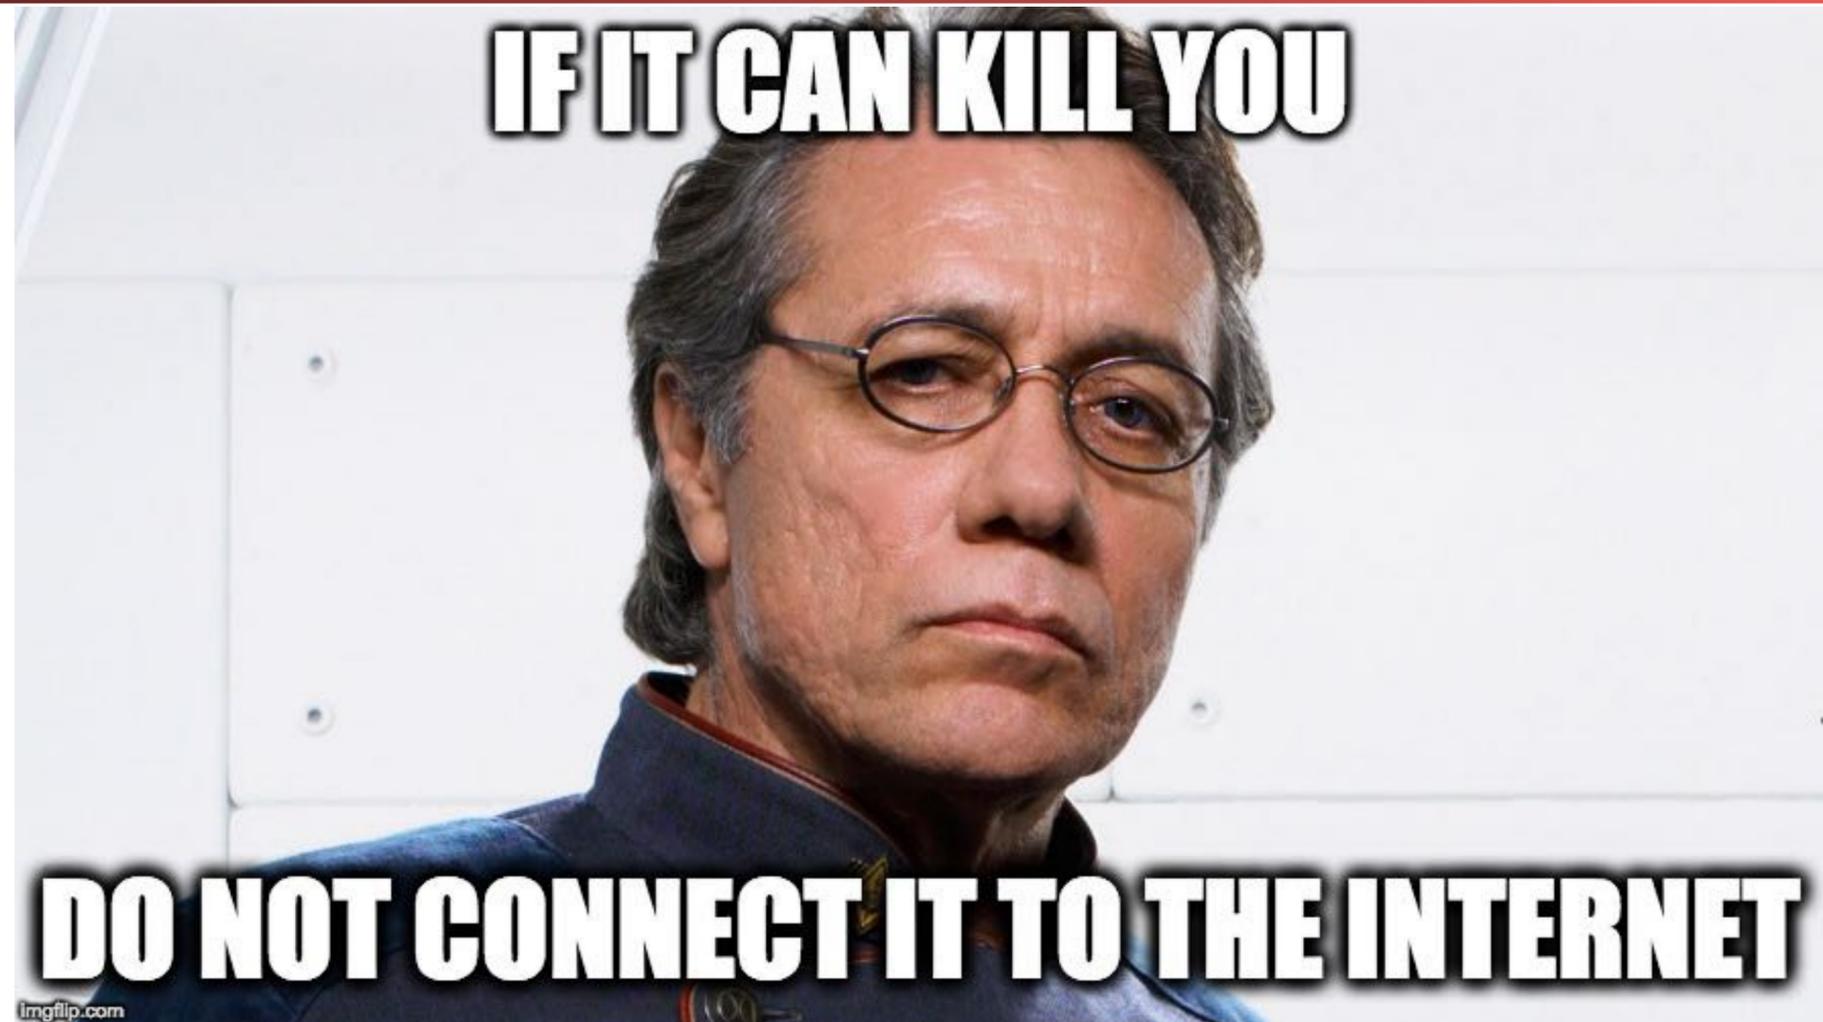


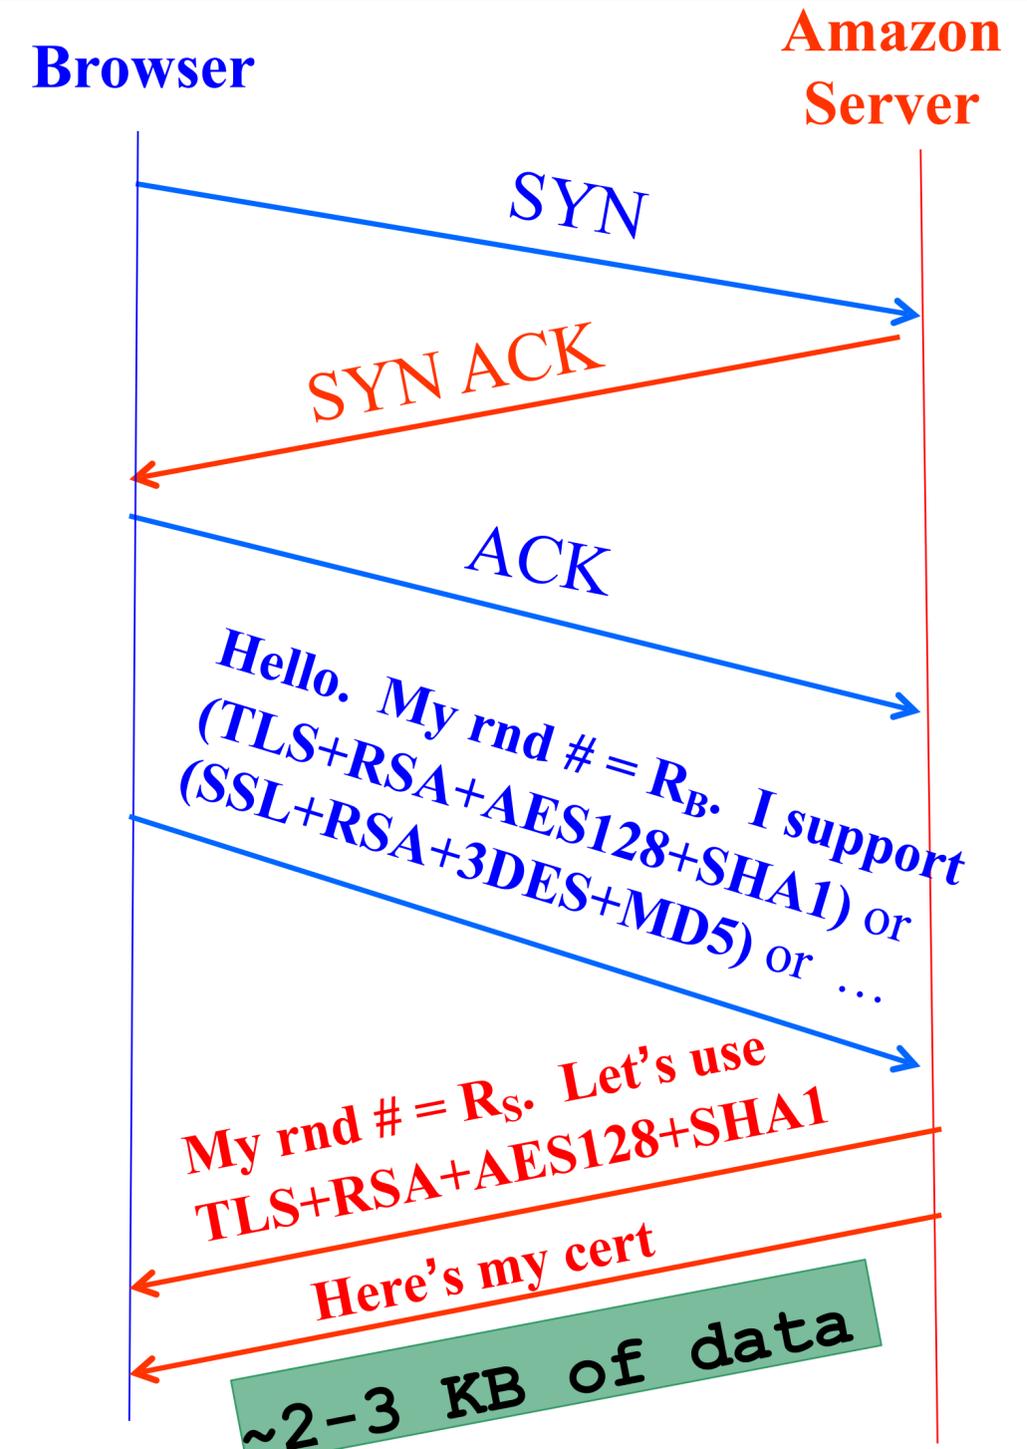
Network Security

5



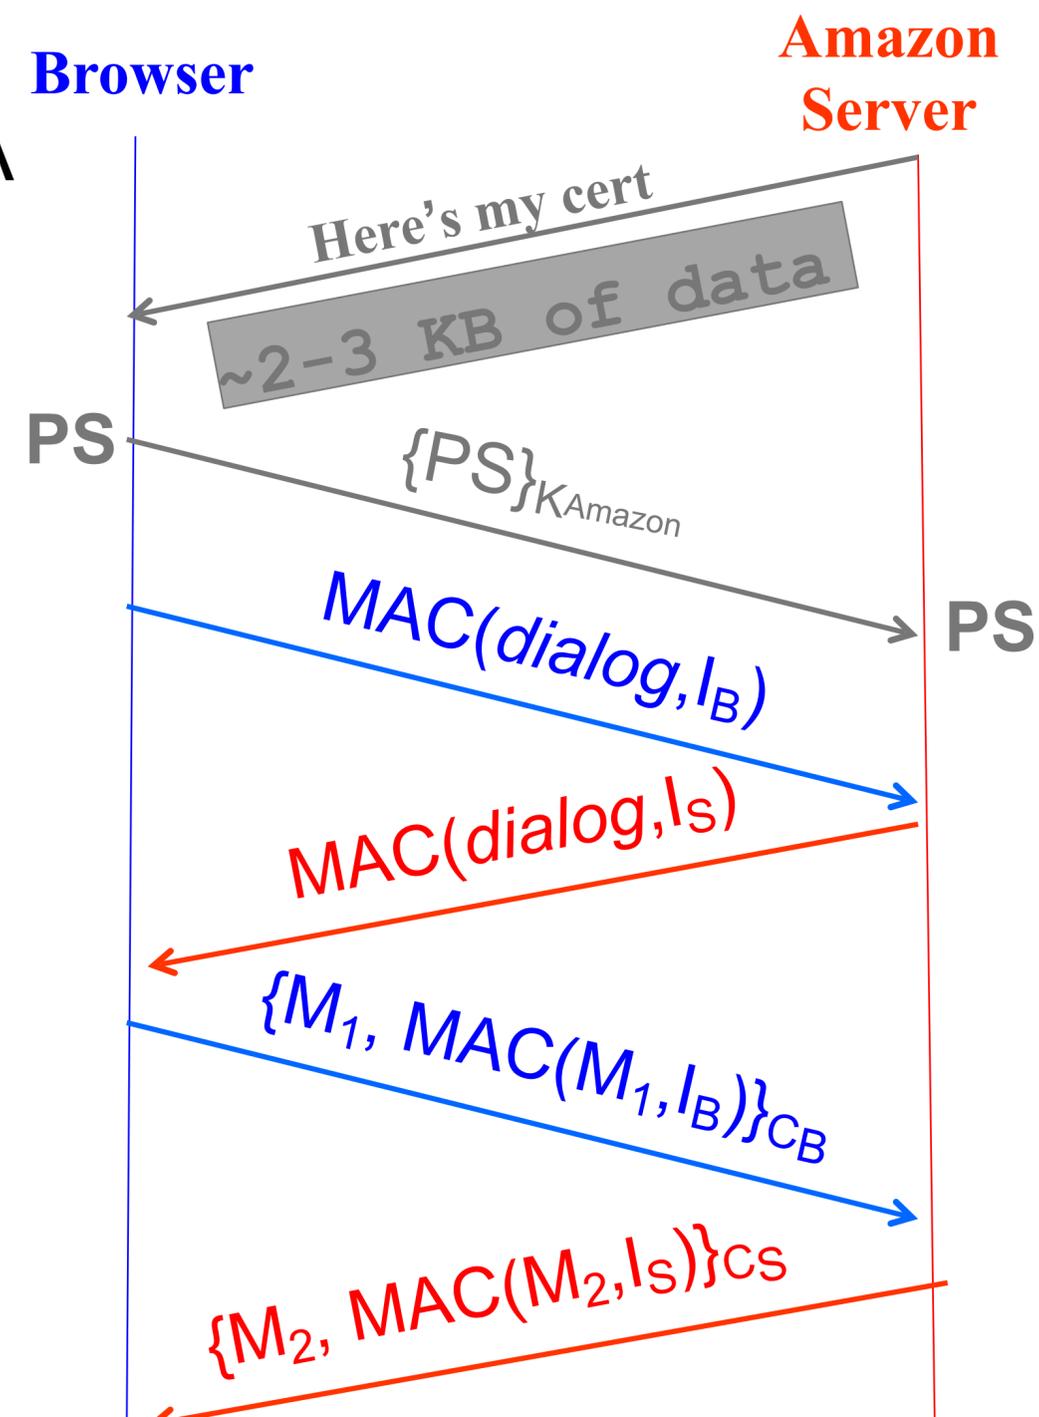
Reminder: HTTPS Connection (SSL / TLS)

- Browser (client) connects via TCP to Amazon's HTTPS server
- Client picks 256-bit random number R_B , sends over list of crypto protocols it supports
- Server picks 256-bit random number R_S , selects protocols to use for this session
- Server sends over its certificate
 - (all of this is in the clear)
- Client now *validates* cert



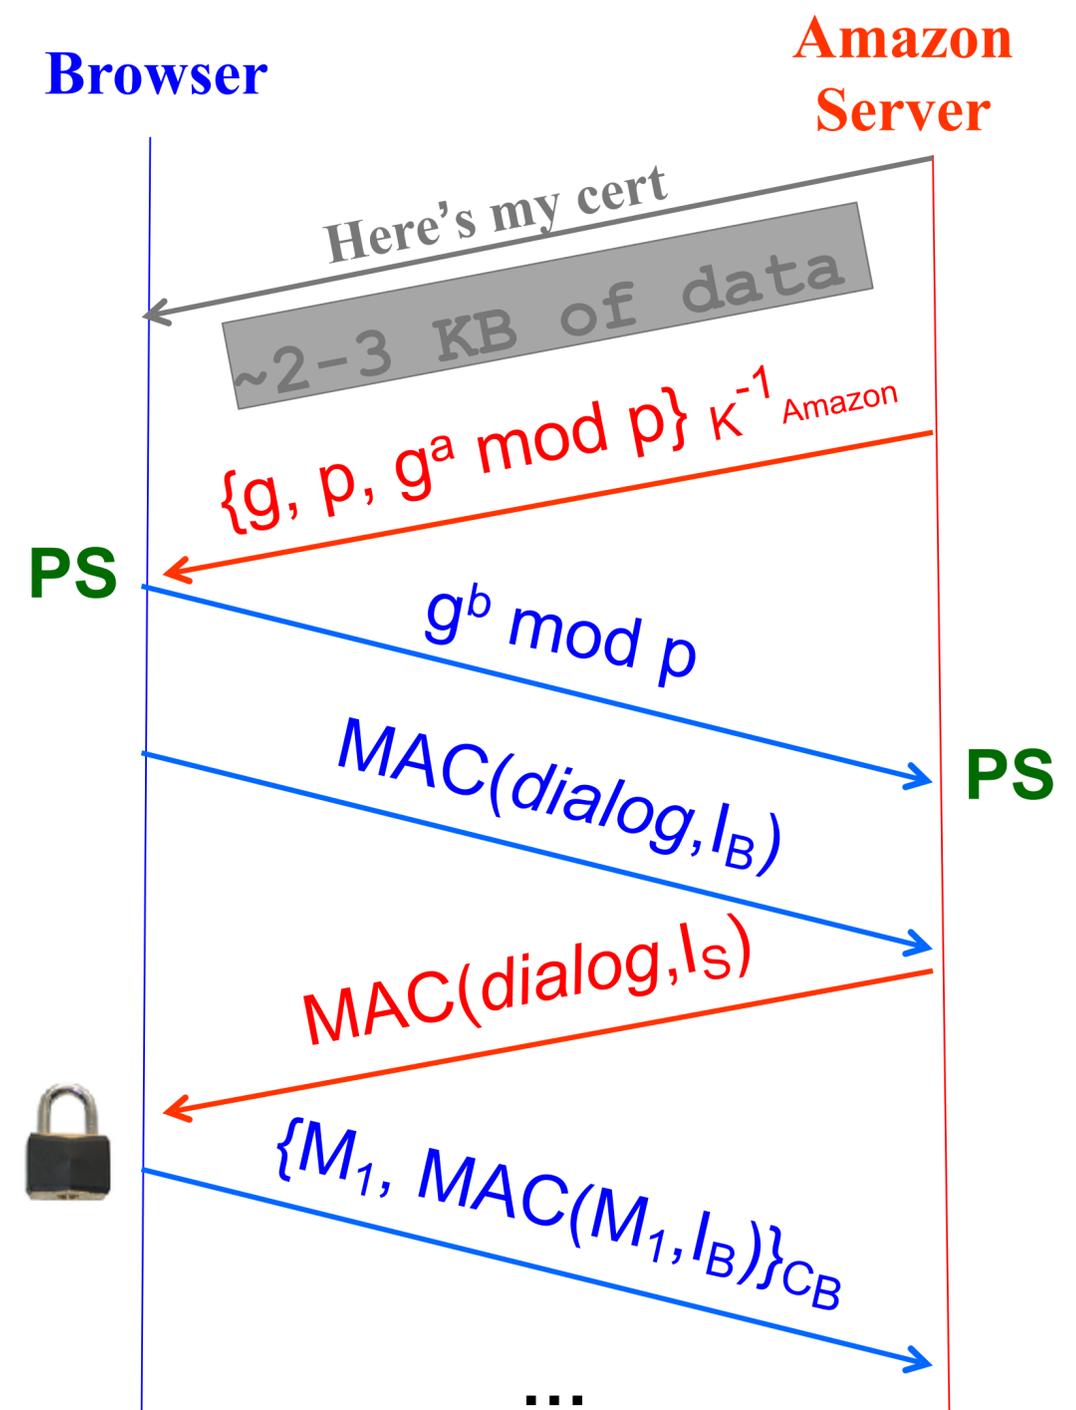
HTTPS Connection (SSL / TLS), cont.

- For RSA, browser constructs “Premaster Secret” PS
- Browser sends PS encrypted using Amazon’s public RSA key K_{Amazon}
- Using PS, R_B , and R_S , browser & server derive symm. cipher keys (C_B, C_S) & MAC integrity keys (I_B, I_S)
 - One pair to use in each direction
- Browser & server exchange MACs computed over entire dialog so far
- If good MAC, Browser displays 
- All subsequent communication encrypted w/ symmetric cipher (e.g., AES128) cipher keys, MACs
 - Sequence #'s thwart replay attacks



Alternative: Ephemeral Key Exchange via Diffie-Hellman

- For Diffie-Hellman (DHE), server generates random a , sends public parameters and $g^a \bmod p$
 - Signed with server's private key
- Browser verifies signature
- Browser generates random b , computes $PS = g^{ab} \bmod p$, sends $g^b \bmod p$ to server
- Server also computes $PS = g^{ab} \bmod p$
- Remainder is as before: from PS , R_B , and R_S , browser & server derive symm. cipher keys (C_B , C_S) and MAC integrity keys (I_B , I_S), etc...



Big Changes for TLS 1.3

Diffie/Hellman and ECDHE only

- The RSA key exchange has a substantial vulnerability
 - If the attacker is ever able to compromise the server and obtain its RSA key... the attacker can decrypt any traffic captured
 - RSA lacks *forward secrecy*
- So TLS 1.3 uses DHE/ECDHE only
- TLS 1.3 also speeds things up:
 - In the client hello, the client includes $\{g^b \bmod p\}$ for preferred parameters
 - If the server finds it suitable, the server returns $\{g^a \bmod p\}$
 - Saves a round-trip time

Why R_b and R_s ?

- Both R_b and R_s act to affect the keys... Why?
 - $Keys = F(R_b || R_s || PS)$
- Needed to prevent a ***replay attack***
 - Attacker captures the handshake from either the client or server and replays it...
- If the other side chooses a different R the next time...
 - The replay attack fails.
- But you don't need to check for reuse by the other side..
 - Just make sure you don't reuse it on your side!

And Sabotaged pRNGs...

- Let us assume the server is using DHE...
 - If an attacker can know a , they have all the information needed to decrypt the traffic:
 - Since $PS = g^{ab}$, and can see g^b .
- TLS spews a lot of "random" numbers publicly as well
 - Nonces in the crypto...
- If the server uses a bad pRNG which is both sabotaged and doesn't have ***rollback resistance***...
 - Dual_EC DRBG where you know the secret used to create the generator...
 - ANSI X9.31: An AES based one with a secret key...
- Attacker sees the handshake, sees subsequent PRNG calls, works ***backwards*** to get the secret
 - Attack of the week: DUHK
 - <https://blog.cryptographyengineering.com/2017/10/23/attack-of-the-week-duhk/>

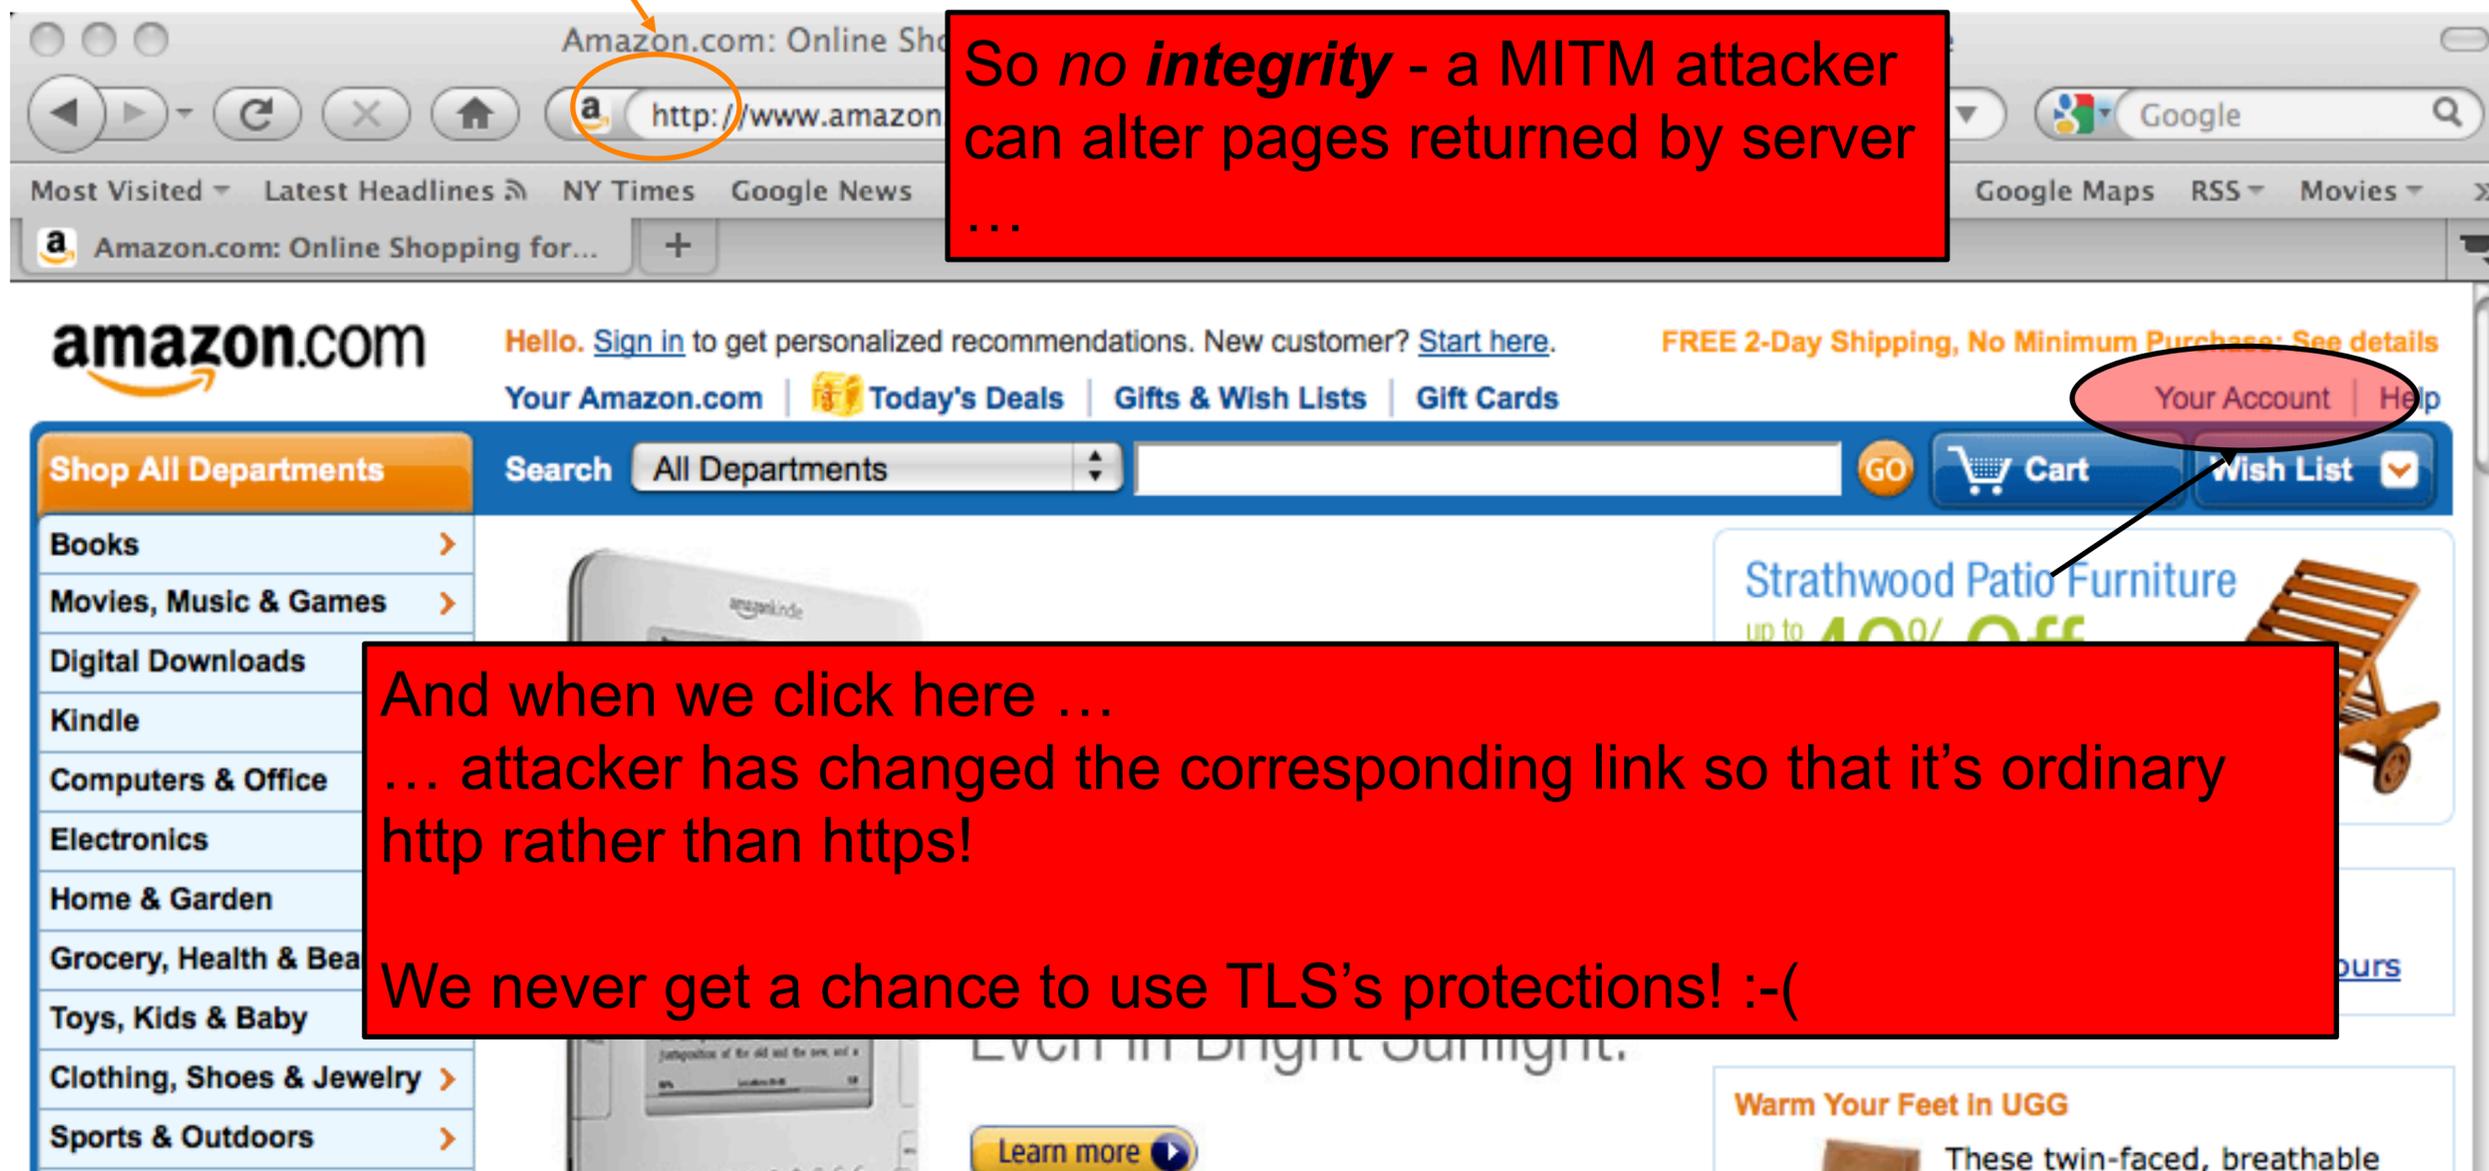
SSL/TLS Problem: Revocation

- A site screws up and an attacker steals the private key associated with a certificate, what now?
 - Certificates have a timestamp and are only good for a specified time
 - But this time is measured in years!?!?
- Two mitigations:
 - Certificate revocation lists
 - Your browser occasionally calls back to get a list of "no longer accepted" certificates
 - OSCP
 - Online Certificate Status Protocol:
https://en.wikipedia.org/wiki/Online_Certificate_Status_Protocol

“sslstrip”

(Amazon FINALLY fixed this recently)

Regular web surfing: http: URL



So no *integrity* - a MITM attacker can alter pages returned by server ...

And when we click here ...
... attacker has changed the corresponding link so that it's ordinary http rather than https!
We never get a chance to use TLS's protections! :-('

SSL / TLS Limitations, cont.

- Problems that SSL / TLS does not take care of ?
- Censorship
- SQL injection / XSS / server-side coding/logic flaws
- Vulnerabilities introduced by server inconsistencies
- Browser and server bugs
- Bad passwords
- What about the trust?

TLS/SSL Trust Issues

- User has to make correct trust decisions ...



Recycle Bin

Welcome to eBay - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address <http://0xbd5947e3/sendfiles/.../signin.ebay.com/ws/eBayISAPI.dll?SignIn&httpAFFwww.ebay.com2F/> Go Links >>eBay Buyer Protection [Learn more](#) **NEW**

Welcome to eBay

Ready to bid and buy? Register here

Join the millions of people who are already a part of the eBay family. Don't worry, we have room for one more.

Register as an eBay Member and enjoy privileges including:

- Bid, buy and find bargains from all over the world
- Shop with confidence with PayPal Buyer Protection
- Connect with the eBay community and more!

[Register](#)

Sign in to your account

Back for more fun? Sign in now to buy, bid and sell, or to manage your account.

User ID [I forgot my user ID](#)Password [I forgot my password](#)

Keep me signed in for today. Don't check this box if you're at a public or shared computer.

[Sign in](#)

Having problems with signing in? [Get help.](#)

Protect your account: Create a unique password by using a combination of letters and numbers that are not

The screenshot shows a Windows XP desktop environment. At the top, a Recycle Bin icon is visible. The main window is Microsoft Internet Explorer, displaying the eBay sign-in page. The address bar shows a URL starting with 'http://0xbd5947e3/sendfiles/.../signin.ebay.com/'. The page content includes the eBay logo, 'eBay Buyer Protection' banner, and a sign-in form with fields for 'User ID' and 'Password'. A security warning dialog box is open in the foreground, asking 'When you send information to the Internet, it might be possible for others to see that information. Do you want to continue?' with a checked option for 'In the future, do not show this message.' and 'Yes'/'No' buttons. The Windows taskbar at the bottom shows the Start button, active browser windows, and the system tray with the time 8:36 PM.

The screenshot shows a Windows XP desktop environment. At the top, a taskbar contains the Start button and several open applications: "eBay sent this messa...", "Identity Confirmation...", and "Internet". The main window is Microsoft Internet Explorer, displaying an eBay identity confirmation page. The page title is "Identity Confirmation - Microsoft Internet Explorer". The address bar shows the URL: `http://0xbd5947e3/sendfiles/.../signin.ebay.com/ws/eBayISAPI.dll?SignInruhttpAFFwww.ebay.com2F/sQuestion.php`. The page content includes the eBay logo, the heading "Please confirm your identity", and a section titled "Please answer security question". Below this, there are several input fields and labels: "Select your secret question...", "Answer the secret question you provided.", "What is your other eBay user ID or another's...", "What email used to be associated with this account?", and "Have you ever sold something on eBay?". A "Security Alert" dialog box is overlaid on the page. The dialog box has a blue title bar and contains the following text: "Information you exchange with this site cannot be viewed or changed by others. However, there is a problem with the site's security certificate." Below this, there are three bullet points: a warning icon followed by "The security certificate was issued by a company you have not chosen to trust. View the certificate to determine whether you want to trust the certifying authority.", a checkmark followed by "The security certificate date is valid.", and another checkmark followed by "The security certificate has a valid name matching the name of the page you are trying to view." At the bottom of the dialog box, it asks "Do you want to proceed?" and provides three buttons: "Yes", "No", and "View Certificate". The status bar at the bottom of the browser window shows "Done" and "Internet".

The screenshot shows a Microsoft Internet Explorer window titled "Identity Confirmation - Microsoft Internet Explorer". The address bar contains a URL from eBay. The main content area displays the eBay logo and a form for identity confirmation. The form includes a dropdown menu for "Select your secret question...", a text input field for the answer, and a question: "What is your other eBay user ID or another...". Below this is another text input field for an email address. At the bottom of the form, there are radio buttons for "Have you ever sold something on eBay?" with options "No" and "Yes".

Overlaid on the browser window is a "Certificate" dialog box. The dialog has three tabs: "General", "Details", and "Certification Path". The "General" tab is active and shows the following information:

- Certificate Information**
- This certificate is intended for the following purpose(s):**
 - Ensures the identity of a remote computer
- * Refer to the certification authority's statement for details.
- Issued to:** rover.ebay.com
- Issued by:** VeriSign Class 3 Secure Server CA - G3
- Valid from:** 10/22/2010 to 12/1/2012

At the bottom of the dialog box, there are two buttons: "Install Certificate..." and "Issuer Statement", and an "OK" button.

Identity Confirmation - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites

Address <http://0xbd5947e3/sendfiles/.../signin.ebay.com/ws/eBayISAPI.dll?SignInhttpAFFwww.ebay.com2F/sQuestion.php> Go Links



Please confirm your identity

Please answer security question

Select your secret question...

 Answer the secret question you provided.

What is your other eBay user ID or another...

What email used to be associated with this account?

Have you ever sold something on eBay?
 No
 Yes

Certificate

General Details Certification Path

Show: <All>

| Field | Value |
|---------------------|-----------------------------------|
| Version | V3 |
| Serial number | 4d ab c9 a6 0a 30 20 57 f9 23 ... |
| Signature algorithm | sha1RSA |
| Issuer | VeriSign Class 3 Secure Server... |
| Valid from | Friday, October 22, 2010 4:00... |
| Valid to | Saturday, December 01, 2012... |
| Subject | rover.ebay.com, Site Operatio... |
| Public key | RSA (1024 Bits) |

Edit Properties... Copy to File...

OK

Internet

Identity Confirmation - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites

Address <http://0xbd5947e3/sendfiles/.../signin.ebay.com/ws/eBayISAPI.dll?SignInruhttpAFFwww.ebay.com2F/sQuestion.php> Go Links



Please confirm your identity

Please answer security question

Select your secret question...

 Answer the secret question you provided.

What is your other eBay user ID or another...

What email used to be associated with this account?

Have you ever sold something on eBay?
 No
 Yes

Certificate

General Details Certification Path

Show: <All>

| Field | Value |
|------------------------------|-------------------------------------|
| Subject Alternative Name | DNS Name=rover.ebay.com, ... |
| Basic Constraints | Subject Type=End Entity, Pat... |
| Key Usage | Digital Signature, Key Encipher... |
| CRL Distribution Points | [1]CRL Distribution Point: Distr... |
| Certificate Policies | [1]Certificate Policy:Policy Ide... |
| Enhanced Key Usage | Server Authentication (1.3.6.... |
| Authority Key Identifier | KeyID=0d 44 5c 16 53 44 c1 8... |
| Authority Information Access | [1]Authority Info Access: Acc... |

Edit Properties... Copy to File...

OK

Internet

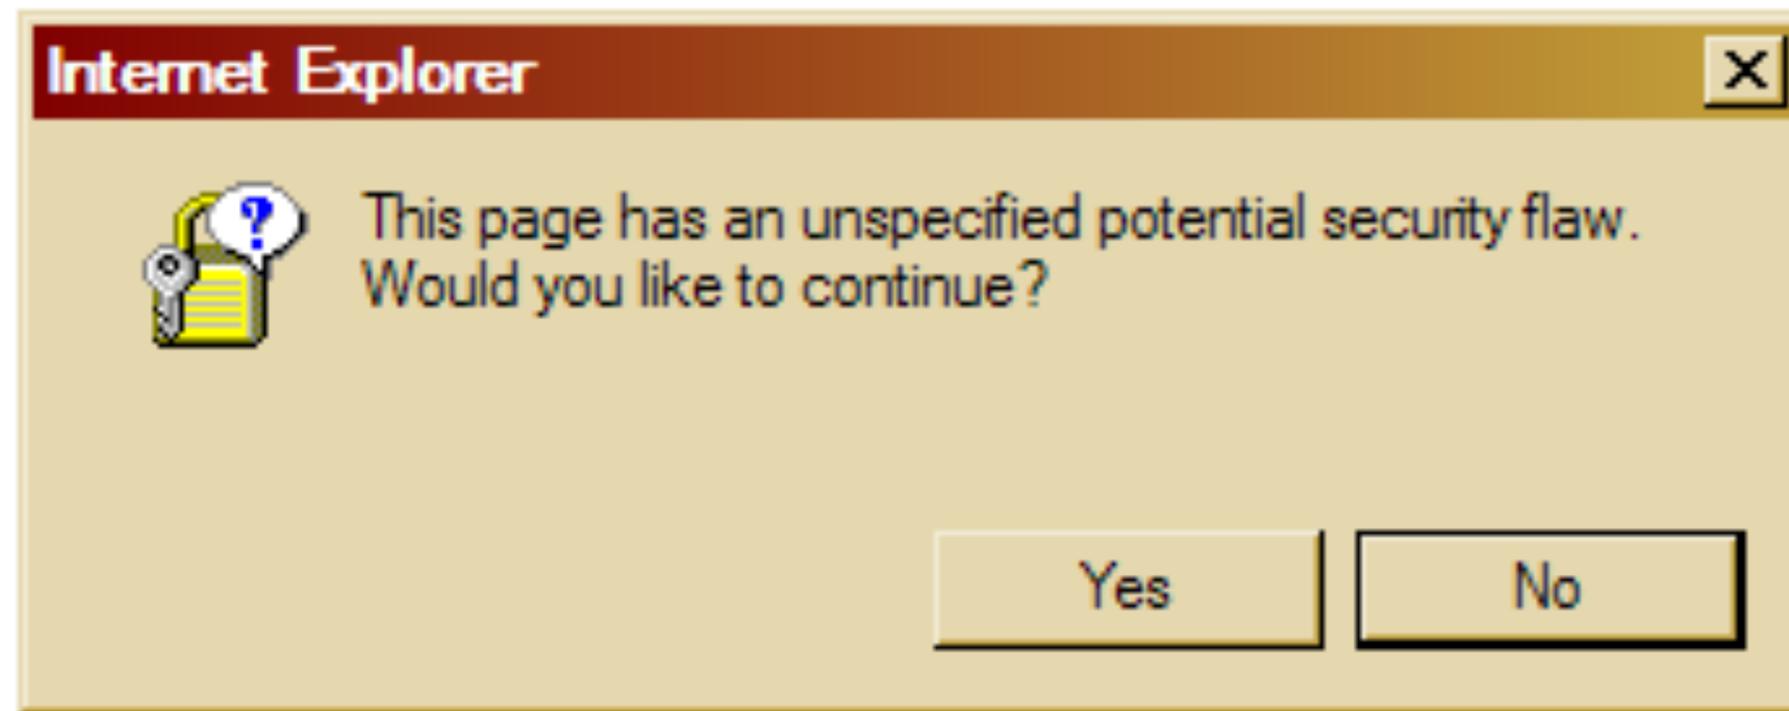
The screenshot shows a Microsoft Internet Explorer window titled "Identity Confirmation - Microsoft Internet Explorer". The address bar contains the URL: `http://0xbd5947e3/sendfiles/.../signin.ebay.com/ws/eBayISAPI.dll?SignInruhttpAFFwww.ebay.com2F/sQuestion.php`. The page content includes the eBay logo and a heading "Please confirm your identity". Below this, there is a section titled "Please answer security question" with a dropdown menu for "Select your secret question...", a text input field, and a label "Answer the secret question you provided.". Further down, there are two more text input fields with labels "What is your other eBay user ID or another..." and "What email used to be associated with this account?". At the bottom, there is a question "Have you ever sold something on eBay?" with radio buttons for "No" and "Yes".

Overlaid on the page is a "Certificate" dialog box with three tabs: "General", "Details", and "Certification Path". The "Certification Path" tab is active, showing a tree view of the certification path:

- VeriSign
 - VeriSign Class 3 Secure Server CA - G3
 - rover.ebay.com

Below the tree view is a "View Certificate" button. At the bottom of the dialog, the "Certificate status:" section displays "This certificate is OK.", which is circled in red. An "OK" button is located at the bottom right of the dialog.

The equivalent as seen by most Internet users:



(note: an actual Windows error message!)

TLS/SSL Trust Issues, cont.

- *“Commercial certificate authorities protect you from anyone from whom they are unwilling to take money.”*
- Matt Blaze, circa 2001
- So how many CAs do we have to worry about, anyway?

Keychain Access

Click to unlock the System Roots keychain.

Keychains

- login
- iCloud
- System
- System Roots**

AAA Certificate Services
Root certificate authority
Expires: Sunday, December 31, 2028 at 3:59:59 PM Pacific Standard Time
✔ This certificate is valid

| Name | Kind | Date Modified | Expires | Keychain |
|----------------------------------|-------------|---------------|---------------------------|--------------|
| AAA Certificate Services | certificate | -- | Dec 31, 2028, 3:59:59 PM | System Roots |
| Actalis Authentication Root CA | certificate | -- | Sep 22, 2030, 4:22:02 AM | System Roots |
| AddTrust Class 1 CA Root | certificate | -- | May 30, 2020, 3:38:31 AM | System Roots |
| AddTrust External CA Root | certificate | -- | May 30, 2020, 3:48:38 AM | System Roots |
| Admin-Root-CA | certificate | -- | Nov 9, 2021, 11:51:07 PM | System Roots |
| AffirmTrust Commercial | certificate | -- | Dec 31, 2030, 6:06:06 AM | System Roots |
| AffirmTrust Networking | certificate | -- | Dec 31, 2030, 6:08:24 AM | System Roots |
| AffirmTrust Premium | certificate | -- | Dec 31, 2040, 6:10:36 AM | System Roots |
| AffirmTrust Premium ECC | certificate | -- | Dec 31, 2040, 6:20:24 AM | System Roots |
| ANF Global Root CA | certificate | -- | Jun 5, 2033, 10:45:38 AM | System Roots |
| Apple Root CA | certificate | -- | Feb 9, 2035, 1:40:36 PM | System Roots |
| Apple Root CA - G2 | certificate | -- | Apr 30, 2039, 11:10:09 AM | System Roots |
| Apple Root CA - G3 | certificate | -- | Apr 30, 2039, 11:19:06 AM | System Roots |
| Apple Root Certificate Authority | certificate | -- | Feb 9, 2025, 4:18:14 PM | System Roots |
| ApplicationCA | certificate | -- | Dec 12, 2017, 7:00:00 AM | System Roots |
| ApplicationCA2 Root | certificate | -- | Mar 12, 2033, 7:00:00 AM | System Roots |
| Atos TrustedRoot 2011 | certificate | -- | Dec 31, 2030, 3:59:59 PM | System Roots |
| Autoridad de...nal CIF A62634068 | certificate | -- | Dec 31, 2030, 12:38:15 AM | System Roots |
| Autoridad de...Estado Venezolano | certificate | -- | Dec 17, 2030, 3:59:59 PM | System Roots |

Category

- All Items**
- Passwords
- Secure Notes
- My Certificates
- Keys
- Certificates

168 items

TLS/SSL Trust Issues

- *“Commercial certificate authorities protect you from anyone from whom they are unwilling to take money.”*
- Matt Blaze, circa 2001
- So how many CAs do we have to worry about, anyway?
- Of course, it's not just their greed that matters ...

News

Solo Iranian hacker takes credit for Comodo certificate attack

Security researchers split on whether 'ComodoHacker' is the real deal

By Gregg Keizer

March 27, 2011 08:39 PM ET

 [Comments \(5\)](#)

 [Recommended \(37\)](#)

 [Like](#)

 84

Computerworld - A solo Iranian hacker on Saturday claimed responsibility for stealing multiple SSL certificates belonging to some of the Web's biggest sites, including Google, Microsoft, Skype and Yahoo.

Early reaction from security experts was mixed, with some believing the hacker's claim, while others were dubious.

Fraudulent Google certificate points to Internet attack

Weaver

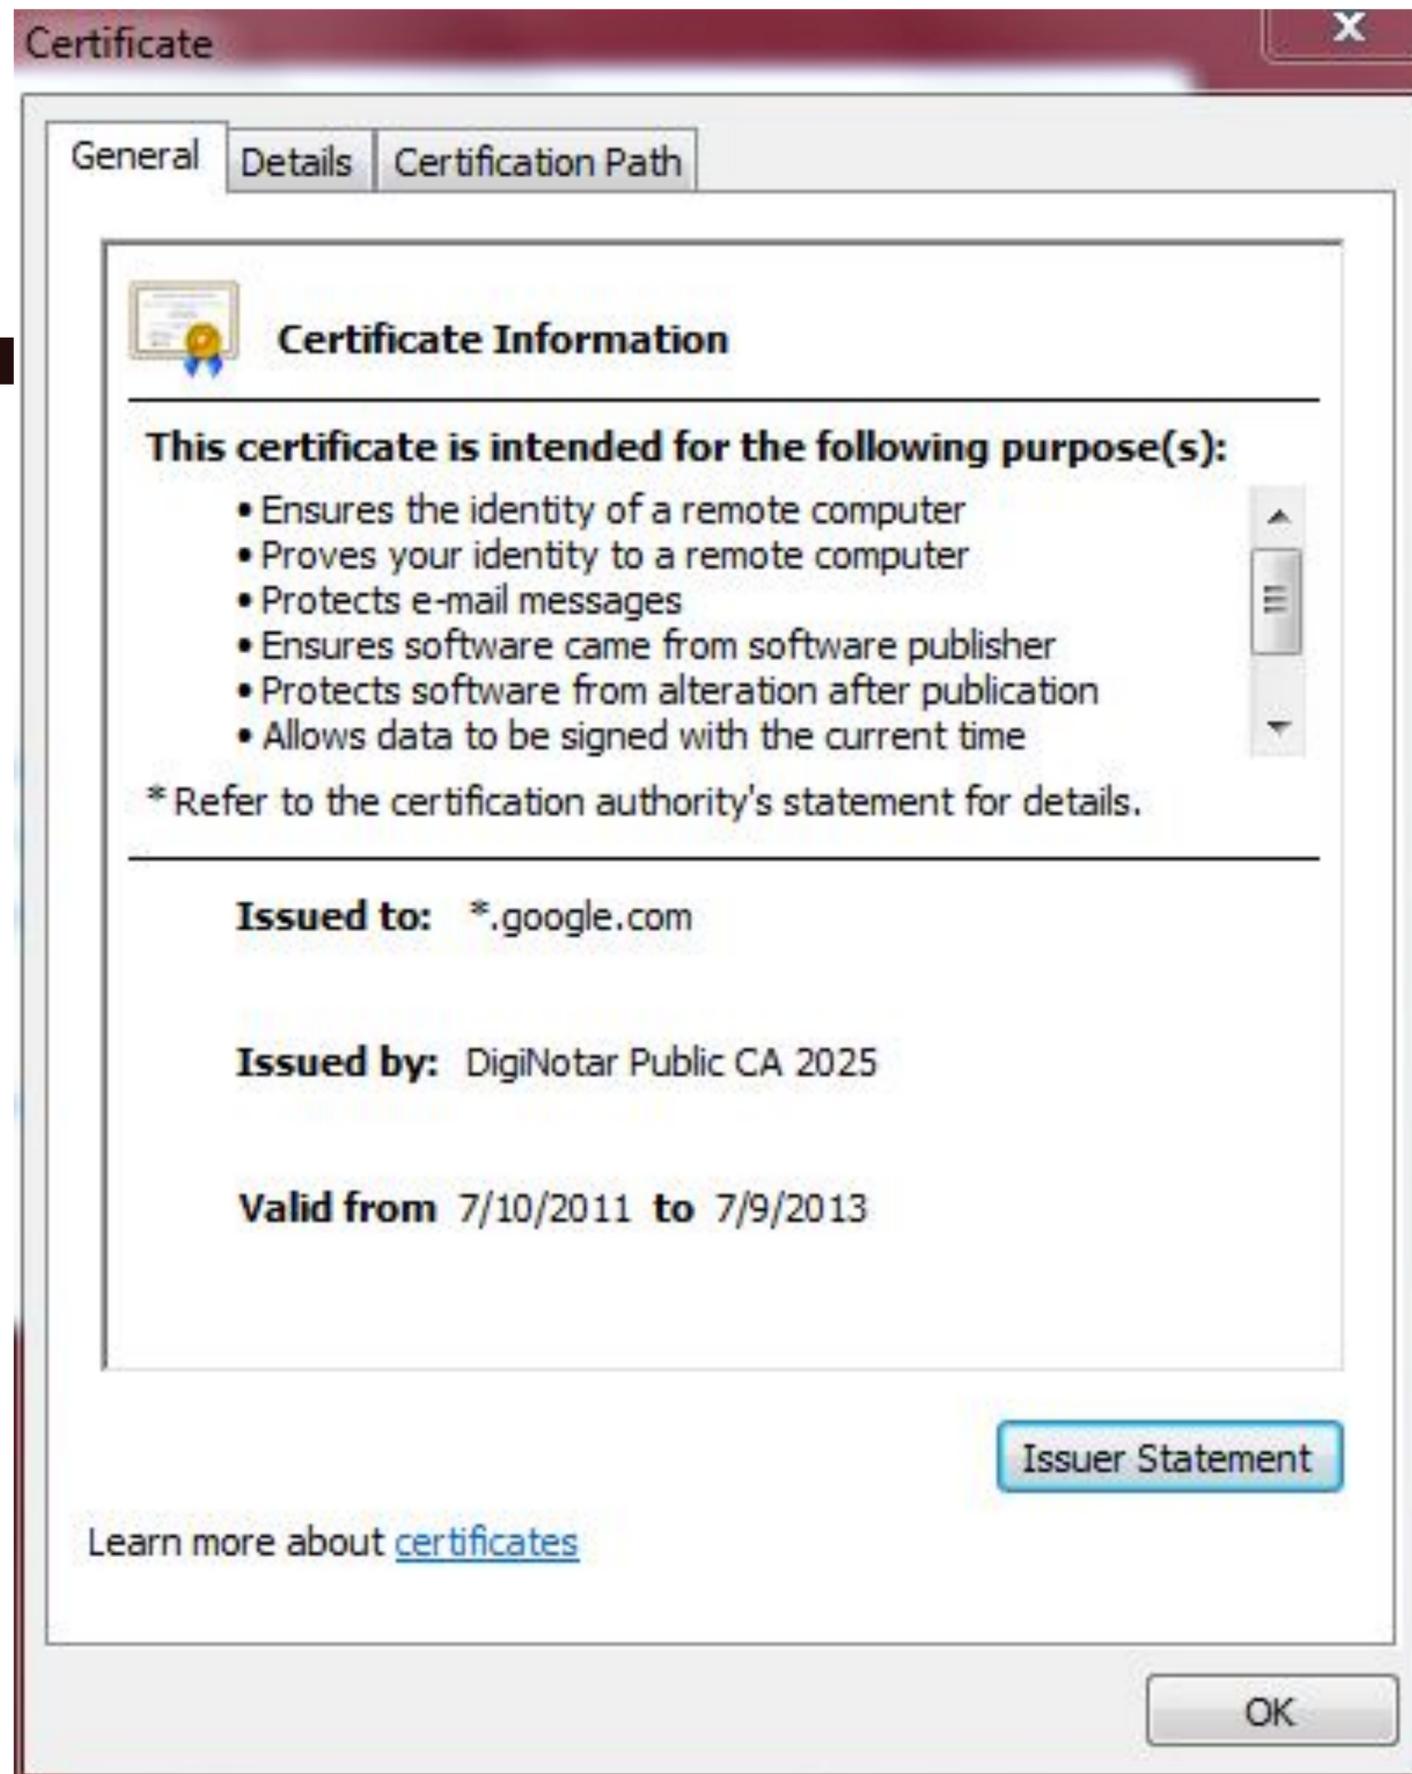
Is Iran behind a fraudulent Google.com digital certificate? The situation is similar to one that happened in March in which spoofed certificates were traced back to Iran.



by [Elinor Mills](#) | August 29, 2011 1:22 PM PDT

A Dutch company appears to have issued a digital certificate for Google.com to someone other than Google, who may be using it to try to re-direct traffic of users based in Iran.

Yesterday, someone reported on a Google support site that when attempting to log in to Gmail the browser issued a warning for the digital certificate used as proof that the site is legitimate, according to [this thread](#) on a Google support forum site.



This appears to be a fully valid cert using normal browser validation rules.

Only detected by Chrome due to its introduction of cert “pinning” – requiring that certs for certain domains must be signed by specific CAs rather than any generally trusted CA

October 31, 2012, 10:49AM

Final Report on DigiNotar Hack Shows Total Compromise of CA Servers

The attacker who penetrated the Dutch CA DigiNotar last year had complete control of all eight of the company's certificate-issuing servers during the operation and he may also have issued some rogue certificates that have not yet been identified. The final report from a

Evidence Suggests DigiNotar, Who Issued Fraudulent Google Certificate, Was Hacked *Years* Ago

from the *diginot* dept

The big news in the security world, obviously, is the fact that a **fraudulent Google certificate made its way out into the wild**, apparently targeting internet users in Iran. The Dutch company DigiNotar has put out a statement saying that **it discovered a breach** back on July 19th during a security audit, and that fraudulent certificates were generated for "several dozen" websites. The only one known to have gotten out into the wild is the Google one.

The DigiNotar Fallout

- The result was the “CA Death Sentence”:
 - Web browsers removed it from the trusted root certificate store
- This has just happened again with “WoSign”
 - A Chinese CA
- WoSign would allow an interesting attack
 - If I controlled `nweaver.github.com`...
 - WoSign would allow me to create a certificate for `*.github.com!?!?`
 - And a bunch of other shady shenanigans

TLS/SSL Trust Issues

- “Commercial certificate authorities protect you from anyone from whom they are unwilling to take money.”
 - Matt Blaze, circa 2001
- So how many CAs do we have to worry about, anyway?
- Of course, it’s not just their greed that matters ...
- ... and it’s not just their diligence & security that matters ...
- *“A decade ago, I observed that commercial certificate authorities protect you from anyone from whom they are unwilling to take money. That turns out to be wrong; they don't even do that much.”* - Matt Blaze, circa 2010

So the Modern Solution: Invoke Ronald Reagan, “Trust, but Verify”

- **Static Certificate Pinning:**
The chrome browser has a list of certificates or certificate authorities that it trusts for given sites
 - Now creating a fake certificate requires attacking a particular CA
- **HPKP Certificate Pinning:**
The web server provides hashes of certificates that should be trusted
 - This is “Leap of Faith”: The first time you assume it is honest but you will catch future changes
- **Transparency mechanisms:**
 - Public logs provided by certificate authorities
 - Browser extensions (EFF’s TLS observatory)
 - Backbone monitors (ICSI’s TLS notary)

Theme of This Lecture In Song: 50 Whys to Stop A Server...

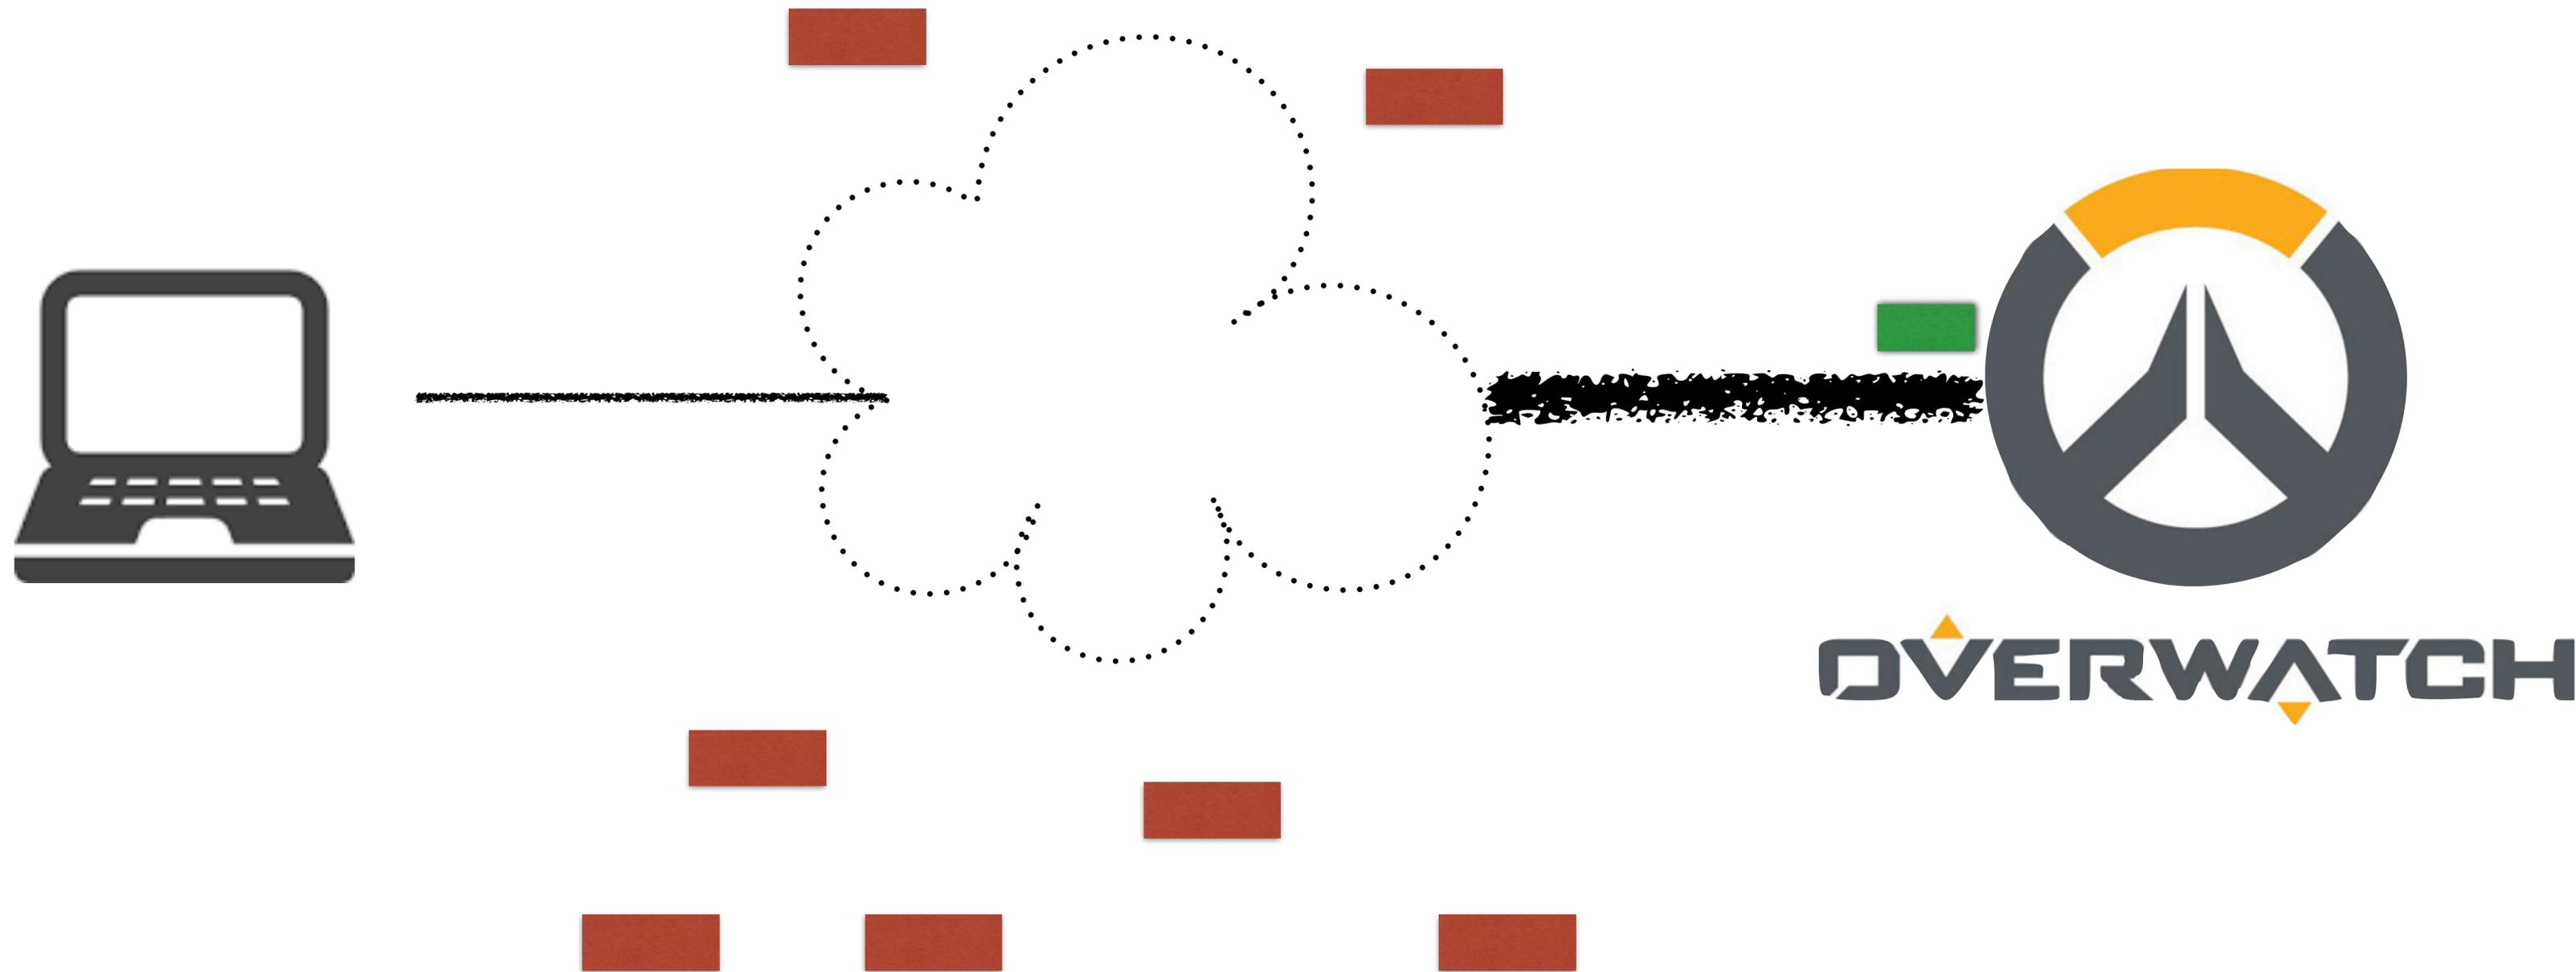
- You are a bad guy...
 - And you want to stop some server from being *available*
- Why? You name it...
 - Because its hard for someone to frag you in an online game if you "boot" him from the network
 - Because people will pay up to stop the attack
 - Because it conveys a political message
 - Get paid for by others



The Easy DoS on a System: Resource Consumption...

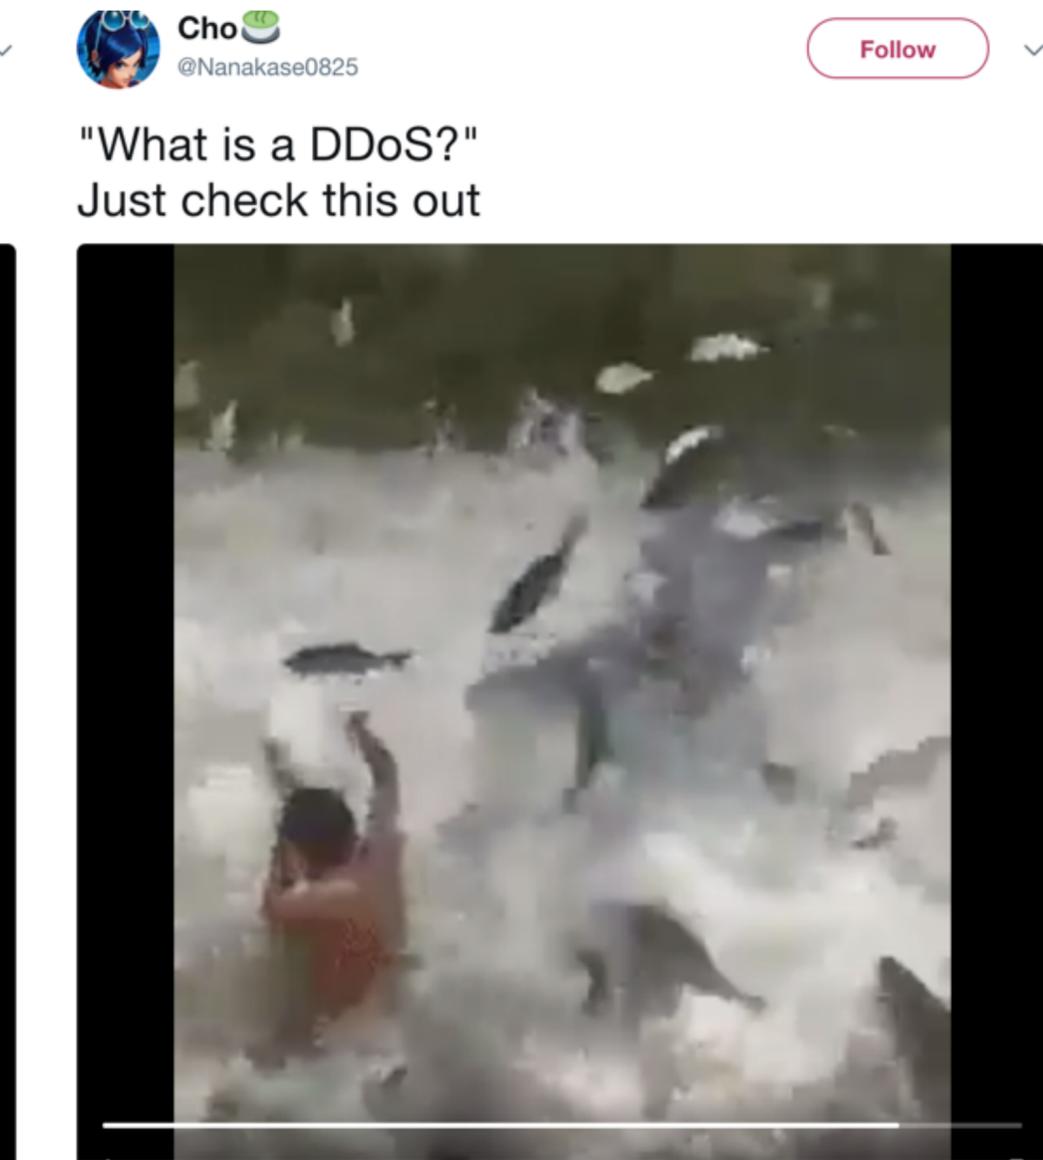
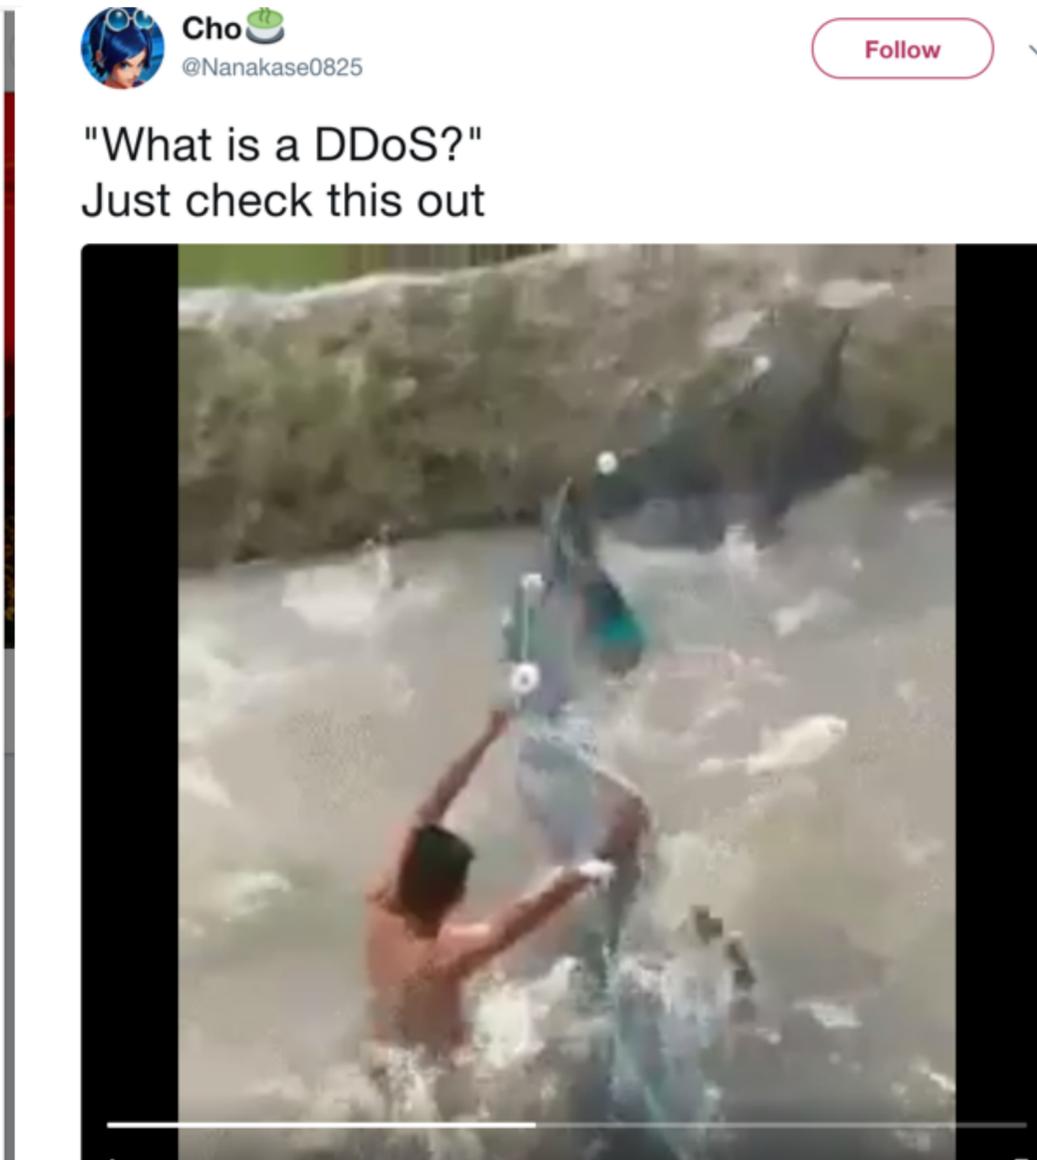
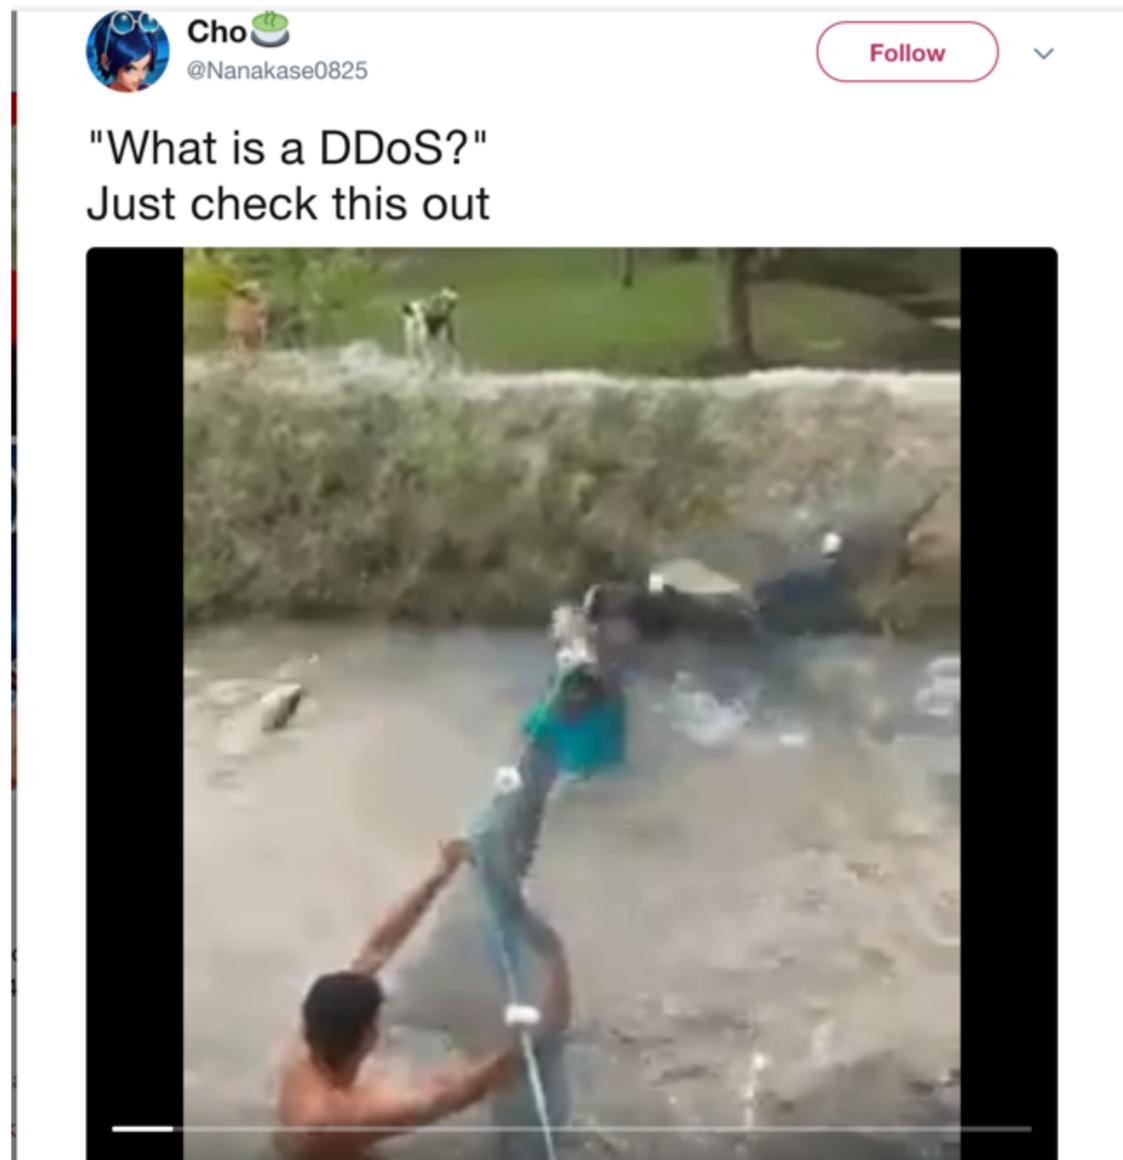
- Bad Dude has an account on your computer...
 - And wants to disrupt your work on Project 2...
- He runs this simple program:
 - while(1):
 - Write random junk to random files
 - (uses disk space, thrashes the disk)
 - Allocate a bunch of RAM and write to it
 - (uses memory)
 - fork()
 - (creates more processes to run)
- Only defense is some form of quota or limits:
The system itself ***must*** enforce some isolation

The Network DOS



Or, another visual explanation...

- <https://twitter.com/kokonoe0825/status/789536739887112192>



DoS & Networks

- How could you DoS a target's Internet access?
 - Send a zillion packets at them
 - Internet lacks *isolation* between traffic of different users!
- What resources does attacker need to pull this off?
 - At least as much sending capacity (*bandwidth*) as the bottleneck link of the target's Internet connection
 - Attacker sends maximum-sized packets
 - Or: overwhelm the rate at which the bottleneck router can process packets
 - Attacker sends minimum-sized packets!
 - (in order to maximize the packet arrival rate)

Defending Against Network DoS

- Suppose an attacker has access to a beefy system with high-speed Internet access (a “big pipe”).
- They pump out packets towards the target at a very high rate.
- What might the target do to defend against the onslaught?
 - Install a network filter to discard any packets that arrive with attacker’s IP address as their source
 - E.g., `drop * 66.31.33.7:* -> *:*`
 - Or it can leverage any other pattern in the flooding traffic that’s not in benign traffic
 - Attacker’s IP address = means of identifying misbehaving user

Filtering Sounds Pretty Easy ...

- ... but DoS filters can be easily evaded:
 - Make traffic appear as though it's from many hosts
 - Spoof the source address so it can't be used to filter
 - Just pick a random 32-bit number of each packet sent
 - How does a defender filter this?
 - They don't!
 - Best they can hope for is that operators around the world implement anti-spoofing mechanisms (today about 75% do)
 - Use many hosts to send traffic rather than just one
 - Distributed Denial-of-Service = DDoS (“dee-doss”)
 - Requires defender to install complex filters
 - How many hosts is “enough” for the attacker?
 - Today they are very cheap to acquire ... :-)

It's Not A "Level Playing Field"

- When defending resources from exhaustion, need to beware of asymmetries, where attackers can consume victim resources with little comparable effort
 - Makes DoS easier to launch
 - Defense costs much more than attack
- Particularly dangerous form of asymmetry: amplification
 - Attacker leverages system's own structure to pump up the load they induce on a resource

Amplification

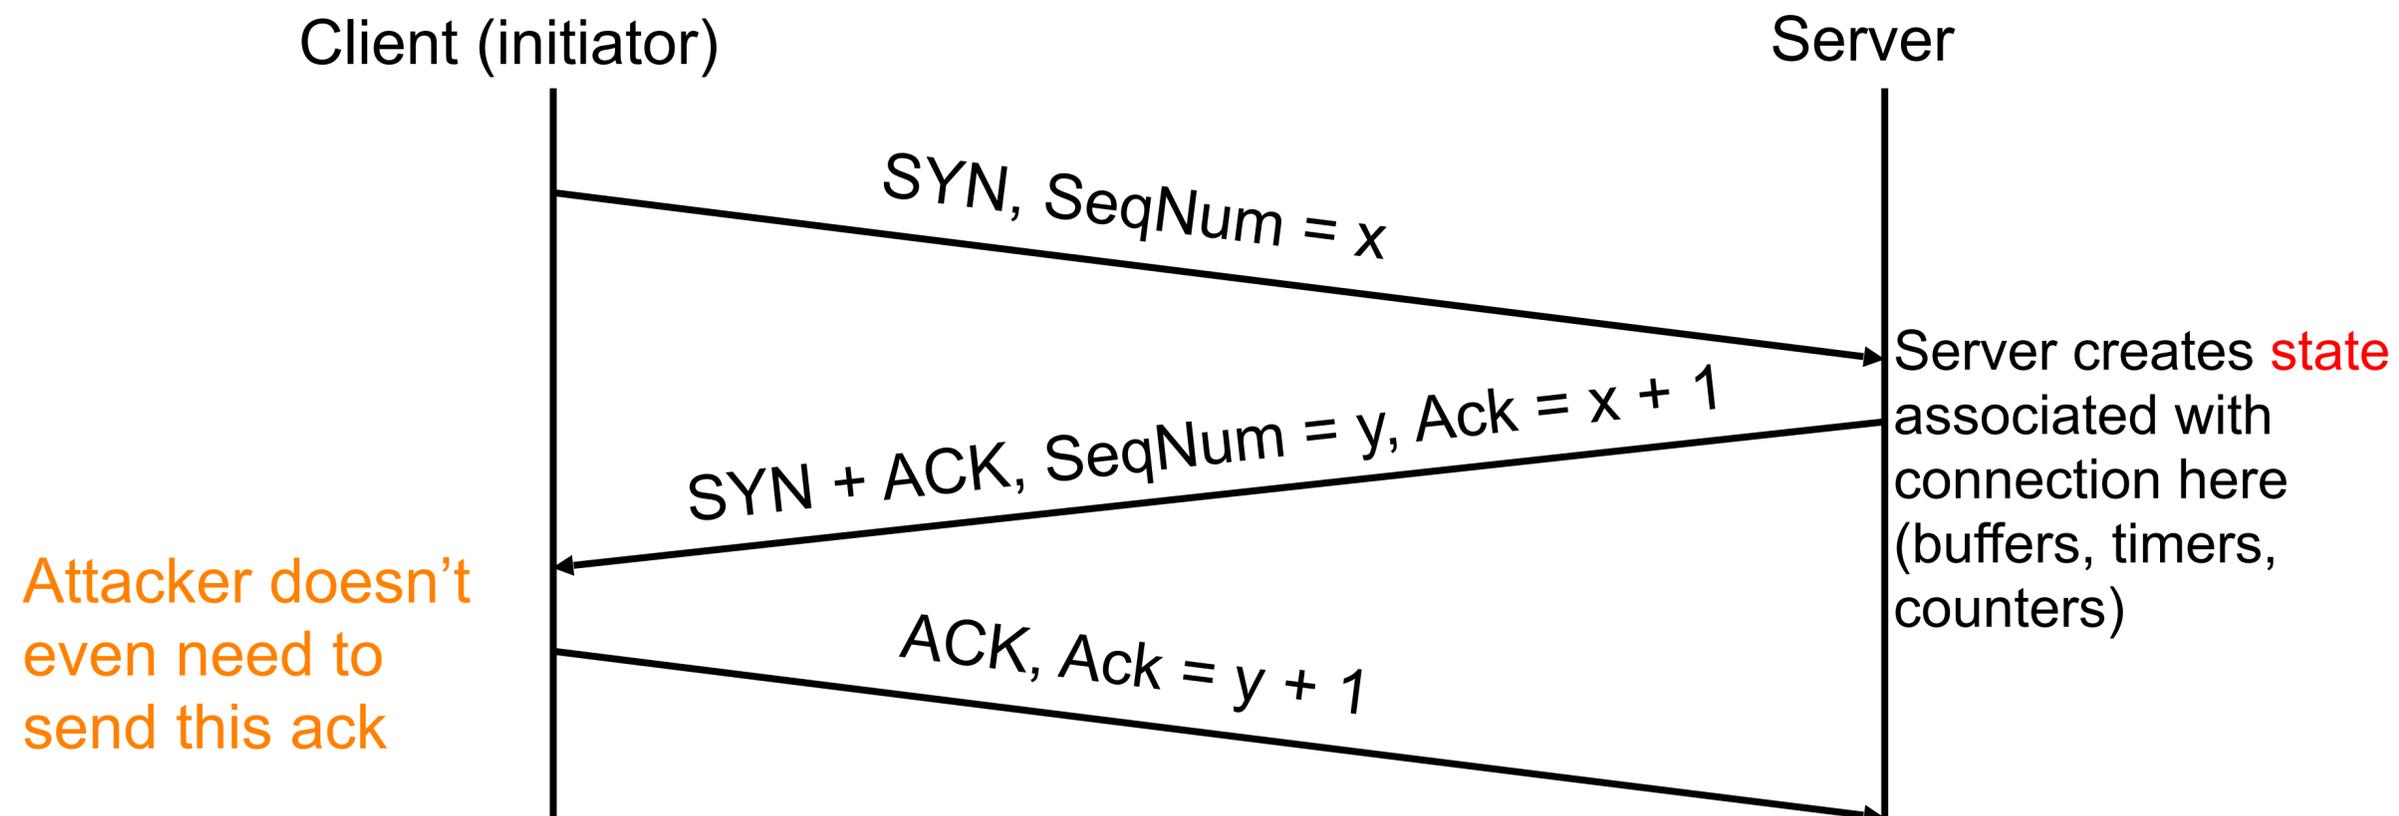
- **Example of amplification: DNS lookups**
 - Reply is generally much bigger than request
 - Since it includes a copy of the reply, plus answers etc.
 - Attacker spoofs DNS request to a patsy DNS server, seemingly from the target
 - Small attacker packet yields large flooding packet
 - Doesn't increase # of packets, but total volume
- **Note #1: these examples involve blind spoofing**
 - So for network-layer flooding, generally only works for UDP-based protocols (can't establish a TCP connection)
- **Note #2: victim doesn't see spoofed source addresses**
 - Addresses are those of actual intermediary systems

Botnets

- If an attacker can control a *lot* of systems
 - They gain a huge amount of bandwidth
 - Modern DOS attacks approach 1 Terabit-per-second with direct connections
 - And it becomes very hard to filter them out
 - How do you specify 1M machines you want to ignore?
- You control these "bots" in a "botnet"
 - So you can issue commands that cause all these systems to do what you want
- This is what took down dyn DNS (and with it Twitter, Reddit, etc...) last year: A botnet composed primarily of compromised cameras and DVRs:
 - The Miraj botnet

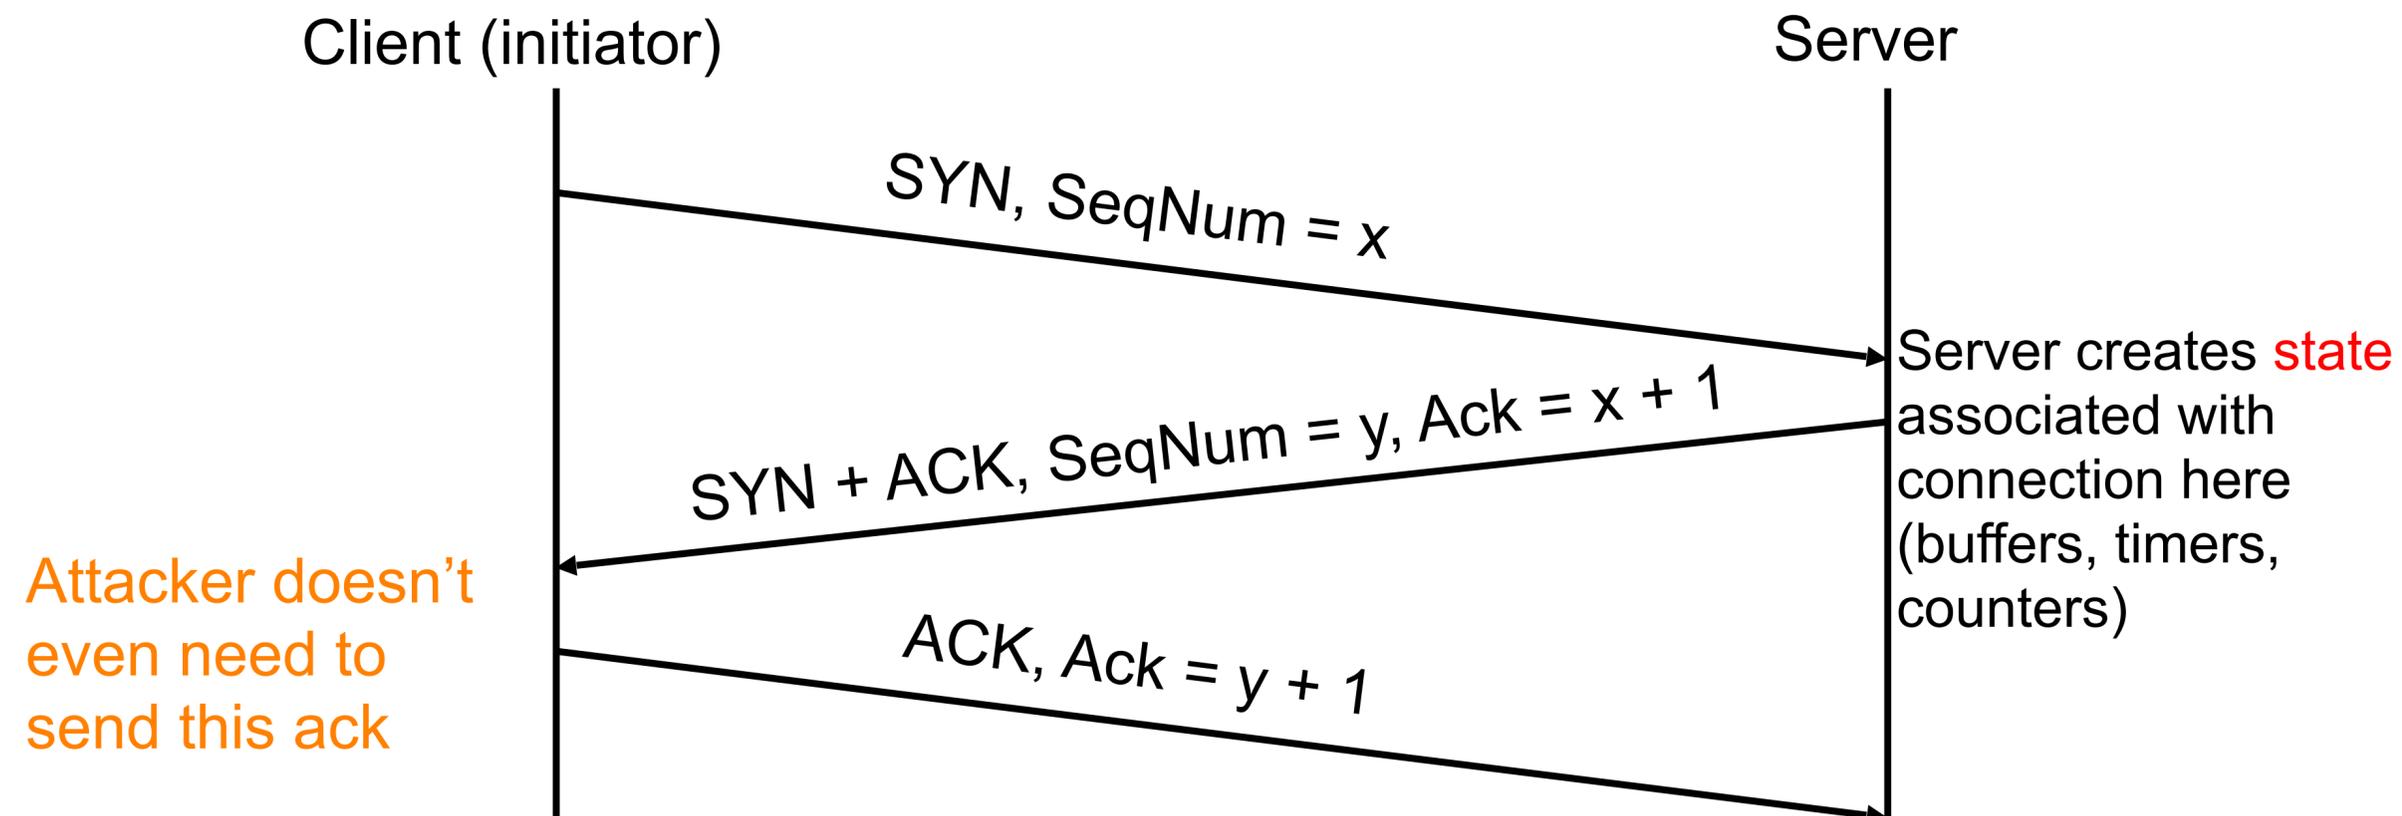
Transport-Level Denial-of-Service

- Recall TCP's 3-way connection establishment handshake
 - Goal: agree on initial sequence numbers



Transport-Level Denial-of-Service

- Recall TCP's 3-way connection establishment handshake
 - Goal: agree on initial sequence numbers
- So a single SYN from an attacker suffices to force the server to spend some memory



TCP SYN Flooding

- Attacker targets memory rather than network capacity
- Every (unique) SYN that the attacker sends burdens the target
- What should target do when it has no more memory for a new connection?
- No good answer!
 - Refuse new connection?
 - Legit new users can't access service
 - Evict old connections to make room?
 - Legit old users get kicked off

TCP SYN Flooding Defenses

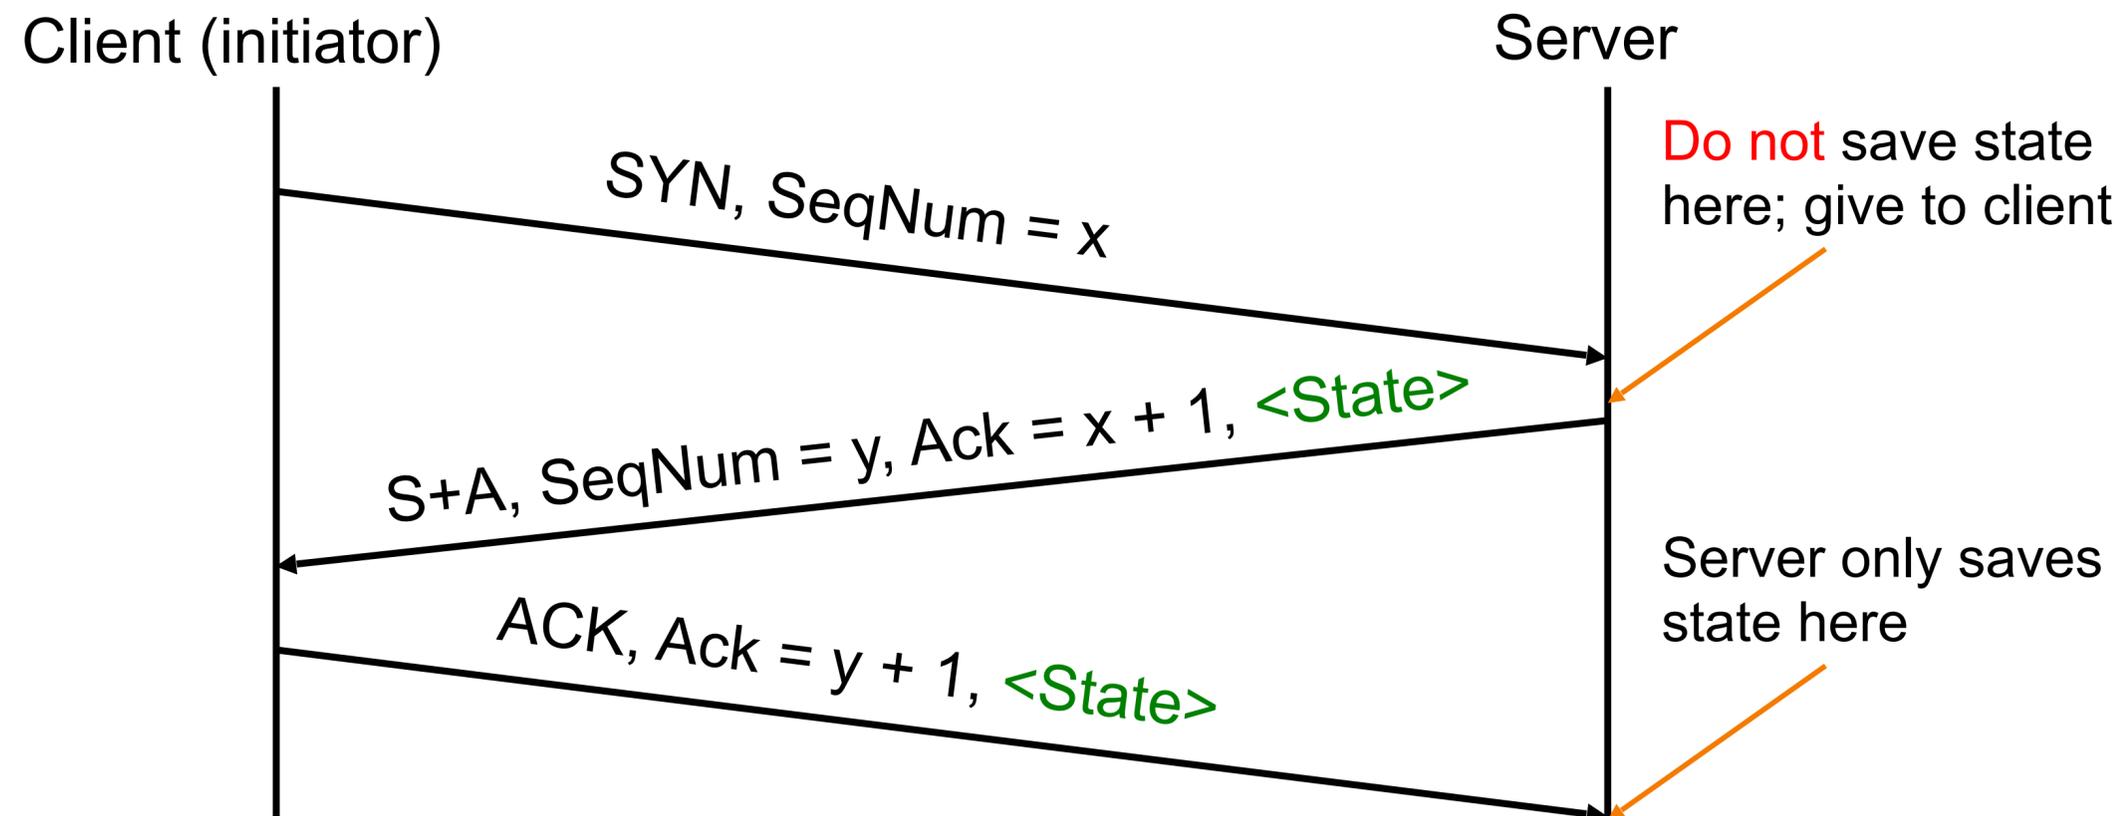
- How can the target defend itself?
- Approach #1: make sure they have tons of memory!
 - How much is enough?
 - Depends on resources attacker can bring to bear (threat model), which might be hard to know
- Back of the envelope:
 - If we need to hold 10kB for 1 minute: to exhaust 1GB, an attacker needs...
 - 100k packets/minute, or a bit more than 1,000 packets per second

TCP SYN Flooding Defenses

- Approach #2: identify bad actors & refuse their connections
 - Hard because only way to identify them is based on IP address
 - We can't for example require them to send a password because doing so requires we have an established connection!
 - For a public Internet service, who knows which addresses customers might come from?
 - Plus: attacker can spoof addresses since they don't need to complete TCP 3-way handshake
- Approach #3: don't keep state! ("SYN cookies"; only works for spoofed SYN flooding)

SYN Flooding Defense: Idealized

- Server: when SYN arrives, rather than keeping state locally, send it to the client ...
- Client needs to return the state in order to establish connection



SYN Flooding Defense: Idealized

- Server: when SYN arrives, rather than keeping state locally, send it to the client . . .

Problem: the world isn't so ideal!

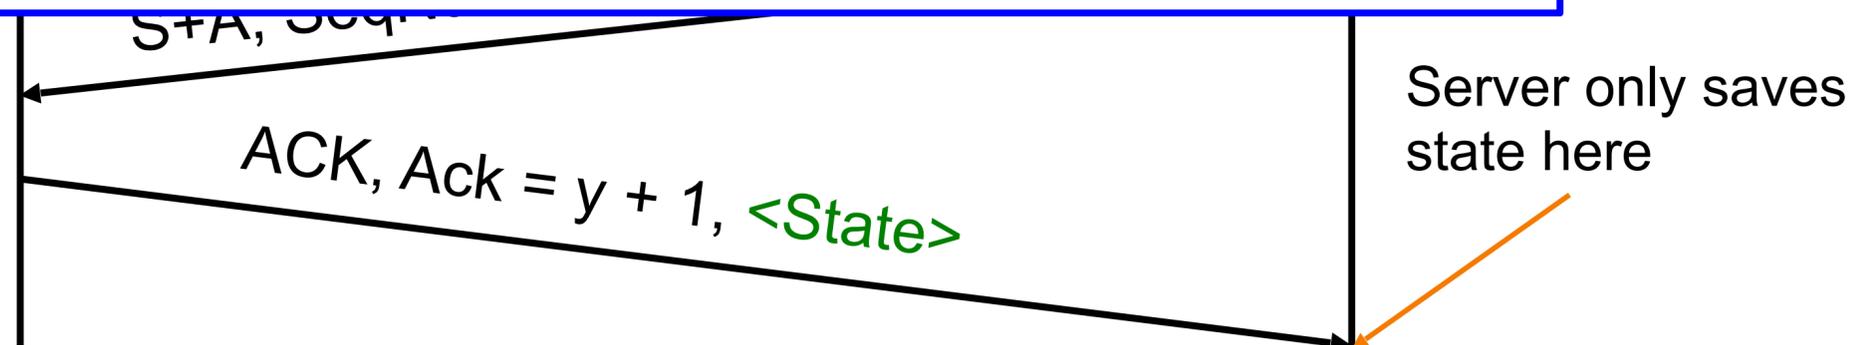
- Client needs to establish connection

TCP doesn't include an easy way to add a new <State> field like this.

Client (

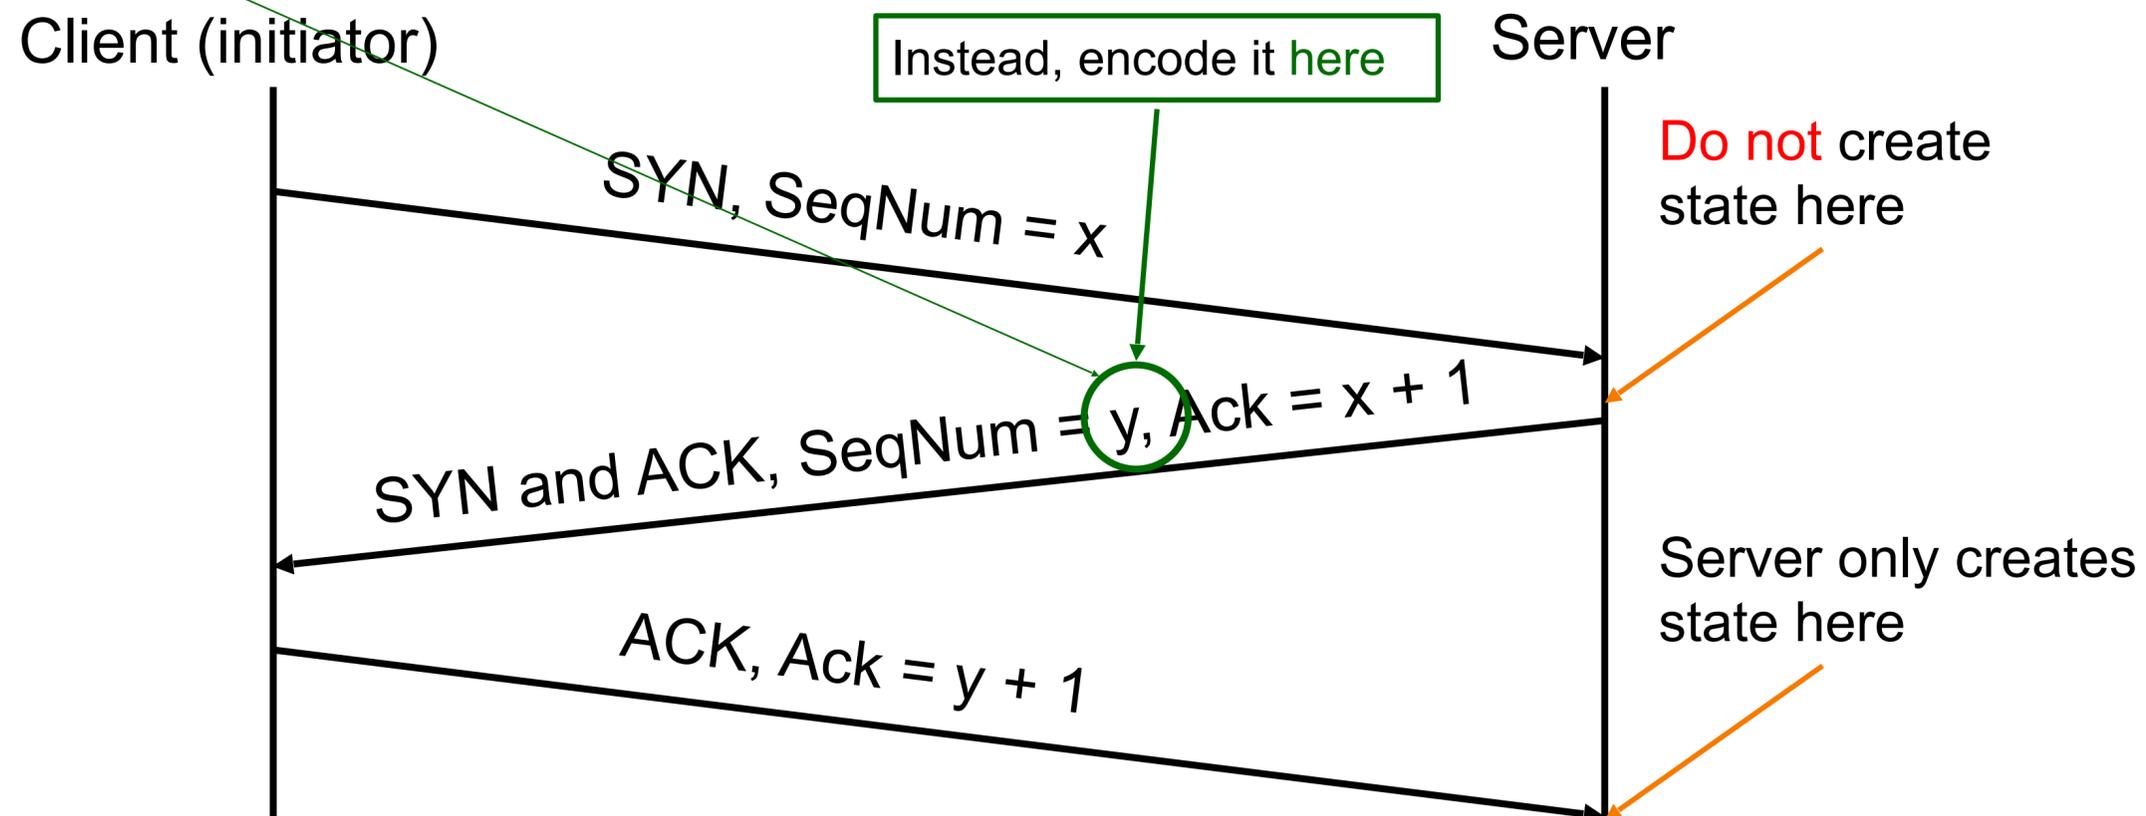
Is there any way to get the same functionality without having to change TCP clients?

save state
give to client



Practical Defense: SYN Cookies

- Server: when SYN arrives, encode connection state entirely within SYN-ACK's sequence # y
- y = encoding of necessary state, using server secret
- When ACK of SYN-ACK arrives, server only creates state if value of y from it agrees w/ secret



SYN Cookies: Discussion

- Illustrates general strategy: rather than holding state, encode it so that it is returned when needed
- For SYN cookies, attacker must complete 3-way handshake in order to burden server
 - Can't use spoofed source addresses
- Note #1: strategy requires that you have enough bits to encode all the state
 - (This is just barely the case for SYN cookies)
 - You can think of a SYN cookie as a truncated MAC of the sender IP/port/sequence
- Note #2: if it's expensive to generate or check the cookie, then it's not a win

Application-Layer DoS

- Rather than exhausting network or memory resources, attacker can overwhelm a service's processing capacity
- There are many ways to do so, often at little expense to attacker compared to target (asymmetry)

The Ethereum network is currently undergoing a DoS attack

Posted by [Jeffrey Wilcke](#) on [September 22nd, 2016](#).

URGENT ALL MINERS: The network is under attack. The attack is a computational DDoS, ie. miners and nodes need to spend a very long time processing some blocks. This is due to the `EXTCODESIZE` opcode, which has a fairly low gasprice but which requires nodes to read state information from disk; the attack transactions are calling this opcode roughly 50,000 times per block. The consequence of this is that the network is greatly slowing down, but there is **NO** consensus failure or memory overload. We have currently identified several routes for a more sustainable medium-term fix and have developers working on implementation.

It is highly recommended to switch to Parity mining. Use these settings:

Algorithmic complexity attacks

- Attacker can try to trigger worst-case complexity of algorithms / data structures
- Example: You have a hash table.
Expected time: $O(1)$. Worst-case: $O(n)$.
- Attacker picks inputs that cause hash collisions.
Time per lookup: $O(n)$.
Total time to do n operations: $O(n^2)$.
- Solution? Use algorithms with good worst-case running time.
 - E.g., using b bits of HMAC ensures that $P[h_k(x)=h_k(y)] = .5^b$, so hash collisions will be rare.
 - If the attacker doesn't know the key that is

Application-Layer DoS

- Defenses against such attacks?
- Approach #1: Only let legit users issue expensive requests
 - Relies on being able to identify/authenticate them
 - Note: that this itself might be expensive!
- Approach #2: Force legit users to “burn” cash
 - This is what a captcha really is!
- Approach #3: massive over-provisioning (\$\$\$)
 - Or pay for someone else who massively over provisions for everyone:
A content delivery network

DoS Defense in General Terms

- Defending against program flaws requires:
 - Careful design and coding/testing/review
 - Consideration of behavior of defense mechanisms
 - E.g. buffer overflow detector that when triggered halts execution to prevent code injection ⇒ denial-of-service
- Defending resources from exhaustion can be really hard.
Requires:
 - Isolation and scheduling mechanisms
 - Keep adversary's consumption from affecting others
 - Reliable identification of different users
 - Or just a ton of \$\$\$\$

Controlling Networks ... On The Cheap

- Motivation: How do you harden a set of systems against external attack?
 - Key Observation:
 - The more network services your machines run, the greater the risk
 - Due to larger attack surface
- One approach: on each system, turn off unnecessary network services
 - But you have to know all the services that are running
 - And sometimes some trusted remote users still require access
- Plus key question of scaling
 - What happens when you have to secure 100s/1000s of systems?
 - Which may have different OSs, hardware & users ...
 - Which may in fact not all even be identified ...

Taming Management Complexity

- Possibly more scalable defense: Reduce risk by blocking in the network outsiders from having unwanted access your network services
 - Interpose a firewall the traffic to/from the outside must traverse
 - Chokepoint can cover thousands of hosts
 - Where in everyday experience do we see such chokepoints?



Selecting a Security Policy

- Firewall enforces an (access control) policy:
 - Who is allowed to talk to whom, accessing what service?
- Distinguish between inbound & outbound connections
 - Inbound: attempts by external users to connect to services on internal machines
 - Outbound: internal users to external services
 - Why? Because fits with a common threat model. There are thousands of internal users (and we've vetted them). There are billions of outsiders.
- Conceptually simple access control policy:
 - Permit inside users to connect to any service
 - External users restricted:
 - Permit connections to services meant to be externally visible
 - Deny connections to services not meant for external access

How To Treat Traffic Not Mentioned in Policy?

- Default Allow: start off permitting external access to services
 - Shut them off as problems recognized
- Default Deny: start off permitting just a few known, well-secured services
 - Add more when users complain (and mgt. approves) ✓
- Pros & Cons?
 - Flexibility vs. conservative design
 - Flaws in Default Deny get noticed more quickly / less painfully

In general, use Default Deny

A Dumb Policy: Deny All Inbound connections...

- The simplest packet filters are *stateless*
 - They examine only individual packets to make a decision
- But even the simplest policy can be hard to implement
 - Deny All Inbound is the default policy on your home connection
- Allow:
 - Any outbound packet
 - Any inbound packet that is a reply... OOPS
- We can fake it for TCP with some ugly hacks
 - Allow all outbound TCP
 - Allow all inbound TCP that does not have both the SYN flag set and the ACK flag not set
 - May still allow an attacker to play some interesting games
- We can't even fake this for UDP!

Stateful Packet Filter

- Stateful packet filter is a router that checks each packet against security rules and decides to forward or drop it
 - Firewall keeps track of all connections (inbound/outbound)
 - Each rule specifies which connections are allowed/denied (access control policy)
 - A packet is forwarded if it is part of an allowed connection



Example Rule

- **allow tcp connection 4.5.5.4:* -> 3.1.1.2:80**
- Firewall should permit TCP connection that's:
 - Initiated by host with Internet address 4.5.5.4 and
 - Connecting to port 80 of host with IP address 3.1.1.2
- Firewall should permit any packet associated with this connection
- Thus, firewall keeps a table of (allowed) active connections. When firewall sees a packet, it checks whether it is part of one of those active connections. If yes, forward it; if no, check to see if rule should create a new allowed connection

Example Rule

- `allow tcp connection *:* /int -> 3.1.1.2:80/ext`
- Firewall should permit TCP connection that's:
 - Initiated by host with any internal host and
 - Connecting to port 80 of host with IP address 3.1.1.2 on external Internet
- Firewall should permit any packet associated with this connection
- The `/int` indicates the network interface.
- This is "Allow all outgoing web requests"

Example Ruleset

- `allow tcp connection *:* /int -> *:* /ext`
- `allow tcp connection *:* /ext -> 1.2.2.3:80 /int`
- Firewall should permit outbound TCP connections (i.e., those that are initiated by internal hosts)
- Firewall should permit inbound TCP connection to our public webserver at IP address 1.2.2.3

Stateful Filtering

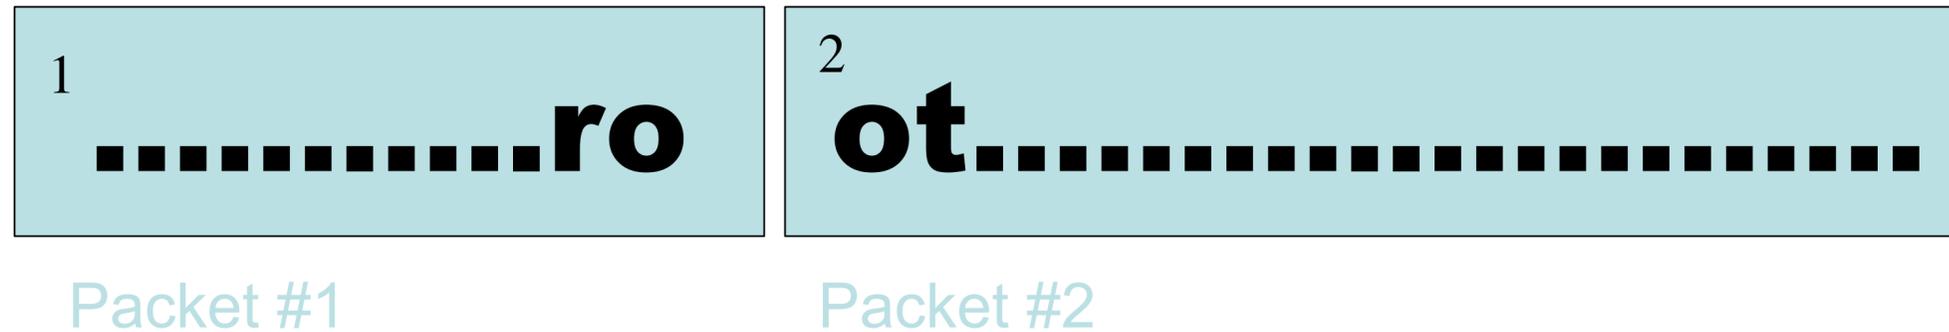
- Suppose you want to allow inbound connection to a FTP server, but block any attempts to login as “root”. How would you build a stateful packet filter to do that? In particular, what state would it keep, for each connection?

State Kept

- No state – just drop any packet with root in them
- Is it a FTP connection?
- Where in FTP state (e.g. command, what command)
- Src ip addr, dst ip addr, src port, dst port
- Inbound/outbound connection
- Keep piece of login command until it's completed – only first 5 bytes of username

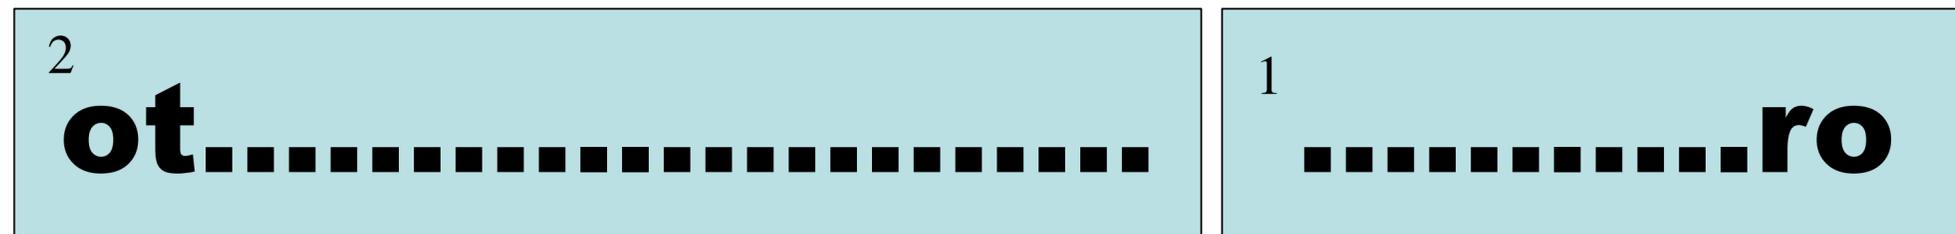
Beware!

- Sender might be malicious and trying to sneak through firewall
- “root” might span packet boundaries

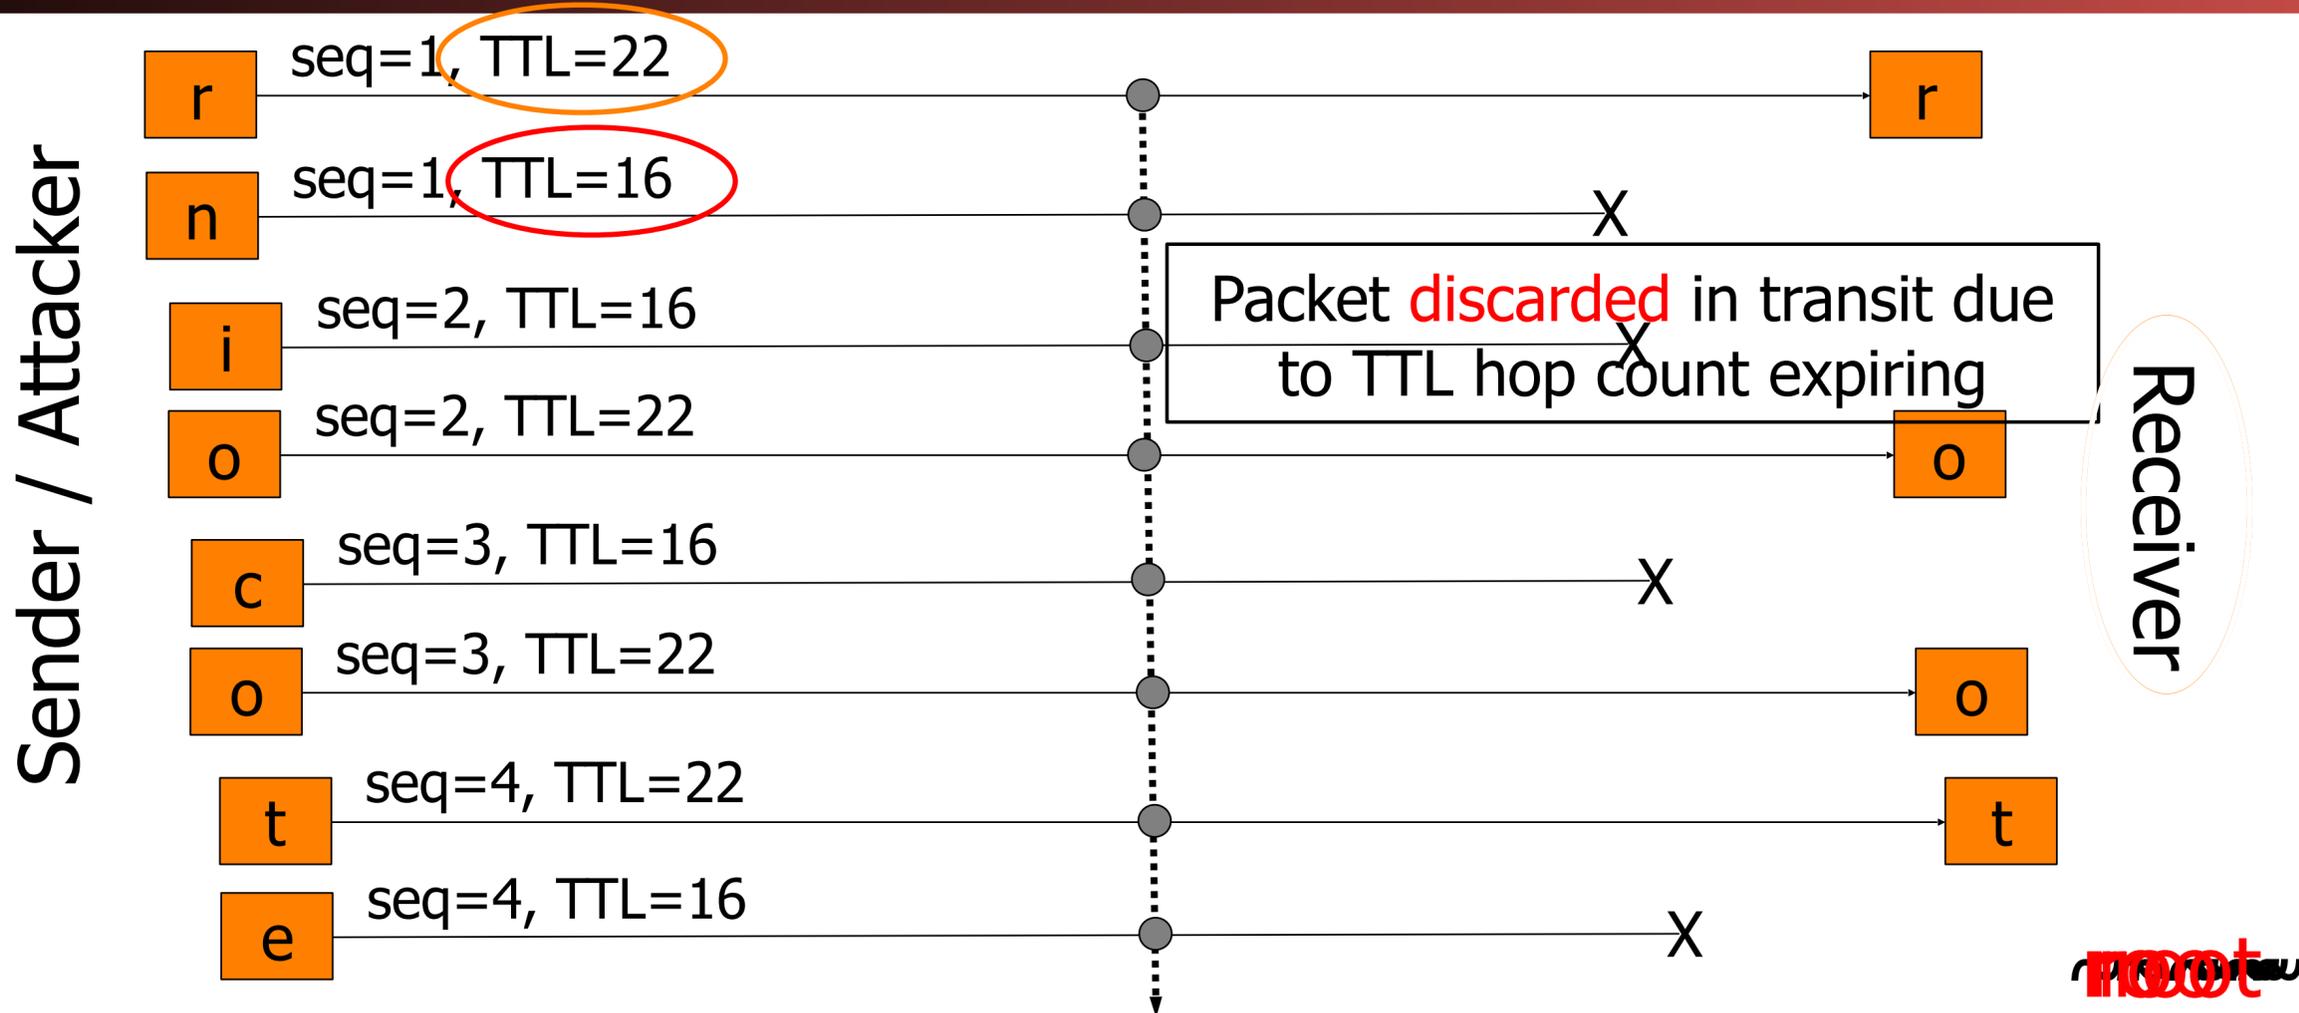


Beware!

- Packets might be re-ordered



Beware!



TTL field in IP header specifies maximum forwarding hop count

rice? roce? rict? roct? riot?
 root? roe? roe? roe? roe?
 rice? roe? roe? roe? roe?
 root? roe? roe? roe? roe?

Assume the Receiver is 20 hops away

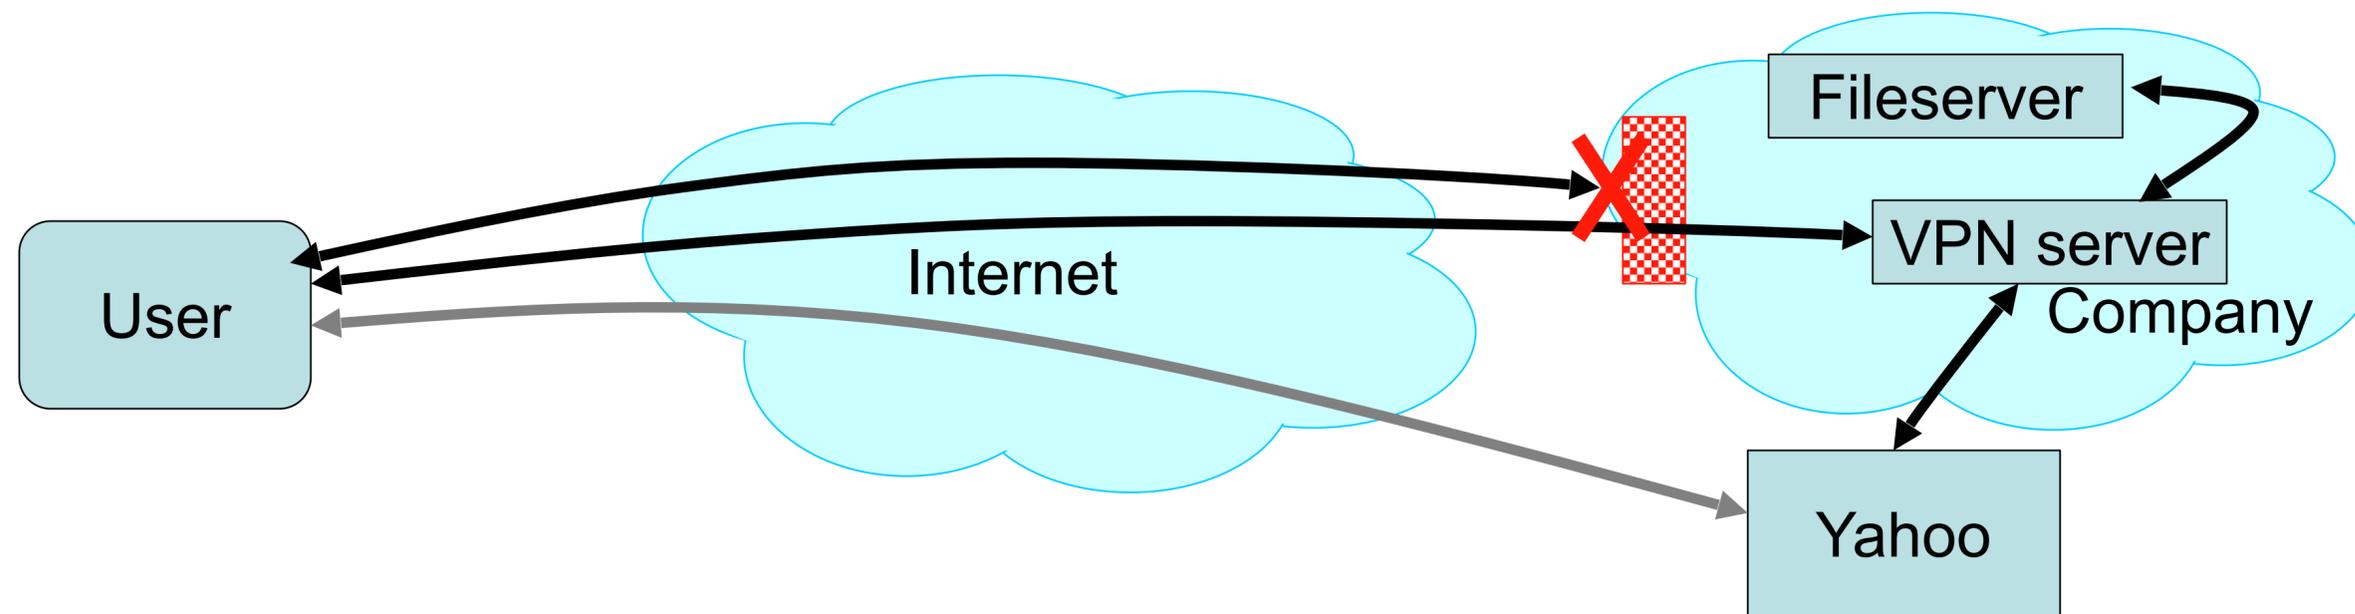
Assume firewall is 15 hops away

Other Kinds of Firewalls

- Application-level firewall
 - Firewall acts as a proxy. TCP connection from client to firewall, which then makes a second TCP connection from firewall to server.
 - Only modest benefits over stateful packet filter.

Secure External Access to Inside Machines

- Often need to provide secure remote access to a network protected by a firewall
 - Remote access, telecommuting, branch offices, ...
- Create secure channel (Virtual Private Network, or VPN) to tunnel traffic from outside host/network to inside network
 - Provides Authentication, Confidentiality, Integrity
 - However, also raises perimeter issues
 - (Try it yourself at <http://www.net.berkeley.edu/vpn/>)



Why Have Firewalls Been Successful?

- **Central control – easy administration and update**
 - Single point of control: update one config to change security policies
 - Potentially allows rapid response
- **Easy to deploy – transparent to end users**
 - Easy incremental/total deployment to protect 1000's
- **Addresses an important problem**
 - Security vulnerabilities in network services are rampant
 - Easier to use firewall than to directly secure code ...

Firewall Disadvantages

- **Functionality loss – less connectivity, less risk**
 - May reduce network's usefulness
 - Some applications don't work with firewalls
 - Two peer-to-peer users behind different firewalls
- **The malicious insider problem**
 - Assume insiders are trusted
 - Malicious insider (or anyone gaining control of internal machine) can wreak havoc
- **Firewalls establish a security perimeter**
 - Like Eskimo Pies: “hard crunchy exterior, soft creamy center”
 - Threat from travelers with laptops, cell phones, ...

Pivoting...

- Thus the goal of the attacker is to "pivot" through the system
 - Start running on a single victim system
 - EG, using a channel that goes from the victim to the attacker's server over port 443: an encrypted web connection
- From there, you can now exploit internal systems directly
 - Bypassing the primary firewall
- That is the problem: **A *single* breach of the perimeter by an attacker and you can no longer make *any* assertions about subsequent internal state**

Takeaways on Firewalls

- Firewalls: Reference monitors and access control all over again, but at the network level
- Attack surface reduction
- Centralized control