


Telescope

Peering Into the Depths of TLS Traffic in Real-Time

Caragea Radu

May 26, 2016

About me

- 
- Bitdefender - Vulnerability Researcher
 - PwnThyBytes: CTF team captain

This presentation

Is not about:

- Faults in the TLS protocol/implementation
- Attacks on the crypto level of TLS

This presentation

Is not about:

- Faults in the TLS protocol/implementation
- Attacks on the crypto level of TLS

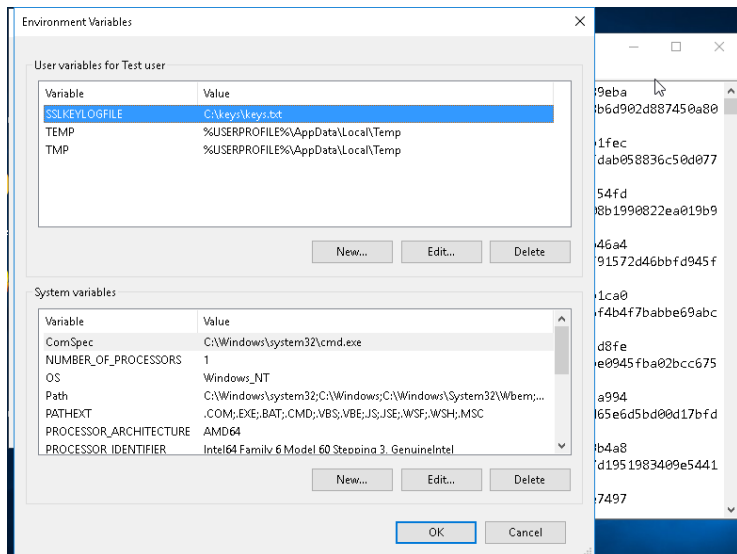
Is about:

- An experiment into what can be done from the hypervisor (powerful adversary; consider cloud providers)
- How far you can go and how noticeable it would be to an end-user
- Some unexpected results


Project Genesis

- Operating a honeypot farm
 - weak root (!) credentials: do what you will
 - analyze traffic (unless it's encrypted)
- Malvertising
 - Automated in-browser crawlers
 - Scour the net with the hope of getting infected
 - Need the infection vector (when served through TLS)

Other solutions: SSLKEYLOGFILE



Other solutions: SSLKEYLOGFILE



Other solutions: SSLKEYLOGFILE

```
3 [SSL segment of a reassembled PDU]
3 [TCP segment of a reassembled PDU]
0000 02 00 00 e4 a5 a6 6c 50 4d ae a9 00 08 00 45 00 .....|P M....E.
0010 1c 7c 87 40 40 00 30 06 73 a9 8d 55 e3 76 05 c4 .|.@@.0. s..U.v..
0020 bd 02 01 bb 9c 08 d9 d0 e4 65 bc ca 96 2e 80 10 .....e.....
0030 00 47 50 01 00 00 01 01 08 0a 0d eb 51 75 05 67 ..GP.....Qu.g
0040 74 c0 17 03 01 01 40 2a be 8f 70 1d b7 44 05 3e t....@* ..p.D.>
0050 f6 2a cf ae e9 43 a9 63 04 9c 18 ce 83 52 0f 9f .*...C.c....R..
0060 6d 1a 84 b3 2a df ea a1 c5 ac cf d0 35 36 11 83 m....*....56..
0070 11 23 4c e7 81 cd c0 10 08 5a d2 d3 88 b0 67 11 .#L.....Z....g.
0080 f0 9b 3e ae 41 e1 f8 fe 3a 9f 97 f7 c1 5a 9d 55 ...>.A... :...Z.U
0090 89 bd ed 88 ae 75 96 f3 a0 4a 04 1e de 9c 65 dd .....u...J....e.
00a0 36 85 b5 82 9f 02 05 22 bd 62 cc b2 99 7e 20 ff 6....." .b....~ .
00b0 97 08 01 f2 3c 45 3b b9 d7 58 b9 69 8c 49 ff 7d ....<E; .X.i.I.}
00c0 09 af 51 a4 f3 f9 01 47 cc 56 74 f8 63 ad c6 11 ..Q....G .Vt.c...
00d0 c2 7c 3f 9a 81 1e 71 d4 d6 94 0e 3f 4d 4c 71 75 .|?...q. ...?MLqu
00e0 d7 40 be aa 9e ac 9d 8d 81 e6 e6 30 35 b7 fc cf .@.....05...
00f0 25 46 1f db 9c 35 2f 73 af 04 73 12 e6 92 5e d0 %F...5/s ..s...^
0100 76 04 6f 49 ba e5 0b f6 c3 f3 58 a5 43 be 76 20 v.oI.... .X.C.v
0110 d4 52 e0 10 84 ac aa f4 8f 43 78 81 2c a8 e5 4f .R.... .Cx.,.0
0120 96 bf e3 db 27 28 cb 77 26 50 36 56 bc d0 10 3b ....'(.w &P6V...;
0130 e7 0a 6d 1e f6 e3 7c 45 82 85 12 cc de 82 27 2b ..m...|E .....!+
```

Frame (7306 bytes)	Decrypted SSL data (296 bytes)
--------------------	--------------------------------

Other solutions: SSLKEYLOGFILE

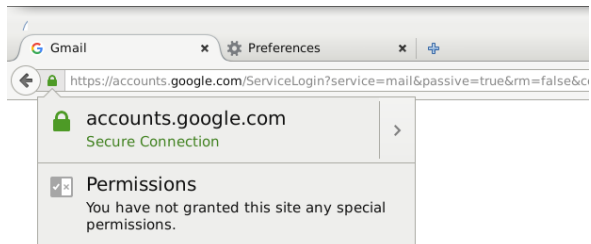
```
3 [SSL segment of a reassembled PDU]
3 [TCP segment of a reassembled PDU]
0000 48 54 54 50 2f 31 2e 31 20 32 30 30 20 4f 4b 0d HTTP/1.1 200 OK.
0010 0a 44 61 74 65 3a 20 46 72 69 2c 20 30 38 20 41 .Date: Fri, 08 A
0020 70 72 20 32 30 31 36 20 31 33 3a 33 35 3a 32 39 pr 2016 13:35:29
0030 20 47 4d 54 0d 0a 53 65 72 76 65 72 3a 20 41 70 GMT..Server: Ap
0040 61 63 68 65 2f 32 2e 32 2e 31 36 0d 0a 4c 61 73 ache/2.2 .16..Las
0050 74 2d 4d 6f 64 69 66 69 65 64 3a 20 54 68 75 2c t-Modified: Thu,
0060 20 33 30 20 41 70 72 20 32 30 31 35 20 31 30 3a 30 Apr 2015 10:
0070 30 35 3a 32 34 20 47 4d 54 0d 0a 45 54 61 67 3a 05:24 GM T..ETag:
0080 20 22 31 33 32 38 30 30 33 2d 39 32 65 39 61 2d "132800 3-92e9a-
0090 35 31 34 65 65 33 62 66 31 36 31 30 30 22 0d 0a 514ee3bf 16100"..
00a0 41 63 63 65 70 74 2d 52 61 6e 67 65 73 3a 20 62 Accept-R anges: b
00b0 79 74 65 73 0d 0a 43 6f 6e 74 65 6e 74 2d 4c 65 ytes..Co ntent-Le
00c0 6e 67 74 68 3a 20 36 30 31 37 35 34 0d 0a 4b 65 ngth: 60 1754..Ke
00d0 65 70 2d 41 6c 69 76 65 3a 20 74 69 6d 65 6f 75 ep-Alive : timeou
00e0 74 3d 31 35 2c 20 6d 61 78 3d 31 30 30 0d 0a 43 t=15, ma x=100..C
00f0 6f 6e 6e 65 63 74 69 6f 6e 3a 20 4b 65 65 70 2d onnectio n: Keep-
0100 41 6c 69 76 65 0d 0a 43 6f 6e 74 65 6e 74 2d 54 Alive..C ontent-T
0110 79 70 65 3a 20 61 70 70 6c 69 63 61 74 69 6f 6e ype: app lication
0120 2f 70 64 66 0d 0a 0d 0a /pdf....
```

Frame (7306 bytes) Decrypted SSL data (296 bytes)


Other solutions: SSLKEYLOGFILE

- Implemented in libNSS (firefox) and openssl (chrome) but not IE/Edge
- The downside: it's blatantly visible

Other solutions: Custom Root CA



Other solutions: Custom Root CA



Other solutions: Custom Root CA

- Typical solution used by AVs/proxies to intercept TLS traffic
- Visible by malware by scanning the disk.
- Moreover, in "Analyzing Forged SSL Certificates in the Wild" Huang et al show how to do this within the browser

Other solutions: PANDA keyfind plugin



534 lines (461 sloc) | 33.6 KB

Raw Blame History

Finding SSL/TLS Master Secrets with PANDA

Introduction

Monitoring SSL/TLS-encrypted traffic is a classic problem for intrusion detection systems. Currently, hypervisor- or network-based IDSes that wish to analyze encrypted traffic must perform a man-in-the-middle attack on the connection, presenting a false server certificate to the client. Not only does this require the client to cooperate by trusting certificates signed by the intrusion detection system, it also takes control of the certificate verification process out of the hands of the client—a dangerous step, given that many existing SSL/TLS interception proxies have a history of certificate trust vulnerabilities.

Instead of a man-in-the-middle attack, we can instead attempt to locate the code that generates SSL/TLS master secret; this secret is sufficient to decrypt any encrypted traffic in a given session, giving us a "man-on-the-inside". Once we have identified the location of the code that generates this secret, we can hook it using any number of standard techniques in order to dump out the master secret. This secret can then be provided to an IDS to decrypt the content of the SSL stream; it may also be provided to a tool like Wireshark to decrypt packet captures after the fact (even if perfect forward secrecy is used).

- really cool solution!
- run the machine under QEMU
- use pre-established "trace points"
- trace memory writes

Other solutions: PANDA keyfind plugin

Place this output into a file named `keyfind_config.txt` in the `panda/qemu` directory. Alternatively, the same information can be derived by hand using a tool like Wireshark and copied into `keyfind_config.txt`, but this is rather more labor intensive.

Locating the Master Key Code

Finally, we can run a replay with the `keyfind` plugin enabled to find out what code generates the master secret. Because the `keyfind` plugin tracks the calling function in order to better identify different memory accesses, we also need to enable the `callstack_instr` plugin, which keeps track of function calls and returns. We'll also use QEMU's VNC output rather than the default SDL because replays don't show any GUI output.

Using `keyfind` can be quite slow! On my machine, this short session, which takes only 12 seconds to replay with no plugins, takes almost 2 hours to run with `keyfind` enabled. This is what the output looks like:

```
brendan@brendantemp:~/git/panda/qemu$ echo "begin_replay ssltut" | \  
x86_64-softmmu/qemu-system-x86_64 -hda debian_squeeze_i386_desktop_tut.qcow2 \  
-m 256 -monitor stdio -vnc :0 -net nic,model=e1000 -net user \  
-panda "callstack_instr;keyfind"  
Initializing plugin callstack_instr  
Initializing plugin keyfind  
Couldn't open keyfind_candidates.txt; no key tap candidates defined.
```

- huge overhead
- non portable

Scavenging the keys from memory?

While connection is still active, the keys must still be in memory.

Scavenging the keys from memory?

While connection is still active, the keys must still be in memory.

- Problem 1: exact key location is unknown
 - need to dump all memory
 - dumping memory takes time (>10 seconds for 4 GB)
 - multiple connections occur one after the other or interspersed
 - space quickly fills up

Scavenging the keys from memory?

While connection is still active, the keys must still be in memory.

- Problem 1: exact key location is unknown
 - need to dump all memory
 - dumping memory takes time (>10 seconds for 4 GB)
 - multiple connections occur one after the other or interspersed
 - space quickly fills up
- Problem 2: we don't even know how to distinguish the correct keys from random memory.

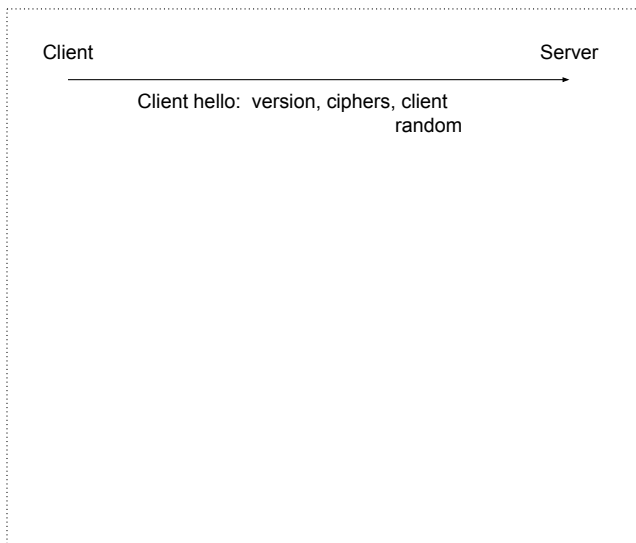
Scavenging the keys from memory?

While connection is still active, the keys must still be in memory.

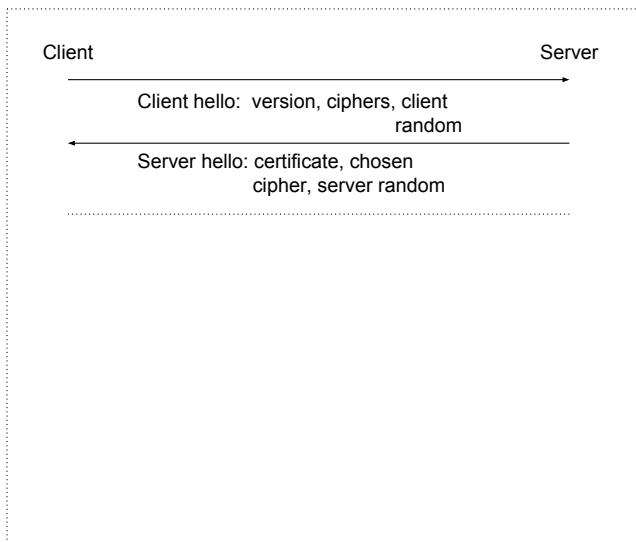
- Problem 1: exact key location is unknown
 - need to dump all memory
 - dumping memory takes time (>10 seconds for 4 GB)
 - multiple connections occur one after the other or interspersed
 - space quickly fills up
- Problem 2: we don't even know how to distinguish the correct keys from random memory.

To understand our approach we must dig deeper!

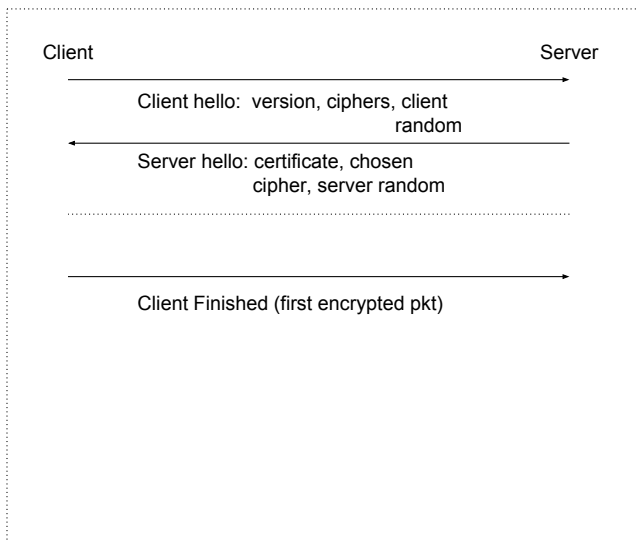
How exactly does TLS work? Client Hello



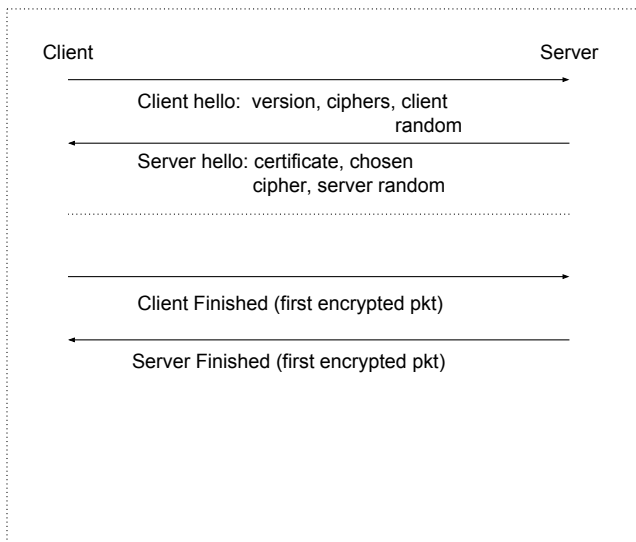
How exactly does TLS work? Server Hello



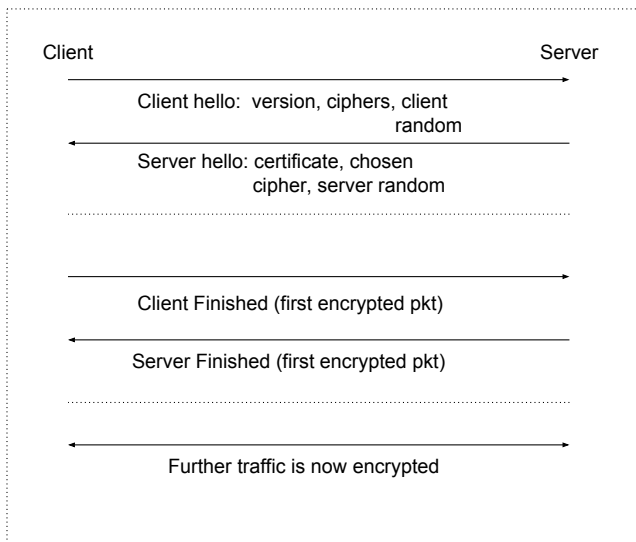
How exactly does TLS work? Client Finished



How exactly does TLS work? Server Finished



How exactly does TLS work? Handshake Complete



But wait!

Excerpt from RFC5246/4346:

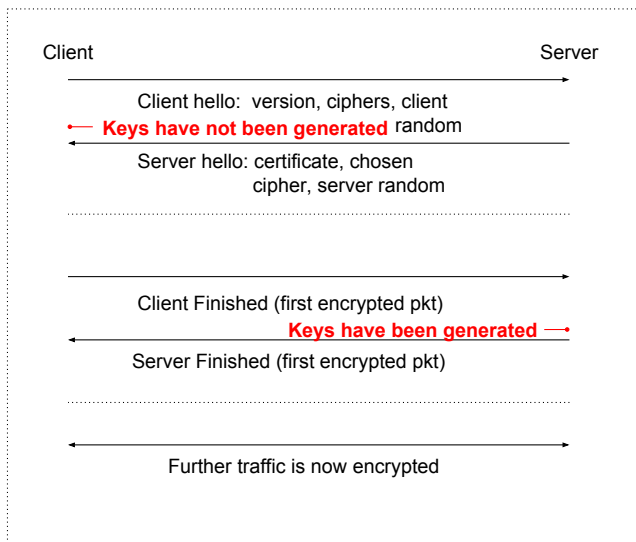
To generate the key material, compute

```
key_block = PRF(SecurityParameters.master_secret,  
                "key expansion",  
                SecurityParameters.server_random +  
                SecurityParameters.client_random);
```

until enough output has been generated.

```
Key material = [client_write_MAC_key] [server_write_MAC_key]  
               [client_write_key] [server_write_key]  
               [client_write_IV] [server_write_IV]
```

Key events in the TLS handshake



Implications

- Only track memory between events
- Events signalled by passing through netfilter queue
- Dramatic decrease in memdump size

Implications

- Only track memory between events
- Events signalled by passing through netfilter queue
- Dramatic decrease in memdump size
- But how do you actually "track" pages?

Think VM live migration

Think VM live migration

Logdirty mechanism

- t_0 : start tracking pages written to from t_0 and flush the RAM to the target on the network

Think VM live migration

Logdirty mechanism

- t_0 : start tracking pages written to from t_0 and flush the RAM to the target on the network
- when this finishes, get the "dirty" pages (at t_1) and send the delta to target again

Think VM live migration

Logdirty mechanism

- t_0 : start tracking pages written to from t_0 and flush the RAM to the target on the network
- when this finishes, get the "dirty" pages (at t_1) and send the delta to target again
- repeat this for every $t_i - t_{i+1}$ until number of pages is under threshold

Think VM live migration

Logdirty mechanism

- t_0 : start tracking pages written to from t_0 and flush the RAM to the target on the network
- when this finishes, get the "dirty" pages (at t_1) and send the delta to target again
- repeat this for every $t_i - t_{i+1}$ until number of pages is under threshold
- stop VM1, do iteration one last time, start VM2

Similar mechanisms exist in most (all) modern hypervisors that support VM migration

- Page fault based (basic)
- EPT A/D

Similar mechanisms exist in most (all) modern hypervisors that support VM migration

- Page fault based (basic)
- EPT A/D
- Recently, a processor extension especially for this: Intel PML.
Convenient, right ???

Putting the TeLeScope together

- Filter target network events and send to netfilter queue
- Start logging on Server Hello
- Stop logging and dump pages on Client Finished
- The result is a micro-memdump
- Can be processed **offline** anytime

TeLeScope results

- on Linux VM per connection: 500K - 10 MB memdump
- on Windows VM per connection: 15 MB - 60 MB memdump
- VM pause time: under 0.5 ms but on average 0.05 ms
- page dump time: 1-10 ms (disguised as packet delay)

Problems revisited

- Problem 1: you don't know where the keys are (**partially solved**)
 - ~~need to dump all memory~~
 - ~~dumping memory takes time (\rightarrow 10 seconds for 4 GB)~~
 - multiple connections occur one after the other or interspersed
 - space quickly fills up
- Problem 2: we don't even know how to distinguish the correct keys from random memory.

Problem 2: key discerning


Apparently multiple unknowns

- key format
- key parameters: IV, nonce, etc
- what to encrypt/decrypt
- what it decrypts to

Known Plaintext Attack

The Client/Server Finished messages have a fixed form:

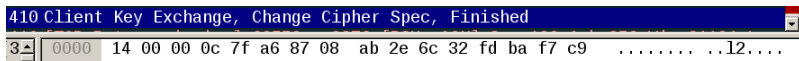
14 00 00 0C [12 random bytes]



Known Plaintext Attack

The Client/Server Finished messages have a fixed form:

14 00 00 0C [12 random bytes]




A screenshot of a Wireshark packet capture. The packet list pane shows a packet of type "410 Client Key Exchange, Change Cipher Spec, Finished". The packet bytes pane shows the hex data "14 00 00 0c 7f a6 87 08 ab 2e 6c 32 fd ba f7 c9" followed by an ellipsis and "...12...". The first four bytes of the hex data are highlighted in grey, indicating they are the known plaintext.

- $\frac{1}{2^{32}}$ chance of a False Positive
- This works if you can decrypt the first 4 bytes (think stream ciphers, AES/CTR, etc)

Alexa top 1000 ciphers


- 5% RC4
- 21% AES CBC
- 73% AES GCM

AES CBC



- Decrypting each block depends on having the previous block
- For the first block you need the IV (not explicit for TLS 1.0)
- The known plaintext is exactly in the first block

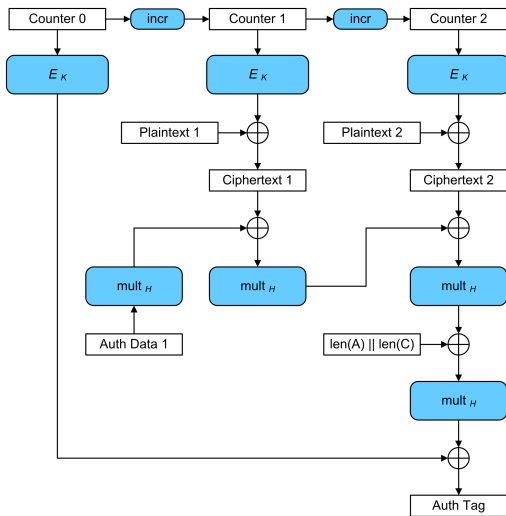
AES CBC




- Decrypting each block depends on having the previous block
- For the first block you need the IV (not explicit for TLS 1.0)
- The known plaintext is exactly in the first block
- We use the last block for the padding

```
0000  14 00 00 0c 7f a6 87 08 ab 2e 6c 32 fd ba f7 c9 .....l2....
0010  9c c5 76 20 da 83 6e b5 27 af ac ce e0 ac 1a e4 ..v ..n.'.....
0020  8c 86 55 fc 0b 0b 0b 0b 0b 0b 0b 0b 0b 0b 0b 0b ..U.....
```

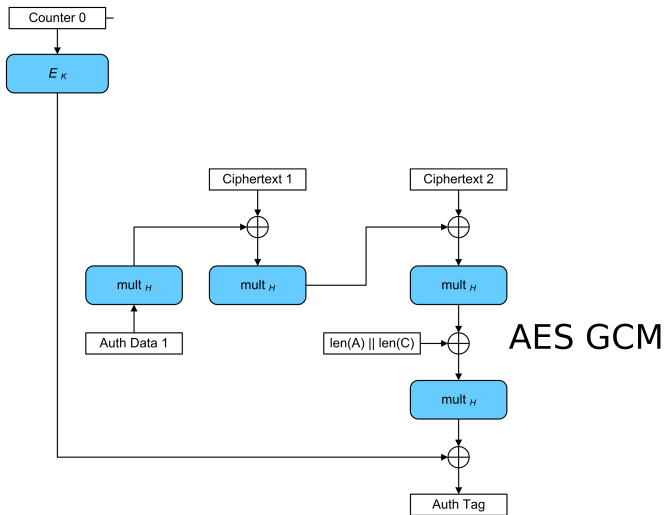
AES GCM/CTR



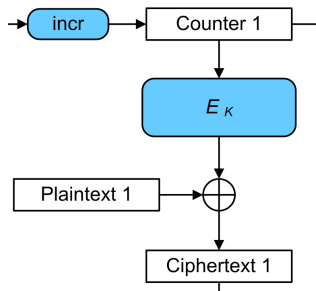
AES GCM/CTR Encryption



AES GCM/CTR Authentication

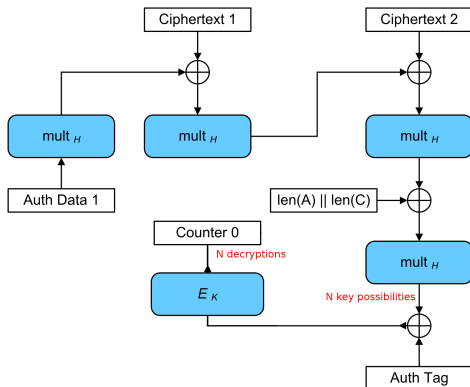


AES GCM bruteforce attempt



- Counter is [8 bytes key material][8 bytes counter]
- Here: 0fddf45e89838e700000000000000001
- The first half is also from the key material
- Implies $O(N^2)$ which we don't like!

AES GCM tag



- Auth Data, the Ciphertexts, Lengths and the Tag Known
- Each key K corresponds to one H
- Reverse the Flow and do one extra decryption
- Known Plaintext Attack on Counter format

Telesync benchmarks

- 1 thread, 240 MB completely random data: 6750 ms
- 6 threads, 240 MB completely random data: 557 ms (12x speedup)
- 6 threads + heuristics, 240 MB typical memdump data: 151 ms (44x speedup)

Telesync benchmarks

- 1 thread, 240 MB completely random data: 6750 ms
- 6 threads, 240 MB completely random data: 557 ms (12x speedup)
- 6 threads + heuristics, 240 MB typical memdump data: 151 ms (44x speedup)

However, a typical memdump size is usually 5-10 times smaller. So we can usually consume as fast as we can produce!

Demo 1 (manual)

Demo 2 (integrated)

Technique genericity

Actually, this can be applied to other protocols that use a similar negotiation technique for the symmetric keys:

- VPN
- SSH
- Tor

Conclusions and more

- Decrypting TLS on current implementations is definitely feasible with a hypervisor-in-the-middle attack
- We developed a fast and efficient PoC
- * You might not observe if you are the one "under scrutiny" on a VPS
- * Actually, if you're not in control of the bare metal all bets are off

Questions

?