

# TLSkex: Harnessing virtual machine introspection for decrypting TLS communication

Benjamin Taubmann, Dominik Dusold, Christoph Frädlich, Hans P. Reiser  
Juniorprofessur für Sicherheit in Informationssystemen  
Universität Passau

31.03.2016



- Encrypted communication (especially TLS) became ubiquitous in the Internet (https, voip, mail, etc.)
- **Why do we need to decrypt encrypted communication?**
  - ▶ Adversaries use encryption to hide attacks
  - ▶ Network intrusion detection systems can not decrypt network traffic (especially Infrastructure-as-a-service based Clouds)
  - ▶ Distinction between legitimate and malicious traffic is hard
  - ▶ Malware analysis

- **Man-in-the-middle attack (MITM)**

- ▶ Force client not to encrypt: proxy replaces `https` with `http` in URLs of websites (`sslstrip` [4])
- ▶ TLS Proxy that uses a fake certificate (`sslsniff` [3], `sslsplit` [1])

- **Disadvantages:**


- ▶ can be detected
- ▶ does not work with certificate pinning
- ▶ (may) reduce the security level

- Decryption of the traffic with the private **RSA key** (ssldump [2])
  - ▶ Key acquisition from hard disk or main memory
  - ▶ Extraction of the RSA keys from main memory (simple pattern matching for the ASN.1 structure);
- **Disadvantages:**
  - ▶ Only applicable when the hard disk or main memory can be accessed
  - ▶ not feasible for malware analysis
  - ▶ Fails when DH or ECDH session keys are used

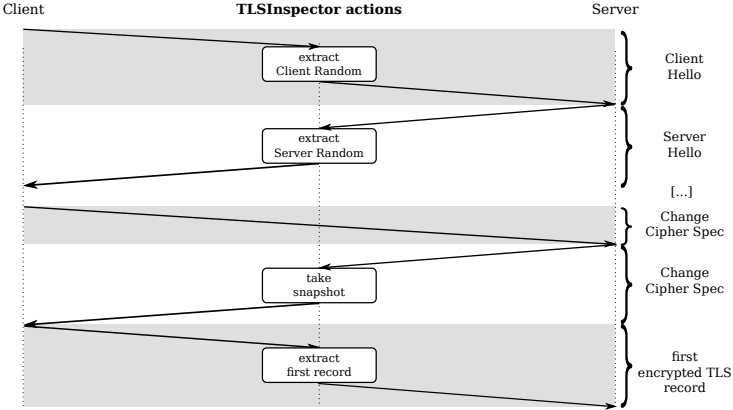
## Approach

- Monitor the network traffic of virtual machines (passively)
- Extract the *TLS session keys* from main memory of virtual machines
- Disadvantage: Works only when access to main memory is given (virtual machines, Firewire, ...)


# TLSkex Architecture



# Background: TLS Key Exchange



## (B) Network Monitoring

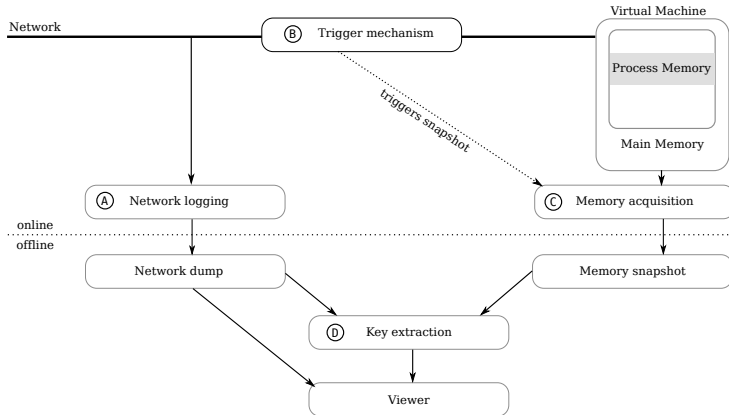


### Timing

- Snapshot must be taken when the key material is exchanged and before the connection is terminated
- When a communication partner has computed the key material it sends a *ChangeCipherSpec* message
- TLSInspector monitors the network traffic and triggers the snapshot when a *ChangeCipherSpec* message is sent



# TLSInspector Architecture




## Snapshot

- The memory of a virtual machine is accessed via libVMM
- Only the memory of the process that handles the connection
- Only *write-able* and *anonymous* pages
- Extract information from kernel memory (`task_struct`)

### **Save only modified pages**

- Set memory event on every page of a process when TCP connections is established
- Save dirty pages
- When hand shake is done, take snapshot of
  - ▶ dirty pages
  - ▶ newly allocated pages

# Background: TLS Key Exchange




# Evaluation - Snapshot time

Process	total	anon&writeable	new	modified	dumped	$t_{snap\_start}$	$t_{snap\_stop}$	$t_{search}$
Apache2	72090	3715	0	26	26	4.3 ms	4.4 ms	30 ms
Curl	38264	3438	15	19	34	3.3 ms	4.0 ms	2 ms
Wget	22813	1378	16	13	29	4.0 ms	3.5 ms	2 ms
s_client	6114	152	9	22	31	0.4 ms	0.6 ms	8 ms

**Table:** Amount of mapped and changed memory pages (4096 bytes) of different processes during the key negotiation procedure and the time to prepare ( $t_{snap\_start}$ ) and take ( $t_{snap\_stop}$ ) a differential snapshot;  $t_{search}$  denotes the time to extract a key from a snapshot

# TLSInspector Architecture



## (D) Key Extraction

### Problem

- There is no standardized way to store TLS session keys in memory

### Brute force approach

- Test every byte sequence as a key for the TLS connection
- Use message authentication code in TLS record to verify if key is corrected and data has been decrypted correctly
- **Advantage:** Implementation independent
- **Problem:** Attacker can fool monitoring tool, e.g., split key

## (D) Key Extraction Improvements

- Test only 4 byte aligned sequences
- Skip null byte areas
- Heuristics - Pretest keys
  - ▶ check if it is ASCII string
  - ▶ compare amount of zeroes and ones, should be uniformly distributed (k: deviation from expected mean)

$$\sum_{\mu-k}^{\mu+k} \binom{n}{k} p^k * (1-p)^{n-k} \geq 0.89$$

$$k = 16, \mu = 192, p = 0.5$$



Process	a						b	c	d
	k=1	k=2	k=4	k=8	k=16	k=32	no string	not all 0 / 1	combined (k=16)
key included	8.12	16.2	31.6	58.5	89.7	99.9	$1 - 10^{-15}$	$1 - 10^{-19}$	87.7
Apache2	0.10	0.28	0.64	1.27	2.33	4.26	85.49	43.54	1.69
Curl	0.15	0.45	1.04	2.11	3.50	4.75	77.53	10.55	3.32
Wget	0.15	0.46	1.06	2.15	3.60	4.91	78.10	10.68	3.38
s_client	0.054	0.18	0.49	0.96	1.89	3.40	56.52	37.35	1.63

**Table:** First row: probability that a key is not eliminated by the heuristic. Other rows: percentage of a memory snapshot that contains a 48 byte long and four byte aligned sequence with: a)  $192 \pm k$  one bits, b) the byte sequence is not an ASCII string c) no 8 byte sequence with only zero or only one bits d) a to c combined

- Easy to circumvent:
  - ▶ Start dedicated crypto process
  - ▶ Store key not in a 48 byte sequence
  - ▶ Mix byte order
  - ▶ Use different protocol or change TLS slightly
- Are kernel structures still trustworthy?
- Potential DoS vector?
- Ethical considerations

## **Non Intrusive**

- No active manipulation of communication
- No modification of application

## **Universal:**

- Independence of specific key exchange
- Independence of encryption algorithm
- Independence of client/server role
- Independence of the implementation

# Questions

?



**P. C. Heckel.**

**Use SSLsplit to transparently sniff TLS/SSL connections – including non-HTTP(S) protocols.**

<http://blog.philippeheckel.com/2013/08/04/use-sslsplit-to-transparently-sniff-tls-ssl-connections/>, 2013.

Accessed: 2015-08-19.



**S. Iveson.**

**Using ssldump to decode/decrypt SSL/TLS packets.**

<http://packetpushers.net/using-ssldump-decode-ssl-tls-packets/>, 2014.

Accessed: 2015-08-19.



**M. Marlinspike.**

**sslsniff.**

<http://www.thoughtcrime.org/software/sslsniff/>, 2002.

Accessed: 2015-08-19.



**M. Marlinspike.**

**sslstrip.**

<http://www.thoughtcrime.org/software/sslstrip/>, 2009.

Accessed: 2015-08-19.