A Perfect CRIME? TIME Will Tell

Tal Be'ery, Web research TL

Agenda

- BEAST
  - + Modes of operation
- CRIME
  - + Gzip compression
  - + Compression + encryption leak data
- TIME
  - + Timing + compression leak data
- Attacking responses

BEAST

BEAST

- Rizzo and Duong - 2011
- Browser Exploit Against SSL/TLS (BEAST)
- Chosen Plaintext Attack
- Targets deterministic Initialization Vectors of Cipher-Block Chaining (CBC)

Chosen Plaintext Attack Model

- **A chosen-plaintext attack (CPA)** is an attack model for cryptanalysis which presumes that the attacker has the capability to choose arbitrary plaintexts to be encrypted and obtain the corresponding ciphertexts.

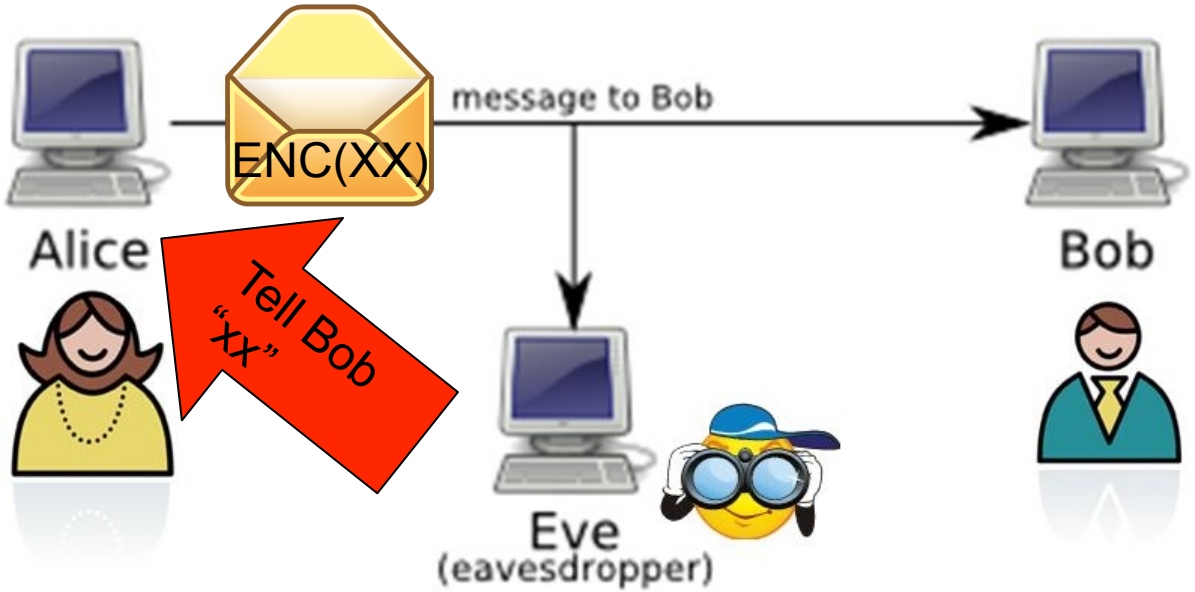

CPA and the web

- Attacker is Eavesdropper – can see ciphered text
- Attacker creates HTTP request interactively (via script)
  - + Full control (almost): URL
  - + Can predict: Most headers
  - + Does not control or see: cookies
    - – Encrypted on wire
    - – Not accessible from script
      - Same Origin Policy
      - "HTTP only"

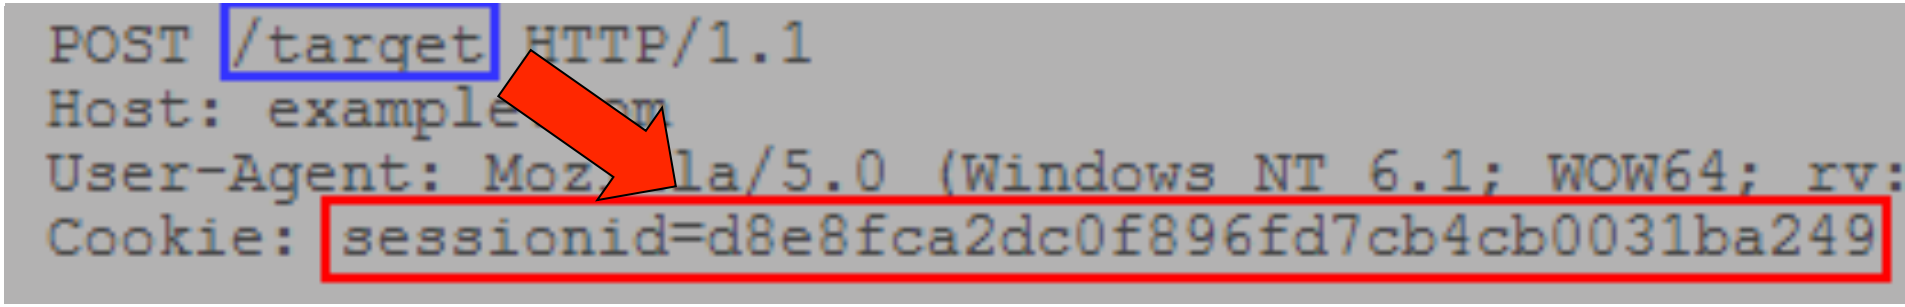

Modes of operation

- procedure of enabling the repeated and secure use of a block cipher under a single key

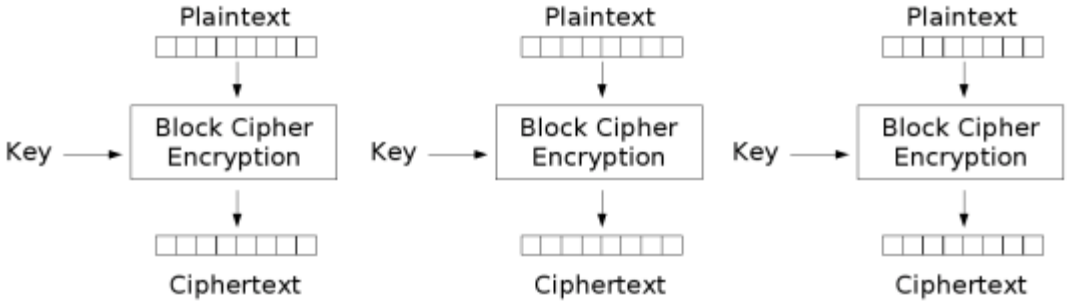

Electronic Codebook (ECB) mode encryption

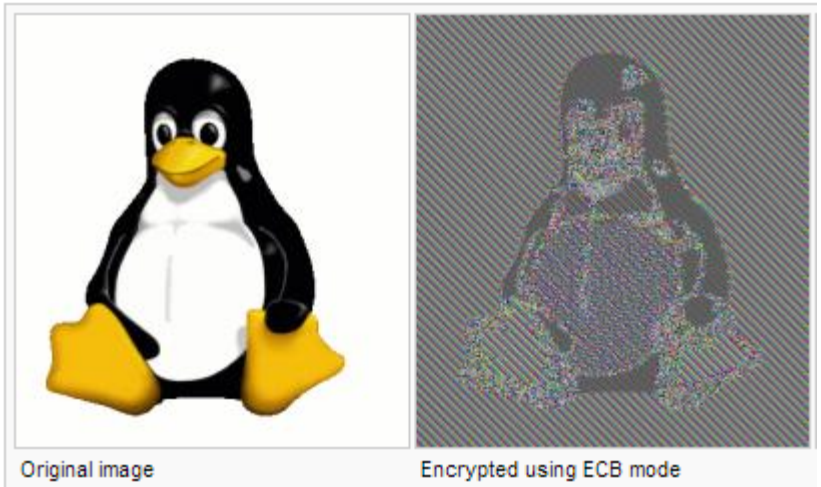

Original image | Encrypted using ECB mode

Modes of operation - CBC

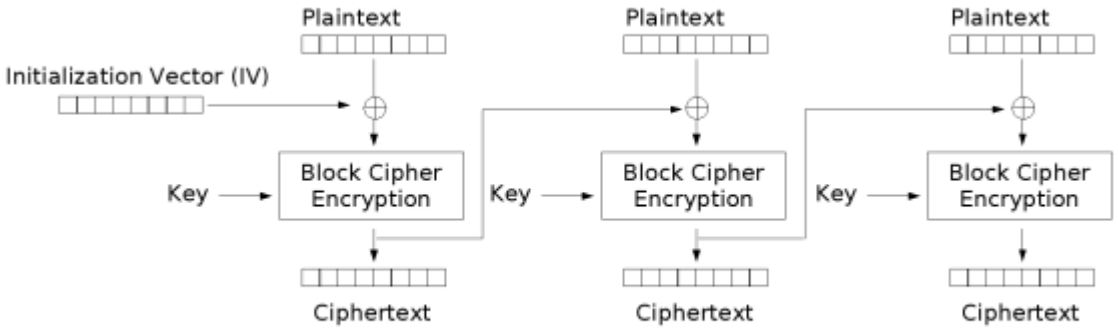

Cipher Block Chaining (CBC) mode encryption

- Previous block encryption result is fed as an IV to the next block
- Encryption becomes “Stateful”

CBC Oracle

- Attacker can verify a guess of any plaintext block

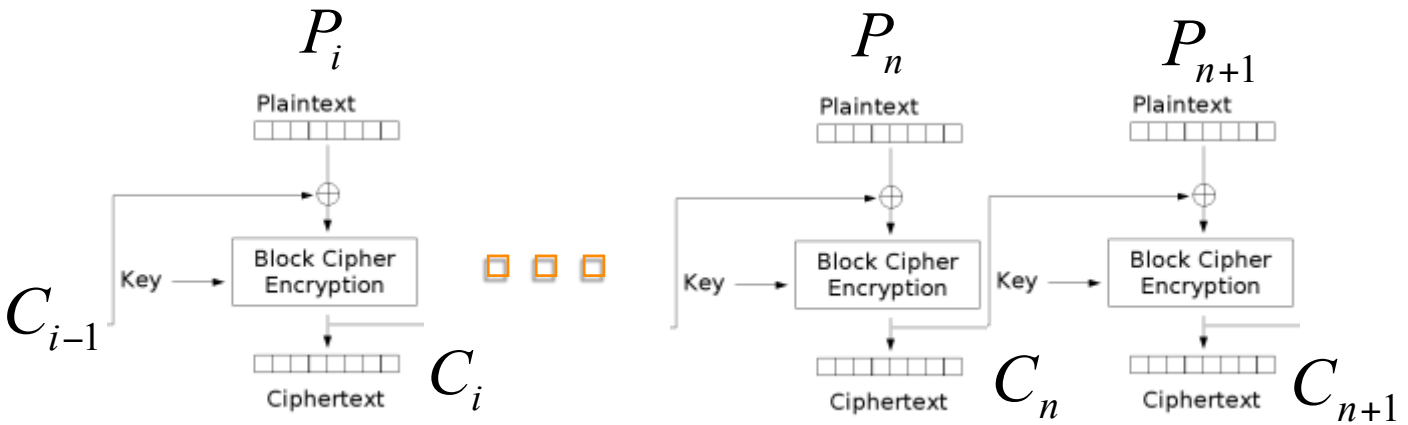

$$P_{n+1} = C_n \oplus C_{i-1} \oplus \tilde{P}_i$$

$$C_{n+1} = Enc(P_{n+1} \oplus C_n) = Enc(C_n \oplus C_{i-1} \oplus \tilde{P}_i \oplus C_n) = Enc(C_{i-1} \oplus \tilde{P}_i)$$

$$\tilde{P}_i = P_i \Rightarrow C_{n+1} = Enc(C_{i-1} \oplus P_i) = C_i$$

$$\tilde{P}_i \neq P_i \Rightarrow C_{n+1} \neq C_i$$

Using the CBC oracle to decrypt the Cookie

- Attacker knows in which block the cookie resides
- Attacker controls the block contents so she can guess only one byte at a time and verify with the oracle
  - 256 guesses on worst case
- Repeat the process to discover all bytes in Cookie

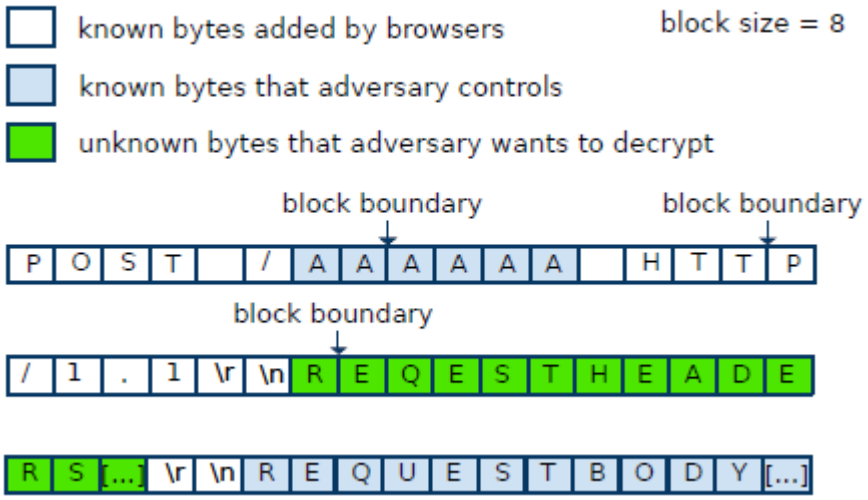

Practical issues

- HTTP requests are [illegible] for BEAST:
  - + New request[illegible]
  - + First bytes [illegible]GET /POST /, e[illegible]
  - + URL can[illegible] Only some cha[illegible]are allowed
- The attac[illegible]eeds [illegible]i-directiona[illegible]ection
  - + Web s[illegible]Java, Sil[illegible]
  - + All of t[illegible]hnologies re[illegible]
- So to expl[illegible]ra vulnerbal[illegible]
  - + SOP bug i[illegible]mentation
  - + XSS in victim

And yet...

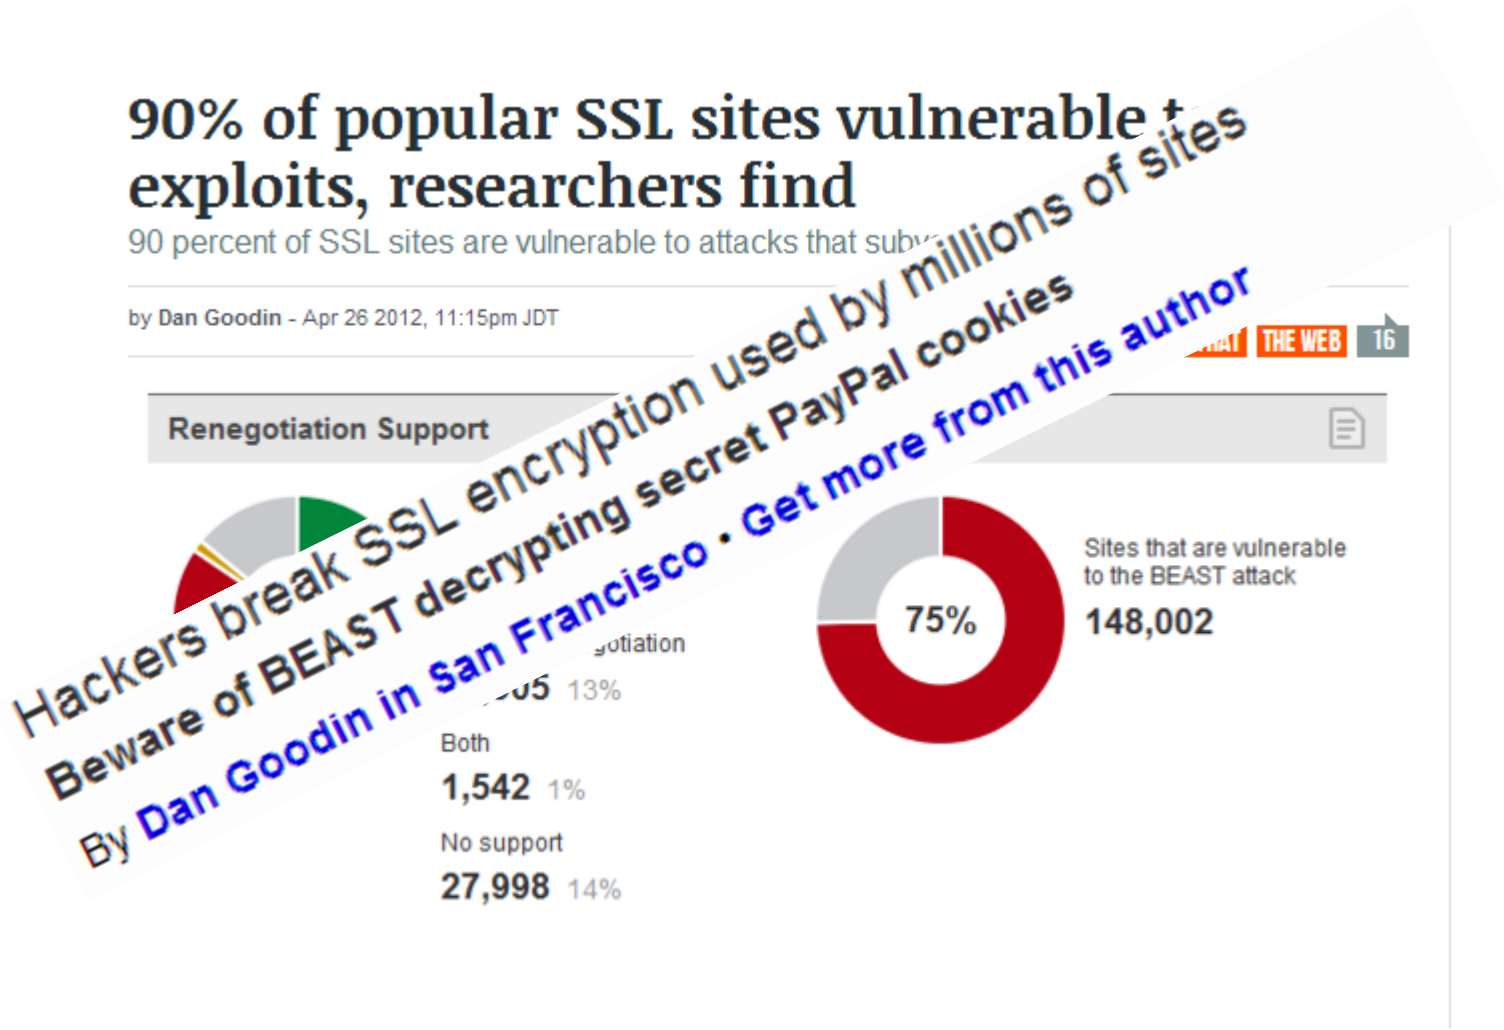

Mitigations

- TLS 1.1 mitigates
  - Explicit IV
  - Not widely adopted
- Some advise to switch to SSL with stream ciphers
  - RC4

CRIME

CRIME

- Rizzo and Duong – 2012
- Compression Ratio Info-leak Made Easy (CRIME)
- Chosen Plaintext Attack
- Targets compression information leakage

Compression – LZ algorithms

- Lempel Ziv, late 70s
- Compress repeating strings
  - + Lossless
  - + Asymptotically optimal
  - + No overhead (No extra dictionary)

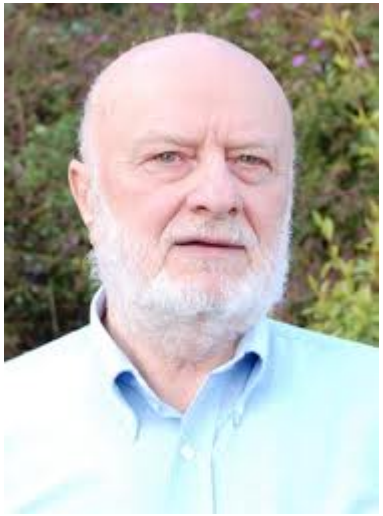

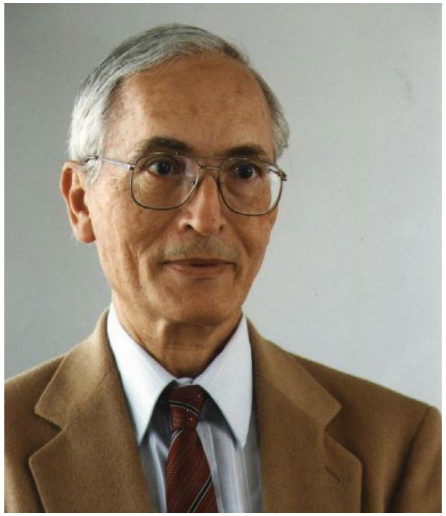

LZ Compression – Example

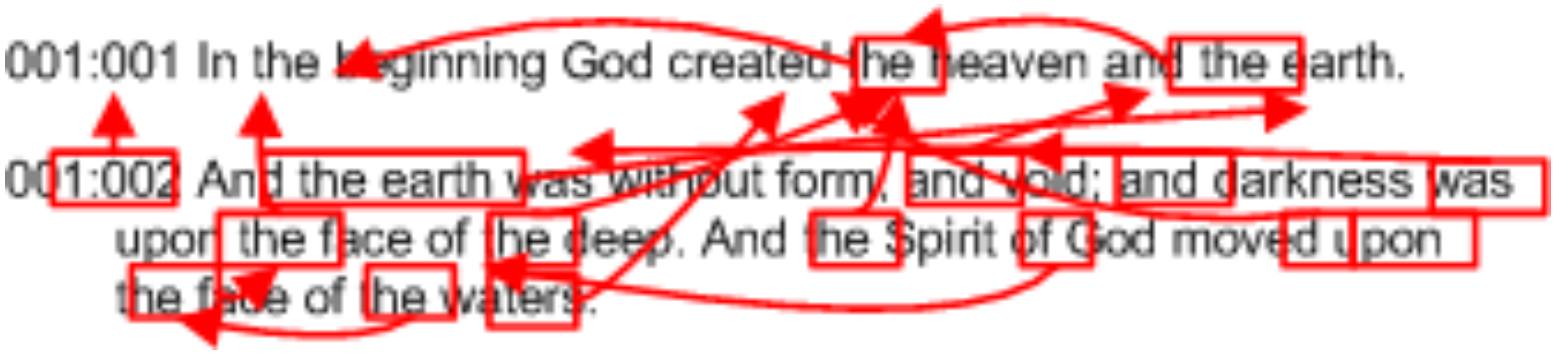

- 001:001 In the beginning God created<25, 5>heaven an<14, 6>earth. 0<63, 5>2 A<23, 12> was without form,<55, 5>void;<9, 5>darkness<40, 4> <0, 7>upo<132, 6>face of<11, 5>deep.<93, 9>Spirit<27, 4><158, 4>mov<156, 3><54, 4><67, 9><62, 16>w<191, 3>rs

Huffman code

- David Huffman - 1952
- Assign shorter codes (in bits) for frequent letters
- Note - Prefix code is a must!
  - Since we cannot rely on length to parse

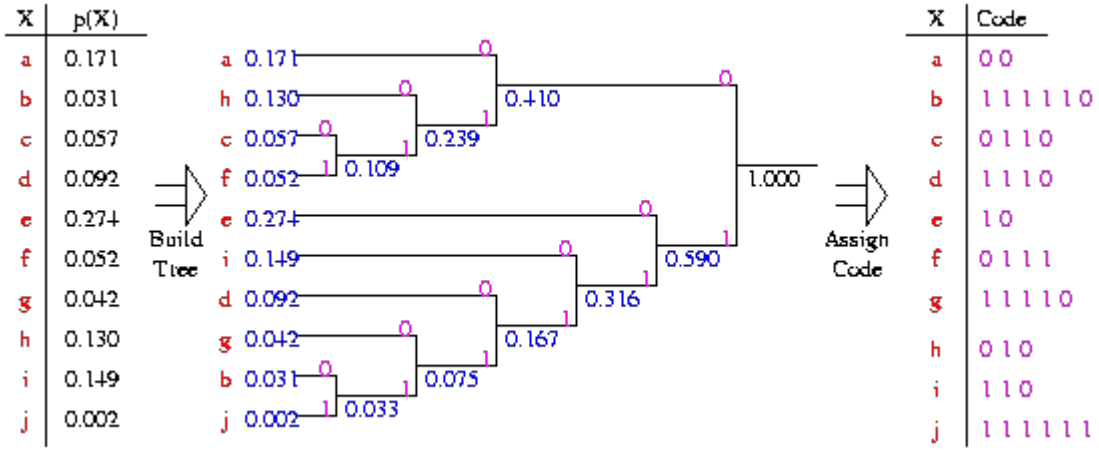

Compression & Encryption

Compression & Encryption

!

Compression on the web

- Content compression
  - + GZIP on response
  - + On request body (Uncommon)
- Header compression
  - + SSL/TLS Compression
    - – Servers: Open SSL, others
    - – Clients: Chrome
  - + SPDY
    - – Servers: Apache MOD_SSL, others
    - – Clients: All but IE

Compression leaks data

- Again
  - + Use the URL attacker controls
  - + Guess byte by byte
  - + Verify with an oracle
    - – If we had guessed correctly then packet size will be shorter

```
POST /sessionid=a HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:
Cookie: sessionid=d8e8fca2dc0f896fd7cb4cb0031ba249
```

```
POST /sessionid=d HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:
Cookie: sessionid=d8e8fca2dc0f896fd7cb4cb0031ba249
```

CRIME in a slide

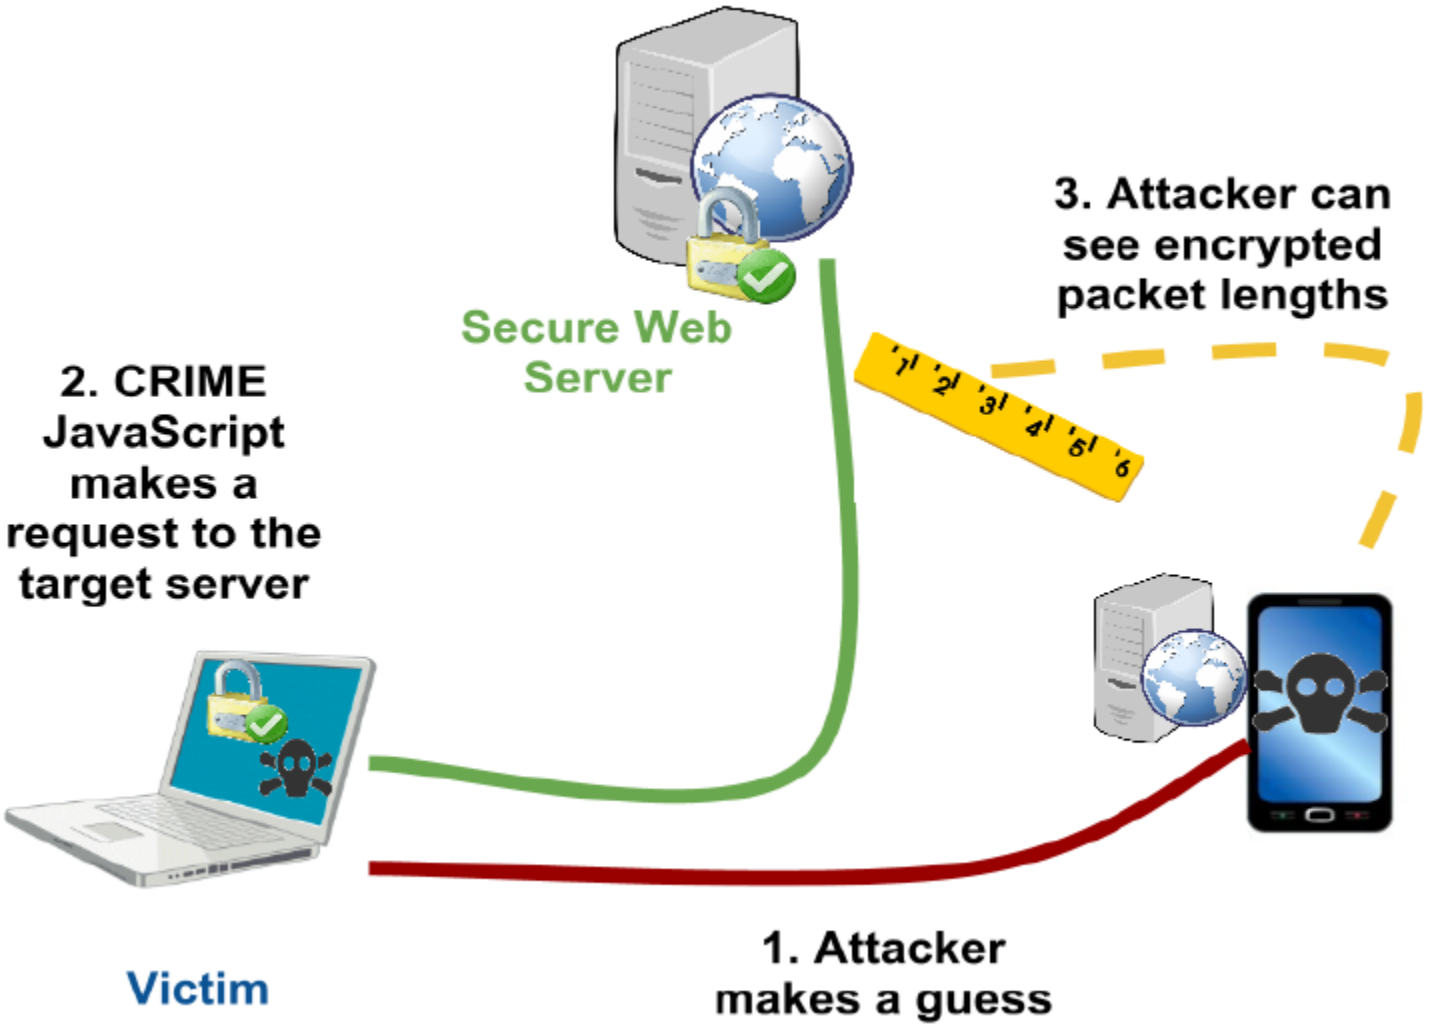

oung & Rizzo original presentation
https://docs.google.com/presentation/d/11eBmGiHbYcHR9gL5nDyZChu_-ICa2GizeuOfaLU2HOU/present#slide=id.g1d134dff_1_157

Practical issues

- HTTP requests are a good vehicle for CRIME:
  - New requests over SPDY use the same SSL connection and compression context
  - The controlled part is “location tolerant”
  - The controlled part can express needed alphabet
- Some issues with Huffman coding
  - Some chars representation < 1 byte
  - Good guess might get unnoticed
- Solutions
  - Mostly tricks to make GZIP compress with not so aggressive Huffman coding

Impact

- Actual impact
  - + SPDY implementations cancel/modify header compression
  - + Chrome disabled SSL compression
- PR Impact
  - + Much less than BEAST
  - + The boy who cried BEAST syndrome

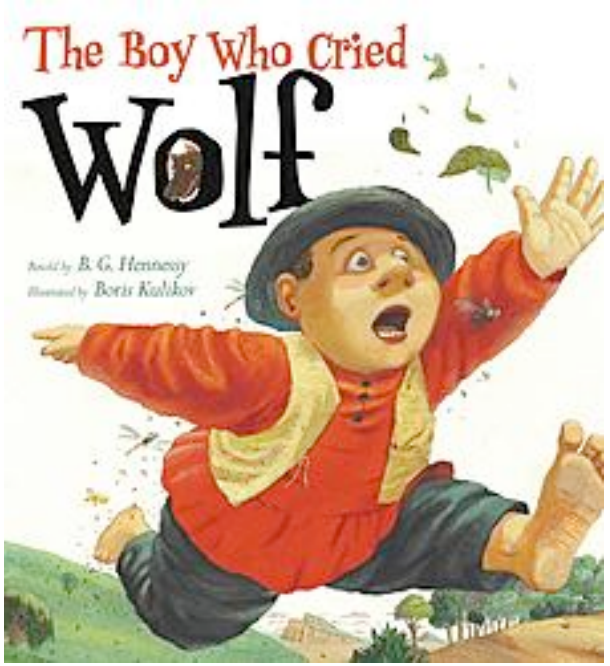

TIME

TIME

- Imperva – 2013
- Timing Info-leak Made Easy (TIME)
- Chosen Plaintext Attack
- Targets compression and timing information leakage

Attack Model

- Attacker has the capability to choose arbitrary plaintexts to be compressed and obtain timing observations on their traffic
- Attacker is no longer an Eavesdropper - attack might be useful against plaintext too!

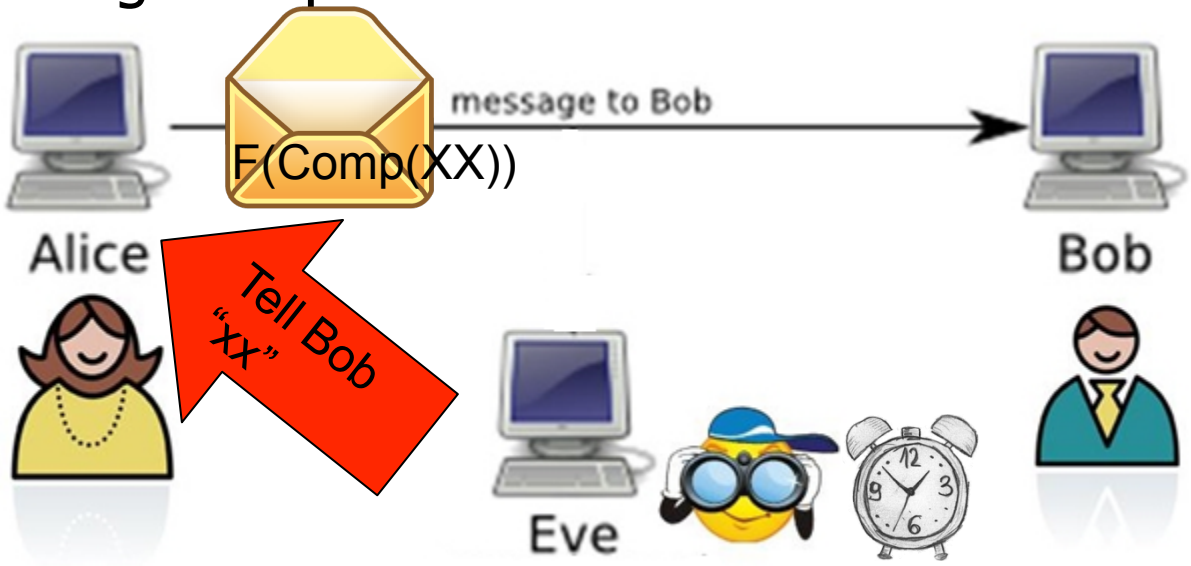

Timing oracle

- Client send a window of TCP packets
- Waits RTT for ACK to send another
- RTT time is noticeable
- attacker can easily distinguish
  - + Size(request) <= window
  - + Size(request) > window
- If payload length is exactly on data boundary, attacker can determine 1 byte differences

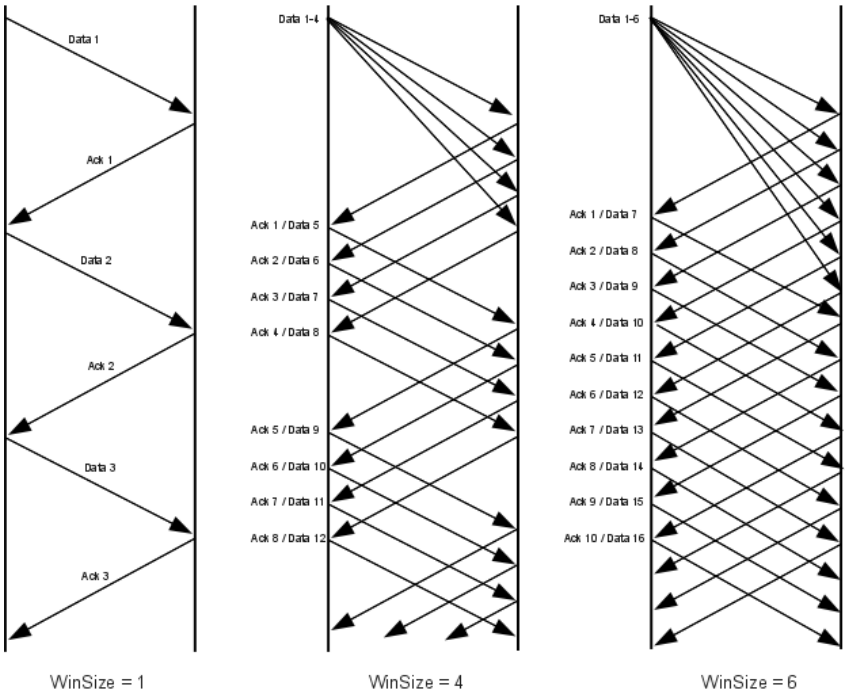

Sliding Windows, bandwidth 6 packets/RTT

http://ulam2.cs.luc.edu/ebook/chap03.html

HTTP Request's Time Measurements

- Create HTTP request with XHR
  - XHR adheres to SOP
  - Allows GET requests to flow
    - If headers allow show response
    - If not, abort
  - We don't care for the response
  - Timing leaks the request size
- Use getTime() on XHR events
  - onreadystatechange
- Noise elimination
  - Repeat the process (say 10 times) and obtain **Minimal** time

Compression leaks data

- Again
  - + Use the URL attacker controls
  - + Guess byte by byte
  - + Verify with an oracle
    - – If we had guessed correctly: packet size will be shorter and so will the time

```
POST /sessionid=a HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:
Cookie: sessionid=d8e8fca2dc0f896fd7cb4cb0031ba249
```

```
POST /sessionid=d HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:
Cookie: sessionid=d8e8fca2dc0f896fd7cb4cb0031ba249
```

RTT Gap in the wild

- Sent with Chrome
- Sends 2 packets and wait
- If you need to send 3 packets – pay extra RTT

No.	Time	Protocol	Length	Info
2284	0.000000000	TCP	66	27983 > http [SYN] Seq=0 Win=8192 Len=0 MSS=1460 W
2298	0.177681000	TCP	66	http > 27983 [SYN, ACK] Seq=0 Ack=1 Win=14480 Len=
2299	0.000092000	TCP	54	27983 > http [ACK] Seq=1 Ack=1 Win=65536 Len=0
2317	0.183176000	TCP	1514	[TCP segment of a reassembled PDU]
2318	0.000016000	TCP	1514	[TCP segment of a reassembled PDU]
2326	0.169969000	TCP	60	http > 27983 [ACK] Seq=1 Ack=1461 Win=8960 Len=0
2327	0.000052000	HTTP	55	GET /?FTYnCuZg9XheUnuABl7mM9aUGk7XtutuTdxsybNa9imA
2328	0.000039000	TCP	60	http > 27983 [ACK] Seq=1 Ack=2921 Win=11776 Len=0
2332	0.167268000	TCP	60	http > 27983 [ACK] Seq=1 Ack=2922 Win=11776 Len=0
2333	0.006509000	TCP	1502	[TCP segment of a reassembled PDU]

RTT Gap in the wild – implementing the Oracle

- HTML with Javascript Sending method is XHR
- Testing cnn.com
- Timing can be correctly captured
- Results are conclusive

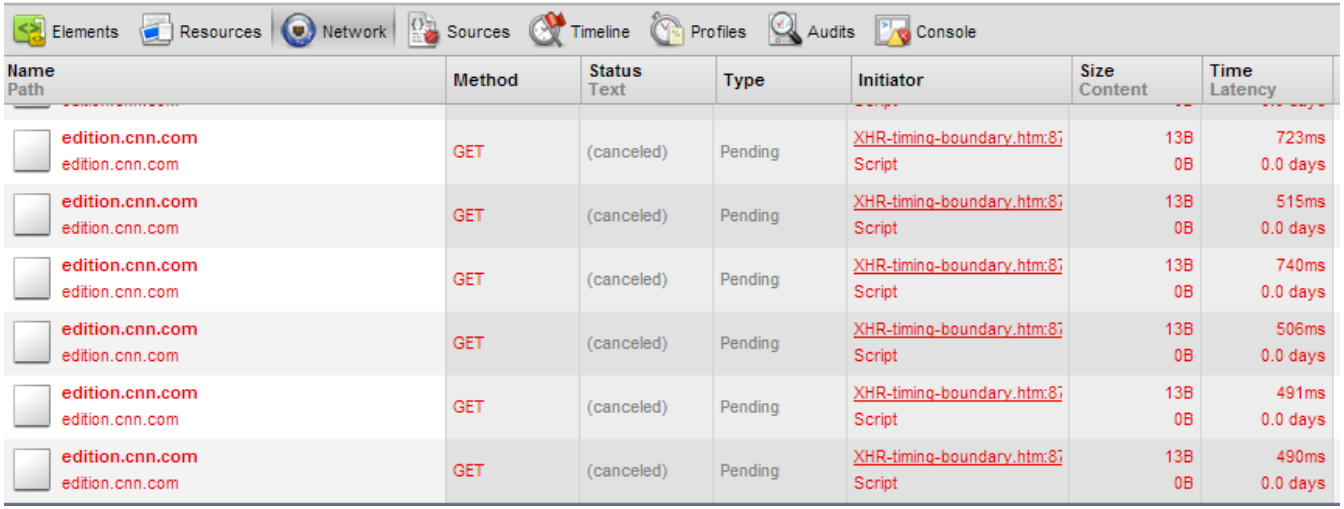

Name / Path	Method	Status / Text	Type	Initiator	Size / Content	Time / Latency
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	723ms 0.0 days
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	515ms 0.0 days
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	740ms 0.0 days
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	506ms 0.0 days
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	491ms 0.0 days
edition.cnn.com edition.cnn.com	GET	(canceled)	Pending	XHR-timing-boundary.htm:87 Script	13B 0B	490ms 0.0 days

Script results

2515,717
2514,512
2515,738
2514,504
2515,490
2514,490
Min first 468
Min Second 246

Attacking responses

Attacking response

- Detecting size – remains the same
- Generating requests – remains the same
- Main change
  - Attacker can only control the response indirectly
  - For example with the search functionality

Attack PoC

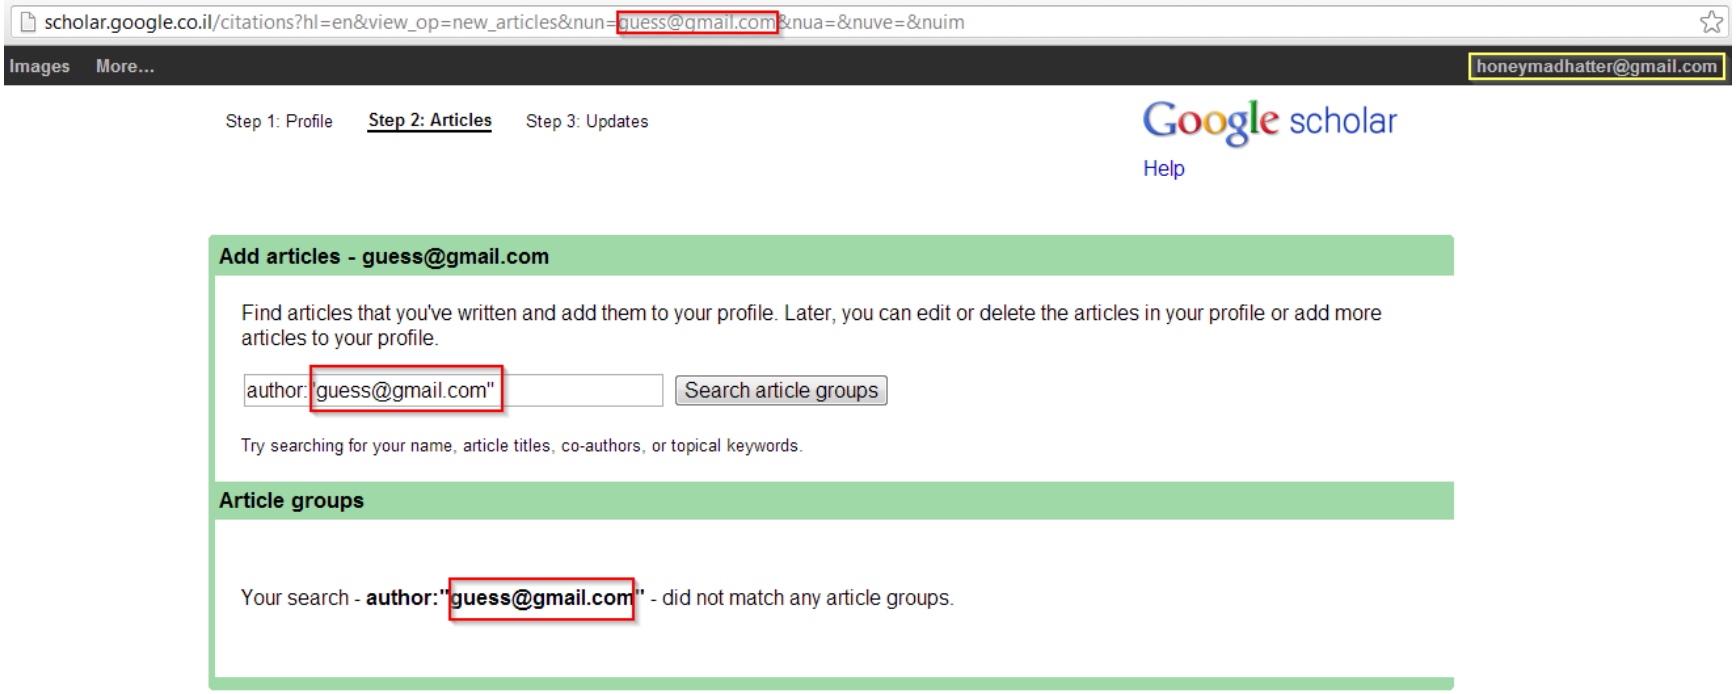

Attack PoC demo

HTTP Response Time Measurements

- Create HTTP request with iframe
  - iframe adhere to SOP
  - Doesn't allow parent to access the response content
  - Timing leaks the response size
- Use getTime() on iframe events
  - onLoad
  - Onreadystatechange (IE)
- Noise elimination – as before

HTTP Response Time Measurements

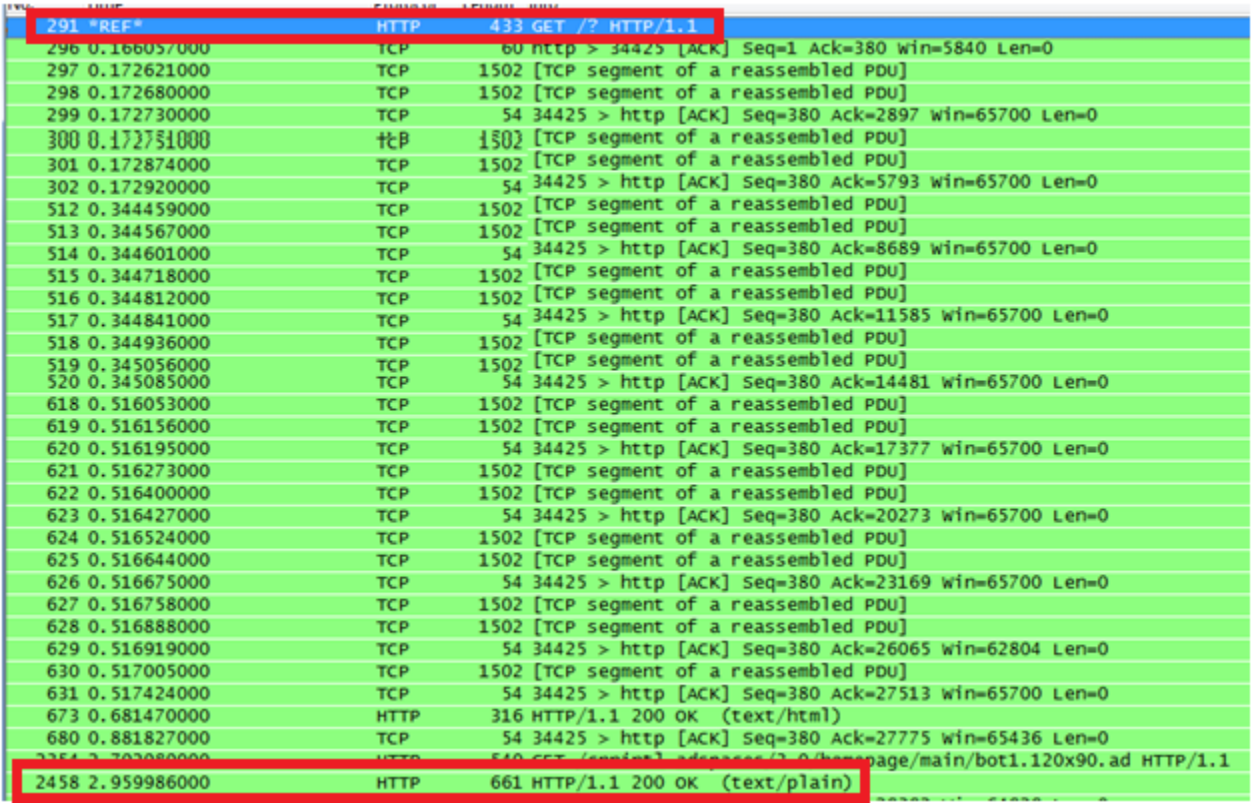

291	*REF*	HTTP	433	GET /? HTTP/1.1
296	0.166057000	TCP	60	http > 34425 [ACK] Seq=1 Ack=380 Win=5840 Len=0
297	0.172621000	TCP	1502	[TCP segment of a reassembled PDU]
298	0.172680000	TCP	1502	[TCP segment of a reassembled PDU]
299	0.172730000	TCP	54	34425 > http [ACK] Seq=380 Ack=2897 Win=65700 Len=0
300	0.172751000	TCP	1502	[TCP segment of a reassembled PDU]
301	0.172874000	TCP	1502	[TCP segment of a reassembled PDU]
302	0.172920000	TCP	54	34425 > http [ACK] Seq=380 Ack=5793 Win=65700 Len=0
512	0.344459000	TCP	1502	[TCP segment of a reassembled PDU]
513	0.344567000	TCP	1502	[TCP segment of a reassembled PDU]
514	0.344601000	TCP	54	34425 > http [ACK] Seq=380 Ack=8689 Win=65700 Len=0
515	0.344718000	TCP	1502	[TCP segment of a reassembled PDU]
516	0.344812000	TCP	1502	[TCP segment of a reassembled PDU]
517	0.344841000	TCP	54	34425 > http [ACK] Seq=380 Ack=11585 Win=65700 Len=0
518	0.344936000	TCP	1502	[TCP segment of a reassembled PDU]
519	0.345056000	TCP	1502	[TCP segment of a reassembled PDU]
520	0.345085000	TCP	54	34425 > http [ACK] Seq=380 Ack=14481 Win=65700 Len=0
618	0.516053000	TCP	1502	[TCP segment of a reassembled PDU]
619	0.516156000	TCP	1502	[TCP segment of a reassembled PDU]
620	0.516195000	TCP	54	34425 > http [ACK] Seq=380 Ack=17377 Win=65700 Len=0
621	0.516273000	TCP	1502	[TCP segment of a reassembled PDU]
622	0.516400000	TCP	1502	[TCP segment of a reassembled PDU]
623	0.516427000	TCP	54	34425 > http [ACK] Seq=380 Ack=20273 Win=65700 Len=0
624	0.516524000	TCP	1502	[TCP segment of a reassembled PDU]
625	0.516644000	TCP	1502	[TCP segment of a reassembled PDU]
626	0.516675000	TCP	54	34425 > http [ACK] Seq=380 Ack=23169 Win=65700 Len=0
627	0.516758000	TCP	1502	[TCP segment of a reassembled PDU]
628	0.516888000	TCP	1502	[TCP segment of a reassembled PDU]
629	0.516919000	TCP	54	34425 > http [ACK] Seq=380 Ack=26065 Win=62804 Len=0
630	0.517005000	TCP	1502	[TCP segment of a reassembled PDU]
631	0.517424000	TCP	54	34425 > http [ACK] Seq=380 Ack=27513 Win=65700 Len=0
673	0.681470000	HTTP	316	HTTP/1.1 200 OK (text/html)
680	0.881827000	TCP	54	34425 > http [ACK] Seq=380 Ack=27775 Win=65436 Len=0
2458	2.959986000	HTTP	661	HTTP/1.1 200 OK (text/plain)

```
LOG: abs 1353340607380
LOG: interactive abs 1353340607771 ref 391
SCRIPT1010: Expected identifier
   ?, line 418 character 146
LOG: complete abs 1353340610285 ref 2905
```

Candidate?

- Get the Twitter username of a logged in user

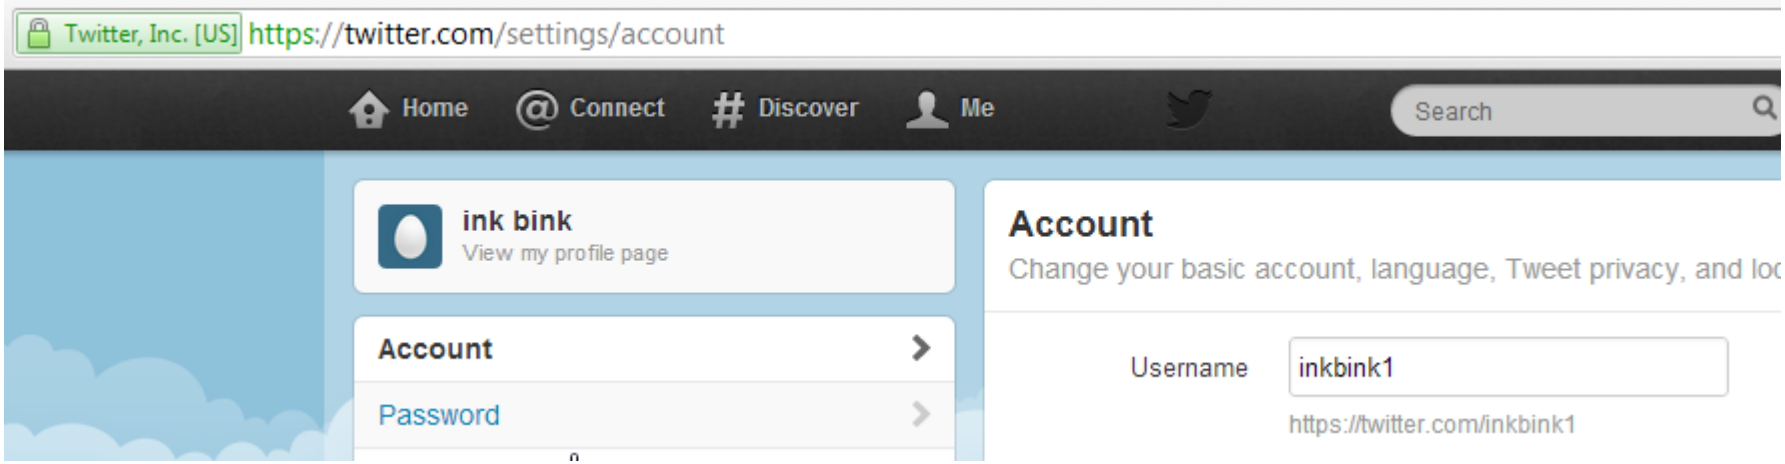


```
<style id="user-style-inkbink1-bg-img" class="js-user-style-bg-img">
  body.user-style-inkbink1 {
          background-image: url(https://twimg0-a.akamaihd.net/images/themes/theme1/bg.png);
  }
</style>
```

Candidate?