



Attacking Highspeed Ethernet Links

Issued by:



ERNW GmbH
Enno Rey
Breslauer Strasse 28
D-69124 Heidelberg
www.ernw.de

Content

1	Introduction	3
1.1	Prerequisites	3
2	Tackling the Problem.....	4
2.1	Tool pcap_extractor	4
2.2	Hardware Needed	4
2.2.1	Identifying the bottleneck	4
2.2.2	Actual lab setup with locally installed hardware	4
2.2.3	Lab setup using the Amazon Elastic Compute Cloud (EC2)	4
3	Results.....	5
3.1	Description of overall test methodology.....	5
3.2	Test Results	5
3.2.1	Results from Tests in Amazon EC2 Cloud Setup	5
4	Conclusions.....	6
5	Appendix	7

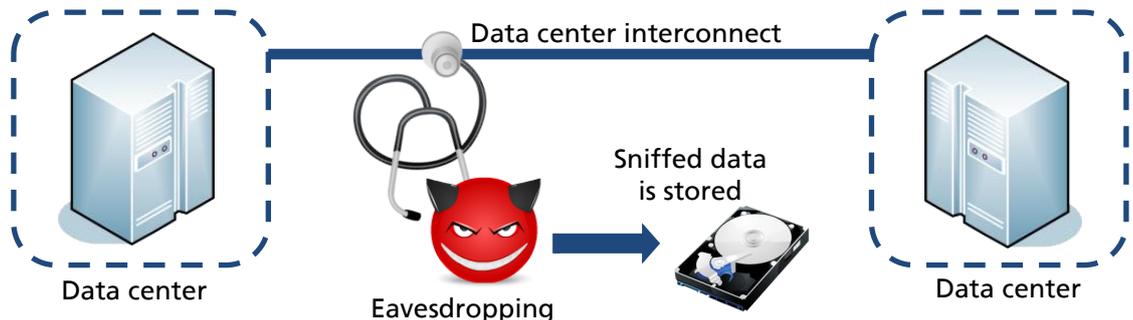
1 Introduction

There is a common misconception that the sheer amount of data coupled with multiplexed channels (e.g. WDM technology) makes successful eavesdropping attacks on high speed Ethernet – like 10 GbE or faster – unlikely. This is mainly based on the assumption that the amount of resources (e.g. RAM, [sufficiently fast] storage or CPU power) needed to process large files of captured data is a limiting factor. No practical evaluation of these assumptions has been performed so far though.

Therefore this paper aims to answer the following questions:

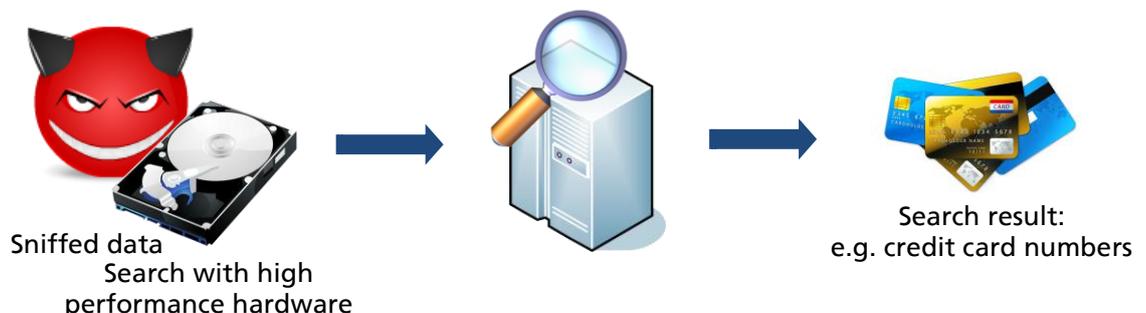
- Can the processing of large¹ amounts of captured data be done “in a feasible way”²?
- How much time and which type of hardware is needed to perform this task?
- Can this be done with readily available tools or is custom code helpful or even required? If so, how should that code operate?
- Can this task be facilitated by means of public cloud services?

To assess if the extraction of specific data from very large sets of captured network traffic can be done in a feasible way in terms of tools, hardware or time needed, a practical approach was undertaken. This included evaluating tools, building a lab with hardware fit for the purpose and performing a number of tests, including a demonstration of extracting credit card data from a file sized 500 GB and containing a virtual machine running an *Oracle* database. Such a file could, for example, derive from an eavesdropping attack against a high speed link connecting two data centers of an organization.



1.1 Prerequisites

It is assumed that the attacker has already gained access enabling her to eavesdrop on the high speed data link. A detailed description how this can be done can be found in [INFOGUARD_WP1]³. It is further assumed the attacker has already performed the packet collection in standard pcap⁴ format. The focus of the present paper is on the subsequent extraction of useful data from the dump file.



¹ For demonstration purposes we limited our research to files in the 500 gigabyte size range (equivalent to the full live migration of 16 virtual machines with 32 GB of virtual memory each) which equates to ~ 400 seconds of sniffing on a fully saturated 10 GbE link.

² Where “feasible” means within some hours and by means of hardware to be purchased for less than 5,000 CHF.

³ http://www.infoguard.com/docs/PDF/IG_Dokumente/WP_Fiber_Optic_Communication-e.pdf

⁴ Pcap API - Libcap (Unix systems), www.tcpdump.org.

2 Tackling the Problem

To execute the data extraction task in a feasible way the potential attacker has to dispose of the right tool(s) and hardware. This section lays out what is needed.

2.1 Tool pcap_extractor

As available command-line tools (like tethereal, tshark or tcpdump) have several limitations a special tool was developed. It's called *pcap_extractor*⁵ and is basically the fastest possible implementation of a pcap file reader. It opens a libpcap file handle for the designated input file, applies a libpcap filter to it and loops through all the filter matching packets, writing them to an output pcap file. Contrary to tcpdump and most other libpcap based analysis tools, it provides the possibility to search for a given string inside of the matching packets, for example a credit card number or a username in easy way.

A sample call to search a pcap file for iSCSI packets which contain a certain credit card number and write them to the output file is shown in the following:

```
# pcap_extractor -i input-file.pcap -o output-file.pcap -f "tcp port 3260" -s "5486123456789012"
```

2.2 Hardware Needed

2.2.1 Identifying the bottleneck

While measuring the performance of multiple pcap analysis tools the profiling of system calls indicated that the tools spend between 85% and 98%⁶ of the search time on waiting for I/O. This means that the fastest tool for 98% of the time doesn't do anything but waiting for dump data.

2.2.2 Actual lab setup with locally installed hardware

The lab system was designed to provide as much I/O bandwidth as possible and was composed of:

- Intel Core i7-990X Extreme Edition, 6x 3.46GHz
- 12GB (3 * 4GB) DDR3 1600MHz, PC3-12800
- ASRock X58 Extreme6 S1366 mainboard
- 4x Intel 510 Series Elm Crest SSD 250GB

The mainboard and the SSDs were chosen to support SATA3 with a theoretical maximal I/O bandwidth of 6 Gbps. FreeBSD was used as operating system.

2.2.3 Lab setup using the Amazon Elastic Compute Cloud (EC2)

The used Amazon EC2 instance was a so called *extra large instance*. Amazon describes this instance type as follows:

- 15 GB memory
- 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each)
- 1,690 GB instance storage
- 64-bit platform
- I/O Performance: High
- API name: m1.xlarge

1 EC2 Compute Unit "provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor"⁷, whereby the I/O performance scale of low, medium and high is not defined in more detail.

⁵ The tool can be downloaded from http://lernw.de/download/pcap_extractor.c.gz.

⁶ Depending on the tool analyzed.

⁷ As stated under <http://aws.amazon.com/ec2/instance-types/>.

3 Results

In the following the test methodology and results are presented.

3.1 Description of overall test methodology

For the tests, five capture files were created using the *mergcap* utility. Different sample traffic dumps⁸ were merged to five large files with different file sizes. All these files consist of several capture files containing a variety of protocols (including iSCSI and FCoE packets). More specifically capture files of ~40, ~80, ~200, ~500, and ~800 Gigabytes were created and were analyzed with different tools. At all tests the filtering expressions for *tcpdump* and *pcap_extractor* were configured to search for a specific source IP address and a specific destination IP address matching for some iSCSI packets inside of the capture file. The UNIX program *time* was used to measure the time of execution. Additionally the tools analyzing the data were started with the highest possible scheduling priority⁹ to ensure execution with the maximum of available resources.

3.2 Test Results

The following table shows the most results achieved with the *pcap_extractor* tool and having the collected data evenly spread across the four SSDs used in the local lab (see above).

Table 1: Results from test lab configured with four individual SSD drives and using custom tool

File size (GB)	Time of custom code (s)	Estimated time for 500 GB (s)	Time of tcpdump (s)	Estimated time for 500 GB (s)
44	113	1284	114	1295
80	203	1269	204	1275
204	521	1277	522	1279
404	1017	1259	1016	1257
516	1290	1250	1290	1250
808	2041	1263	2043	1264

This means that, in the demonstration setup, the search for specific credit card data could be performed in about 21 minutes, employing readily available tools and using common-of-the-shelf hardware for about 3500 CHF.

Furthermore the time needed scales in a linearly fashion with the file size, so that processing a 1 TB data volume presumably would have taken ~ 42 minutes, a 2 TB file would have taken ~ 84 minutes and so on. In addition, SSD prices are constantly declining, too.

3.2.1 Results from Tests in Amazon EC2 Cloud Setup

During the benchmarks, the performance was significantly lower than the one of the system described above even though eight different *Elastic Block Storage* (EBS) volumes were used to avoid the bottleneck of a single storage volume. Searching a 500 GB file with *pcap_extractor* took 4242 seconds (~ 71 minutes).

The overall performance of the test was apparently limited by I/O performance limitations within virtualized instances and virtualized storage systems.

⁸ <http://wiki.wireshark.org/SampleCaptures>

⁹ This was done using the *Unix nice* command.

4 Conclusions

It could be shown that the extraction of specific data from very large network traffic dumps can be achieved within a rather short time period. This was even possible using COTS hardware available for about 3500 CHF (as of March 2011). This means that an attacker disposing of (large) data sets resulting from previous eavesdropping attacks will most likely succeed in getting the exact data she's going after. In the lab setup, it took about 21 minutes to find a certain credit card number within a file sized 500 gigabyte which equates to a live migration of 16 virtual machines with 32 GB of virtual memory each.

The widespread perception that the sheer amount of data transferred over high speed network links prevents the extraction of data could thereby be proven wrong.

The entropy of a given file has practically no influence on the search/processing time needed. Furthermore it should be noted that the tests show a high correlation/degree of linearity between the size of the files (potentially derived by splitting one very large file into several chunks, each of them still of huge size) and the search time. So it seems there's a simple trade-off between the amount of storage provided and the search time. Given the ever-declining price of storage this will lead to even lower costs for such an attack in the future.

For all these reasons the protection of information, even when using optical networks, is vital and can be undertaken without restriction. InfoGuard is the leading manufacturer of layer 2 encryption solutions. These offer all-round and secure information exchange in MAN, WAN and SAN networks with 100% encryption throughput and minimum latency. Thus they are suitable for use in heavily utilized links and for time-critical applications. Data encryption is undertaken by the Advanced Encryption Standard (AES) with a key length of 256 bits.

5 Appendix

Demonstration of the Extraction of credit card data from an Oracle database running on virtual machine being live migrated

To show the impact of capturing and filtering data in business environments, an Oracle database was set up on a virtual machine. Figure 2 shows some sample content of the database, set up for simulating an environment storing sensible credit card data.

name	ccnum	valid	code
Enno Rey	1111222233334444	08/2012	123
Daniel Mende	5555666677778888	01/2012	456

Figure 1: Contents of sample Oracle database

At first a transfer of the whole virtual machine running the Oracle database was initiated. For simplicity reasons, FTP was used instead of VMotion. By default both approaches do not encrypt the transferred data.

This capture file was then analyzed with `pcap_extractor` to search for a specific string (for example the name of a credit card holder). This is shown in the following figure.

```
droelf# ./pcap_extractor -i 0/vm_transfer.pcap -o out_ennorey.pcap -f "host 192.168.0.13" -s "Enno Rey"
pcap_highspeed_extractor version 1 by Daniel Mende - dmende@ernw.de
Found 5 matching packets in 103 seconds.
droelf#
```

Figure 2: Console output of `pcap_extractor`

The resulting output file contains the packets with the data searched for, as shown below.

No.	Time	Source	Destination	Protocol	Info
1	10:31:14.431042	192.168.0.13	192.168.0.12	FTP-DA1	FTP Data: 1460 bytes
2	10:32:03.917055	192.168.0.13	192.168.0.12	FTP-DA1	[TCP Previous segment lost] FTP Data: 1460 bytes
3	10:40:30.813727	192.168.0.13	192.168.0.12	FTP-DA1	FTP Data: 1460 bytes
4	10:40:30.819194	192.168.0.13	192.168.0.12	FTP-DA1	[TCP Previous segment lost] FTP Data: 1460 bytes
5	11:12:47.929407	192.168.0.13	192.168.0.12	FTP-DA1	FTP Data: 1460 bytes

Offset	Bytes	Hex	ASCII
0440	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
0480	65 6e 64 65 10 35 35 35	35 36 36 36 36 37 37 37 Daniel M
0490	37 38 38 38 07 30 31 2f	32 30 31 32 03 c2 05	ende.555 5666777
04a0	39 2c 01 04 08 45 6e 6e	6f 20 52 65 79 10 31 31	9...Enno Rey.11
04b0	31 31 32 32 32 32 33 33	33 33 34 34 34 34 07 30	11222233 334444.0
04c0	38 2f 32 30 31 32 03 c2	02 18 02 06 e5 a8 06 a2	8/2012..
04d0	00 00 0c 02 00 01 d8 a8	0f 00 00 00 01 04 6e 51PQ
04e0	00 00 01 00 00 00 4e 23	01 00 d8 a8 0f 00 00 00N#
04f0	00 00 02 00 32 00 08 02	00 01 00 00 00 00 00 002.....
0500	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
0510	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
0520	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
0530	00 00 00 00 00 00 ff ff	0e 00 98 1f 8a 1f 8a 1f
0540	00 00 00 00 00 00 74 6c	e6 4d 74 6c e6 4d 00 00t1.Mt1.M.
0550	00 00 00 00 00 00 00 00	00 00 00 00 00 00 71 eeq.
0560	80 00 00 00 00 00 00 00	00 00 04 00 00 00 08 00
0570	00 00 74 ee 80 00 00 00	00 00 00 00 00 00 00 00t.....
0580	00 00 01 00 00 00 00 00	00 00 01 00 00 00 58 0aS.....P.....X.
0590	01 00 53 7e 0f 00 00 00	00 00 70 ee 80 00 08 00
05a0	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
05b0	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00
05c0	00 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00

Figure 3: Outfile of `pcap_extractor` in Wireshark