



## PASSIVE SCAN RESEARCH

for

**Stichting Internet Domein Registratie Nederland**

V1.1  
Amsterdam,  
November 8th, 2015

## Document Properties

Client	Stichting Internet Domein Registratie Nederland
Title	Passive Scan Research
Target	
Version	1.1
Researchers	Koen Jeukendrup, Matt Erasmus, Eireann Leverett
Authors	Peter Mosmans, Christine Rose, Matt Erasmus, Boi Sletterink
Reviewed by	Melanie Rieback
Approved by	Melanie Rieback

## Version control

Version	Date	Author	Description
0.1	November 3rd, 2015	Peter Mosmans	Initial draft
0.2	November 3rd, 2015	Peter Mosmans, Christine Rose	Editing content; formatting
0.3	November 3rd, 2015	Matt Erasmus, Boi Sletterink, Christine Rose	CVE content; appendix; images; conclusion
0.4	November 6th, 2015	Boi Sletterink, Peter Mosmans	Major rewrite of intro, results, and conclusion to better match work and findings. More clarification on methodology.
1.0	November 6th, 2015	Boi Sletterink, Peter Mosmans	Rewrites.
1.1	November 8th, 2015	Boi Sletterink, Peter Mosmans	Added future work, rewrite of introduction and conclusion.

## Contact

For more information about this Document and its contents please contact Radically Open Security BV.

Name	Melanie Rieback
Address	Radically Open Security BV

	Overdiemerweg 28, 1111 PP Diemen
Phone	+31 6 10 21 32 40
Email	melanie@radicallyopensecurity.com

## Table of Contents

1	Introduction .....	5
1.1	Scope .....	5
1.2	Assumptions .....	6
2	Methodology .....	7
2.1	Passive Scanning .....	7
2.2	IP Address Selection .....	7
2.3	Analysis Method .....	7
3	Mathematical Approach .....	9
4	Datasets and Tools .....	10
4.1	Scans.io .....	10
4.2	Shodan .....	11
4.3	Archive.org .....	12
4.4	Circl.lu Passive SSL Database .....	13
4.5	National Vulnerability Database .....	14
5	Results .....	15
5.1	Top Ten Scans.io .....	17
5.2	Top Ten Shodan .....	19
5.2.1	Differences Between Scans.io and Shodan .....	20
5.3	Heartbleed .....	21
5.4	HTTP Headers .....	21
5.5	SSL Research .....	23
6	Future Work .....	25
7	Conclusion .....	26
	Appendix 1 CVE Vulnerability Entries .....	28
	Appendix 2 PassiveScanning Tool .....	39
	App 2.1 Compilation .....	39
	App 2.2 Usage .....	39
	Appendix 3 Researcher Biographies .....	41

# 1 Introduction

Stichting Internet Domeinregistratie Nederland (hereafter “SIDN”) requested Radically Open Security B.V. (hereafter “ROS”) to perform a passive scan on websites hosted in the Netherlands.

The main purpose of the report is to introduce a new tool and methodology for passively surveying vulnerabilities on the Internet: Using passive, readily available data to perform an unobtrusive "scan".

In this report, we present the following:

- How this new tool works
- How the tool selects a sample for scanning
- What data sources we use

In addition, we also show the preliminary results from an example scan on web servers in the Dutch IP address space. These results showed that about 25% of hosts in our sample set are detected as containing vulnerabilities. The majority of HTTP hosts that responded with a banner (60%) contained vulnerabilities found in 2014 and earlier, for which updates have been available for quite some time now.

This allows us to conclude that overall, software is not regularly being updated, which results in these hosts being vulnerable to attackers.

Additionally, we conclude that certain HTTP headers, that actually can increase the overall security of a website were rarely used. For example, we found that only 1.6% of the web servers (22 out of 1380) sent X-Frame-Options headers.

## 1.1 Scope

The initial motivation behind this request was to evaluate the security of small- and medium enterprises (SME, Dutch: MKB) webshops in the .NL domain. The research aimed to obtain an overall picture of this group's infrastructure's safety situation. This group is interesting because it is a large group with limited resources to spend on IT security, and little knowledge and experience to organize this in a professional manner.

This scope was quickly extended to "any web server in the .NL domain" since there is currently no realistic way to determine if a web server is from a SME or another entity (e.g. large corporation, government, or consumer). This actually made it easier to implement.

The scope extended even further to include findings on other services besides web servers, e.g. MySQL and FTP. Although the tool now includes these protocols, its current focus is still web servers. This could be changed in future versions (see the [Future Work](#) (page 25) chapter).

## 1.2 Assumptions

As we only used passive scanning techniques, we made a number of assumptions. It is impossible to verify the results without active 'interrogation' techniques, or without actively trying to exploit the vulnerability.

- We assume that all banners returned by the server are accurate.

Server administrators can turn off banners, or alter the banner returned by a server, to obfuscate or mislead scanners. In practice, not all detected versions will be accurate.

- We assume that the detected software has a vulnerable configuration.

Sometimes vulnerabilities are only applicable when a certain configuration condition are met. Not all servers will have such a configuration. Some software does not update version numbers when security updates are applied (e.g. Microsoft IIS, some Linux packages).

- We assume that our sample size is sufficient to portray an accurate picture of the whole population.

Using sampling is prone to accuracy issues, so we selected a sample size that we think is sufficient to get results that are representative of the entire population.

## 2 Methodology

### 2.1 *Passive Scanning*

The research has been completely done on passive datasets. No single IP address has been actively scanned.

The reason behind performing a passive scan rather than an active scan is two-fold:

1. Active scanning is illegal without a waiver, and we're just scanning random hosts.
2. Hackers can find vulnerability information about targets without sending a single packet, which makes a passive scan risk-free and undetectable.

So even if the method is called passive "scanning", we didn't scan anything by ourselves; we used existing data sources that already scan the entire internet on a regular basis.

### 2.2 *IP Address Selection*

The first step for the research was obtaining a list of all IP addresses geographically located in The Netherlands. This list was downloaded from the Nirsoft Country IP Database website (<http://www.nirsoft.net/countryip/nl.html>) and contained 43,926,528 unique IP addresses.

This full list of Dutch IP addresses include devices from anything from small to medium and large enterprises, governmental organizations, NGO's, to consumer IP addresses; there is currently no simple and reliable method to distinguish between them. The devices behind each IP address can be anything from webshops, mail servers, other business-related servers, but also network devices such as (home- and business-) routers and firewalls, SCADA equipment, webcams, home automation controllers and more.

Next we looked how many of these IP addresses ran a webserver that responded with a version number. 551,768 Hosts (1,26 %) met this criteria. All statistics presented in this report are based purely on these servers. Servers that did not run a webserver responding with a version number have not been considered in the data below.

### 2.3 *Analysis Method*

For the purpose of this analysis, we used different data sources to find out security-relevant properties for HTTP, FTP, SSH, and SMTP. These services are very commonly available. Services

like IMAP and POP3, which are about equally common as SMTP, were not considered in the analysis because these services rarely announced their version number.

Next, we ran our PassiveNLSurvey program on the list of all 43M Dutch IP addresses. The tool combines passive data from the sources listed in the chapter [Datasets and Tools](#) (page 10), filters the IP addresses for hosts that run a webserver that reports a version as described above, and collects and compiles everything it can find across these sources. The output of the program shows which vulnerabilities have been detected.

Unfortunately, querying these data sources for all properties on all 43M addresses would take prohibitively long. This depends mostly on the speed of the servers answering the queries. It may also have been possible to optimize the queries, or run more queries in parallel, but the budget did not allow such optimization.

To get representative results in a reasonable time, we sampled the Dutch IP address space, performing detail queries for only a limited number of hosts. For each service, we calculated the appropriate sample size given the total number of host running this kind of service, using the formula described in the [Mathematical Approach chapter](#) (page 9).

Once the versions were found, we created a quantitative list of software versions used by the population. We compared these versions against other data sources, like Shodan and archive.org, which were used to try and find version numbers using comments in the source code.

Then, we conducted a simple SSL investigation using circl.lu to determine possible issues with the certificate. We combined this SSL research with the Heartbleed results obtained from Scans.io.

This list was then combined with a vulnerability database to determine how many of these servers contained vulnerabilities.



### 3 Mathematical Approach

From the total population of eligible 551,768 hosts as described in [Address Pool Selection](#) (page 7), we drew a random sample for each service from hosts that were running that service and reported version information. This sample size was determined using the estimation detection theory, which holds a low margin of error and a very high accuracy rate.

We used an sample size estimation method described in [Larry Green's course in statistics and sample size selection](#). The estimation theory calculation is formulated as the following:

$$n = p(1 - p) \left( \frac{Z_c}{E} \right)^2$$

Where in our case:

- $n$  is the number of samples needed for the required confidence.
- $p$  is the chance of a service being vulnerable. Since we don't know this in advance, we estimate that there is a 50% chance of a service being vulnerable.
- $Z_c$  is the cut-off value for the head and tail in the normal distribution, which represent the error margins (see also [this explanation of confidence intervals](#)). In our case, we accept a 5% error margin (see also  $E$ , below); in this case,  $Z_c$  needs to be 1.96.
- $E$  is the error margin. Since we accept a 5% error margin, we use 0.05 for  $E$ .

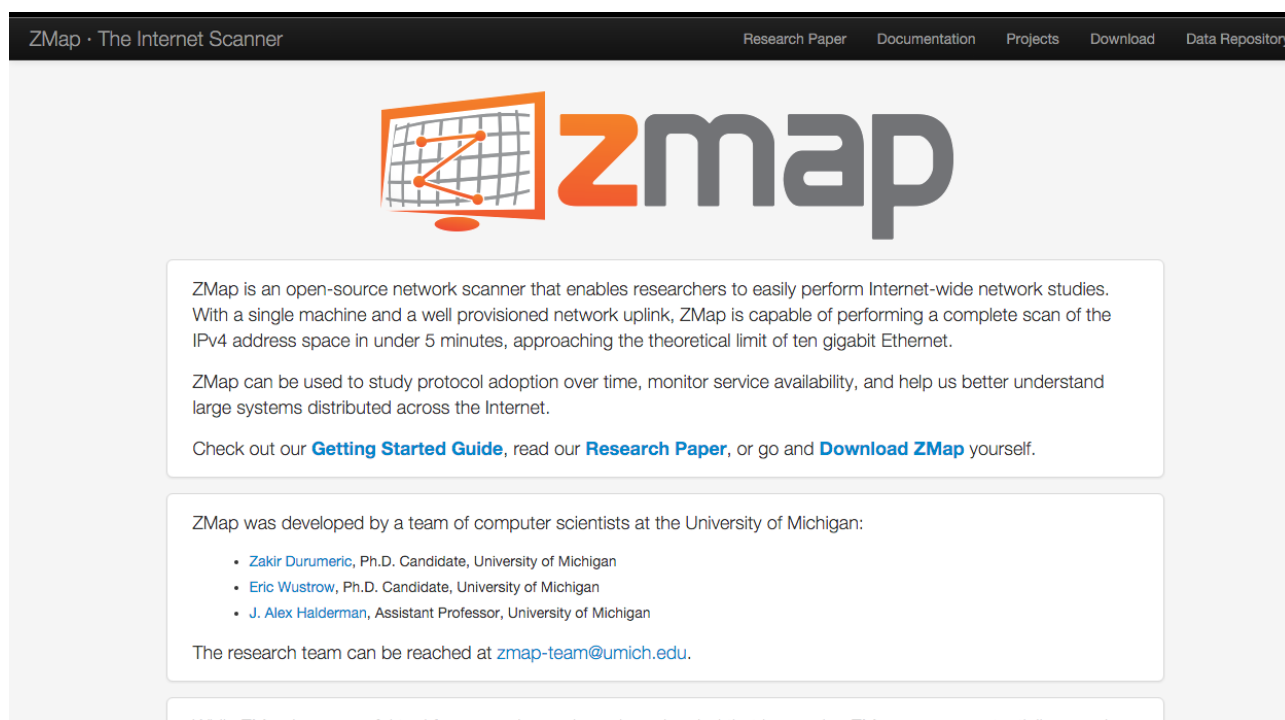
Each of these IP hosts run services. An example of a service would be a process that allows a browser like Google Chrome to view a webpage using the HTTP protocol. The random sample (1,380 hosts) included only hosts that ran an HTTP server. From there, we obtained "banner information", which is a way to gather information about a network system and services. That gave us further insight into the overall security of the web environment.

The sample size of the HTTP protocol was 1380 hosts, for SMTP 1011 hosts, for SSH 638 hosts, for MySQL 389 hosts and for FTP 152 hosts.

## 4 Datasets and Tools

We used the following tools and sources of data for this research.

### 4.1 Scans.io



Scans.io is a public archive with results of scans performed with ZMAP.

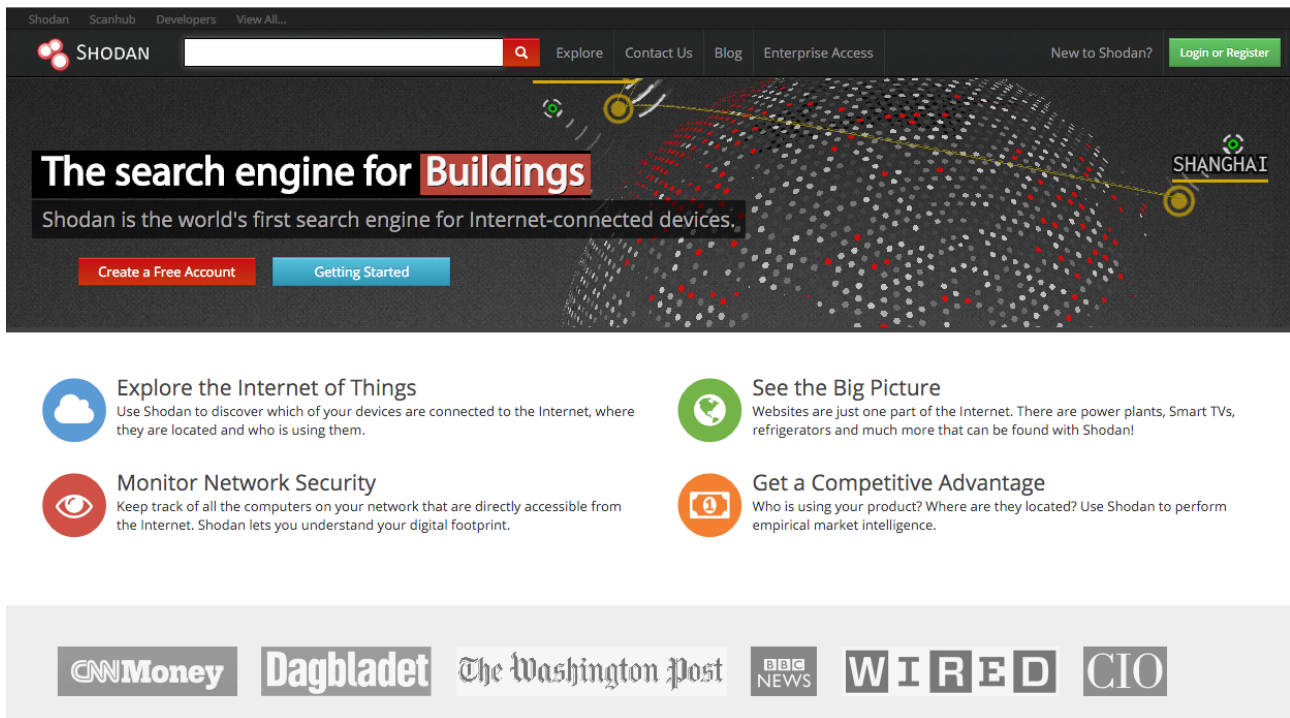
ZMAP is an open-source tool that enables researchers to perform a complete scan of the IPv4 address space with a single machine in mere minutes. Following the limitations of NMAP, scientists from the University of Michigan developed ZMAP. ZMAP keeps congestion of the host PC and the network to a minimum with state-less scanning and by bypassing the complications of the TCP stack. This allows it to scan large subnets within seconds and with minimal impact. To scan the complete Ipv4 address space with NMAP would be a massive undertaking. Plus, it would be outdated the moment the scan finished. A ZMAP scan can be done quickly and continuously. ZMAP is not only used to study protocol adoption over time, but it can also monitor service availability and help gain insight into large systems across the Internet.

The ZMAP team performs several scans for their own research. The results of these scans are published on <https://scans.io/> (Note: this site is moving to <https://www.censys.io/>). We used this scan database in our research to find past and current services, versions, and vulnerabilities.

Scans.io allows researchers to download the entire set of scan data. We used this to reduce query time, which was not just useful during the final passive scan on which the results below are based,

but especially during the testing. The dataset used during this scan was about 60 old, which we considered recent enough to be relevant and meaningful.

## 4.2 Shodan



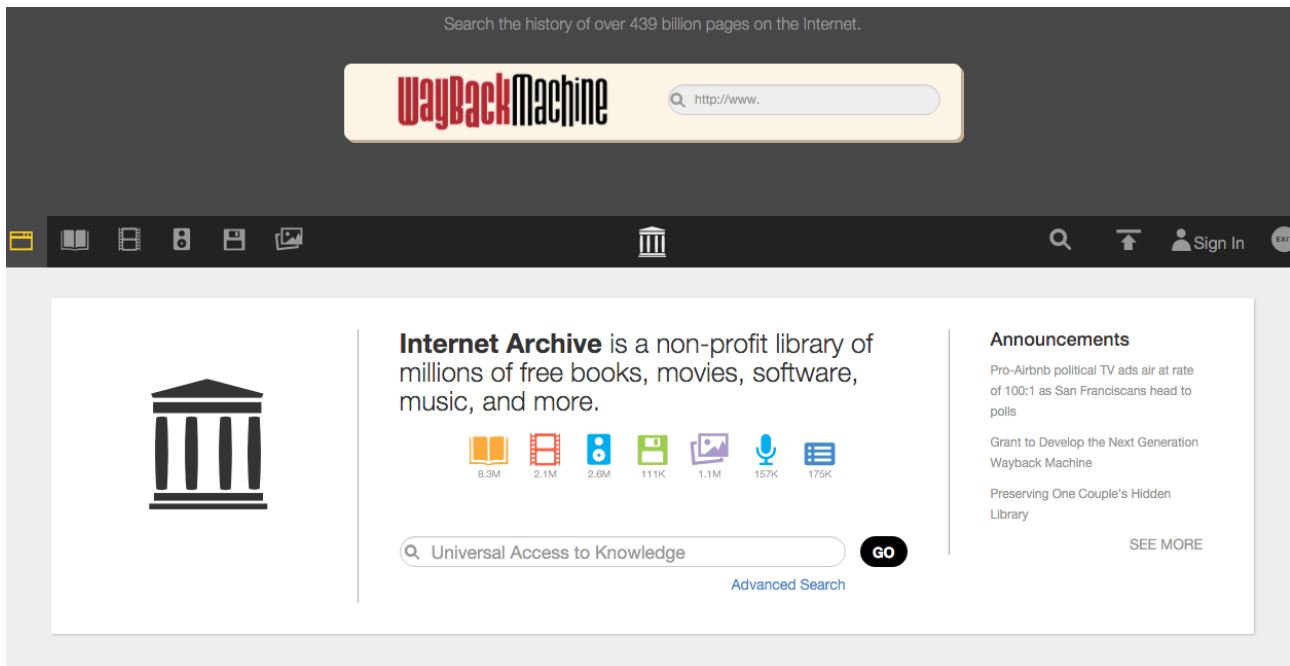
Shodan (<https://www.shodan.io/>) is a search engine that gathers data on computers (routers, servers, etc.) connected to the Internet. This includes web servers (HTTP port 80), FTP (21), SSH (22) Telnet (23), SNMP (161), and SIP (5060) services.

Shodan takes the following information into consideration:

- SSL
- FTP
- DNS
- Banners (HTTP, SMTP, etc.)
- Stack traces
- Open ports
- Debug and administration pages
- Public information

Shodan performs its own scans independently, and as a result, those scans were not initiated by ROS. Therefore, all data reported is historic data, but search results are available immediately. Shodan is constantly scanning the Internet, so the Shodan results are very recent.

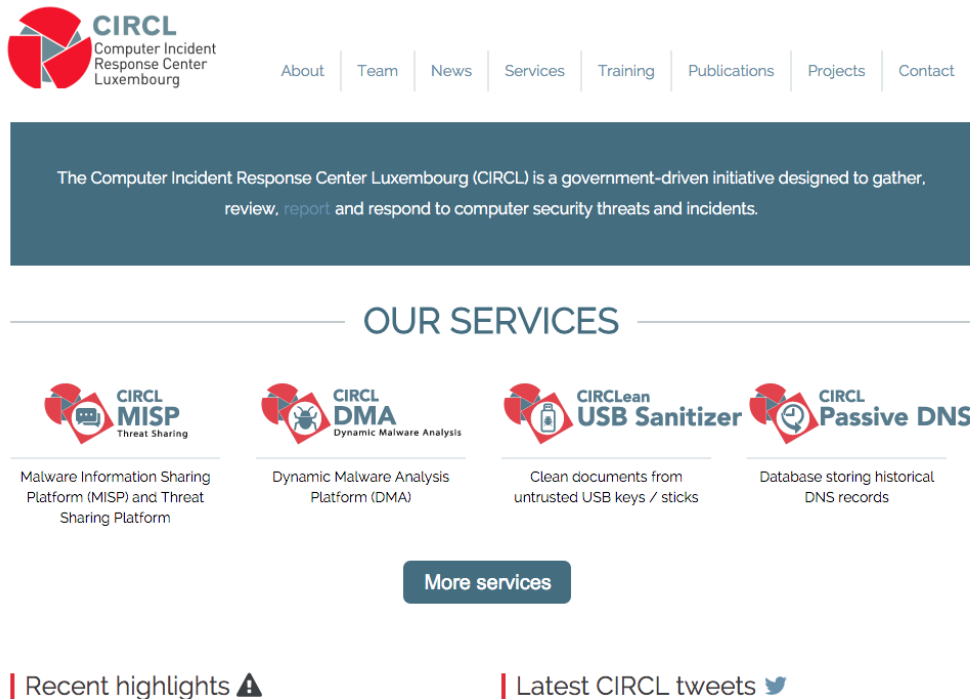
### 4.3 Archive.org



<https://archive.org/> (also called the Way Back Machine) archives websites. Sites are crawled, as is done by a regular search engine, but historic information is stored and older versions of websites can be re-visited using archive.org.

We looked at this service to find historic version information and known vulnerabilities for web-based services in the Netherlands.

## 4.4 Circl.lu Passive SSL Database




The screenshot displays the CIRCL (Computer Incident Response Center Luxembourg) website. The header features the CIRCL logo and a navigation menu with links: About, Team, News, Services, Training, Publications, Projects, and Contact. Below the header, a dark blue banner contains the text: "The Computer Incident Response Center Luxembourg (CIRCL) is a government-driven initiative designed to gather, review, report and respond to computer security threats and incidents." The main section is titled "OUR SERVICES" and lists four services, each with a red icon and a brief description:

- CIRCL MISP Threat Sharing**: Malware Information Sharing Platform (MISP) and Threat Sharing Platform
- CIRCL DMA Dynamic Malware Analysis**: Dynamic Malware Analysis Platform (DMA)
- CIRCLClean USB Sanitizer**: Clean documents from untrusted USB keys / sticks
- CIRCL Passive DNS**: Database storing historical DNS records


A "More services" button is located below the service list. At the bottom, there are two links: "Recent highlights" with a warning icon and "Latest CIRCL tweets" with a Twitter icon.

CIRCL Passive SSL (<https://circl.lu/services/passive-ssl/>) is a database that stores historical X.509 certificates seen per IP address. The Passive SSL historical data is indexed and searchable by IP address. This can be used to detect past and current insecure or revoked certificates, as well as changing certificates and apparent (legitimate or illegitimate) changes in the CA used by an organization.

## 4.5 National Vulnerability Database



Sponsored by  
DHS/NCCIC/US-CERT



NIST  
National Institute of  
Standards and Technology

### Mission and Overview

NVD is the U.S. government repository of standards based vulnerability management data. This data enables automation of vulnerability management, security measurement, and compliance (e.g. FISMA).

### Resource Status

**NVD contains:**

- 73386 [CVE Vulnerabilities](#)
- 321 [Checklists](#)
- 249 [US-CERT Alerts](#)
- 4393 [US-CERT Vuln Notes](#)
- 10286 [OVAL Queries](#)
- 106239 [CPE Names](#)

Last updated: 10/30/2015 5:41:22 PM

CVE Publication rate: 24.17


### Email List

NVD provides four mailing lists to the public. For information and subscription instructions

The US National Vulnerability Database (NVD, <https://nvd.nist.gov/download.cfm>) was used to provide more information on the CVEs reported by the search tools. This information also contains the CVSS (Common Vulnerability Scoring System) base score, which is an indicator of the level of threat towards the vulnerable system. The CVSS score was used to filter for the most severe threats found.


## 5 Results

In the passive scan, we investigated which vulnerabilities the 1,380 researched hosts contained. The next table represents the scan results categorized per protocol:



*Figure 1: Number of vulnerabilities in sample size*

A CVSS score above 5 indicates that those systems have an above average security risk. As shown on the table above, a very high percentage of the systems contain serious vulnerabilities. For the HTTP protocol, 30% of the sample size showed significant vulnerabilities: 408 out of 1380 tested. From the other protocols, only FTP fared much better (4% significant vulnerabilities), followed by MySQL (11%), SSH (22%) and SMTP (23%).



*Figure 2: Number of old and new vulnerabilities in sample size*

Figure 2 shows the distribution between "old" (pre-2015) and "new" (from 2015) CVE's. Outdated software can contain more than one known vulnerability. Generally speaking, the older the software, the more vulnerabilities it will contain, as new ones are discovered as time goes by. The older the software, the more susceptible it leaves the system - it is not only at risk from the new vulnerabilities and threats, but it is also at risk from more dated and widely known hacks.

Figure 2 shows how 60% of the HTTP services that responded with a banner are vulnerable only because of vulnerabilities older than 2015. For these vulnerabilities, an update or patch has usually been available for some time. SMTP and SSH show a similar but less dramatic picture.

Out of thousands of vulnerabilities found, the sections below show the top ten versions of vulnerabilities and their respective CVSSv2 ratings, categorized per dataset and protocol. The versions and the CVEs are independent.



## 5.1 Top Ten Scans.io

HTTP CVE Frequencies			HTTP Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
85	CVE-2010-3972	10	85	Microsoft-IIS/7.5	5 10
35	CVE-2010-0425	10	56	Apache/2.2.22	10 7.5
18	CVE-2008-0075	10	55	Apache/2	- -
3	CVE-2003-0224	10	46	Apache/2.4.7	9 6.8
85	CVE-2010-2730	9.3	37	Apache/2.2.15	24 7.8
21	CVE-2009-3023	9.3	36	Microsoft-IIS/8.0	- -
34	CVE-2008-1446	9	24	Microsoft-HTTPAPI/2.0	- -
116	CVE-2010-1256	8.5	23	nginx/1.4.6	- -
87	CVE-2011-3192	7.8	19	Microsoft-IIS/8.5	- -
21	CVE-2009-1535	7.6	18	Microsoft-IIS/6.0	15 10

table 2: Top ten vulnerabilities found while scanning HTTP using Scans.io

The following tables show the top ten for the FTP as well as the HTTP protocol.

FTP CVE frequencies			FTP Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
25	CVE-2015-3306	10	25	ProFTPD 1.3.5	1 10
2	CVE-2010-4221	10	12	vsFTPD 3.0.2	1 5
10	CVE-2011-4130	9	12	ProFTPD 1.3.4a	- -
10	CVE-2010-3867	7.1	10	ProFTPD 1.3.4b	- -
2	CVE-2008-2375	7.1	7	ProFTPD 1.3.1	8 9
10	CVE-2010-4652	6.8	6	ProFTPD 1.3.5a	- -
7	CVE-2009-0543	6.8	6	vsFTPD 2.3.5	- -
8	CVE-2009-3639	5.8	5	ProFTPD 1.3.3c	- -
12	CVE-2015-1419	5	5	ProFTPD 1.3.4c	- -
10	CVE-2011-1137	5	5	NASFTPD Turbo station 1.3.2e	- -

table 3: Top ten vulnerabilities found while scanning FTP using Scans.io

Each CVE ID listed is a discovered vulnerability. Those vulnerabilities with a CVSS score higher than 7 can potentially allow an attacker to compromise the service. Needless to say, these vulnerabilities should be patched by their respective system administrators.

A passive scan can only attempt to find CVE's on the basis of banners that the server returns. If a banner hasn't been sent or if a banner doesn't contain information which can be correlated to a CVE identifier, a CVE cannot be discovered. This can cause number of afflicted services to be higher. On the other hand, some Linux distributions (typically, stable or long term support versions) backport security fixes to older software versions for stability reasons. This may lead to over-reporting of CVE's, since the reported version is not precise enough to determine if it really is vulnerable to a specific CVE. Yet another reason CVE's can be over-reported when looking only at versions, is that some CVE's only apply to specific configurations. Also, some admins have set up specific countermeasures to mitigate high-risk CVE's.

To truly verify a vulnerability, the only way to make sure is by triggering the actual bug, which runs counter to the idea of a passive scan and involves the risk of causing damage to the owner of the service.




figure 3: Ratio of CVSS score for all vulnerabilities found using the HTTP protocol




figure 4: Ratio of CVSS score for all vulnerabilities found using the FTP protocol

The vulnerabilities found range from low to high severity, as shown in the tables and charts above. Each vulnerability basically amounts to a security risk where a malicious attacker could alter the integrity of the data, access confidential information or somehow influence the availability, or a

combination of the three. Some low impact examples include how a user name can be enumerated and the possibility of exploitation due to the lack of authentication.

High impact examples include a Denial of Service (DoS) attack, in which a malicious actor attempts to make a network unavailable to its users to temporarily (or indefinitely) interrupt services of a host connected to the Internet. Another high risk possibility is privilege escalation where an attacker could use cookies to violate configuration policy and cause a malicious client to be trusted.

A list explaining each CVSS vulnerability code in detail is included in the Appendix. These explanations above as well as the CVSS vulnerability codes in the Appendix also apply to the next section of the report.

## 5.2 Top Ten Shodan

Out of the thousands of vulnerabilities found, the following tables are the top ten versions and CVEs, as reported by Shodan, for each protocol: FTP, HTTP, MySQL, SMTP, and SSH.

HTTP CVE Frequencies			HTTP Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
19	CVE-2008-0075	10	167	Apache httpd 2	- -
94	CVE-2010-3972	10	94	Microsoft IIS httpd 7.5	5 10
21	CVE-2010-0425	10	75	Apache Tomcat/Coyote JSP eng	- -
2	CVE-2003-0224	10	64	Apache httpd 2.2.22	10 7.5
94	CVE-2010-2730	9.3	48	Apache httpd 2.4.7	9 6.8
21	CVE-2009-3023	9.3	47	Apache httpd 2.2.29	- -
33	CVE-2008-1446	9	38	nginx 5.3	- -
125	CVE-2010-1256	8.5	29	Apache httpd 2.4.6	10 6.8
50	CVE-2011-3192	7.8	29	nginx 1.4.6	- -
2	CVE-2009-1122	7.6	28	Apache httpd 2.2.15	24 7.8

table 4: Top ten vulnerabilities found while scanning HTTP using Shodan

FTP CVE frequencies			FTP Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
2	CVE-2010-4221	10	47	ProFTPD 1.3.5	1 10
47	CVE-2015-3306	10	34	ProFTPD 1.3.5a	- -
15	CVE-2011-4130	9	27	ProFTPD 1.3.4a	- -
15	CVE-2010-3867	7.1	27	ProFTPD 1.3.4b	- -
15	CVE-2010-4652	6.8	24	ProFTPD 1.3.3e	- -
13	CVE-2009-0543	6.8	20	ProFTPD 1.3.4c	- -
15	CVE-2009-3639	5.8	16	ProFTPD 1.3.3c	- -
15	CVE-2011-1137	5	14	Microsoft ftpd 7.5	- -
13	CVE-2008-7265	4	14	MiniServ 1.760	- -
15	CVE-2012-6095	1.2	13	ProFTPD 1.3.1	8 9

table 5: Top ten vulnerabilities found while scanning FTP using Shodan

MySQL CVE Frequencies			MySQL Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
9	CVE-2012-2750	10	43	MySQL 5.1.73	- -
22	CVE-2012-3163	9	41	MySQL 5.5.44	3 7.2
73	CVE-2014-6507	8	23	MySQL 5.5.45	11 4
73	CVE-2014-6491	7.5	22	MySQL 5.5.43	7 4.3
73	CVE-2014-6500	7.5	21	MySQL 5.6.26	17 4
22	CVE-2012-3158	7.5	20	MySQL 5.0.96	- -
22	CVE-2013-1492	7.5	17	MySQL 4.76	- -
22	CVE-2012-0553	7.5	13	MySQL 5.5.42	5 5.7
16	CVE-2012-0882	7.5	12	MySQL 5.5.31	48 8
3	CVE-2015-0411	7.5	12	MySQL 4.72	- -

table 6: Top ten vulnerabilities found while scanning MySQL using Shodan

SMTP CVE frequencies			SMTP Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
60	CVE-2010-4344	9.3	179	Exim smtpd 4.85	- -
131	CVE-2011-1764	7.5	129	Exim smtpd 4.76	3 6.8
71	CVE-2011-1407	7.5	77	Exim smtpd 4.86	- -
89	CVE-2010-4345	6.9	60	Exim smtpd 4.84	- -
89	CVE-2011-0017	6.9	46	Exim smtpd 4.69	8 9.3
294	CVE-2014-2957	6.8	42	Exim smtpd 4.73	5 7.5
200	CVE-2012-5671	6.8	34	Exim smtpd 4.80.1	2 6.8
26	CVE-2011-0411	6.8	29	Exim smtpd 4.72	7 7.5
26	CVE-2011-1720	6.8	29	Microsoft Exchange smtpd 7.5	- -
26	CVE-2008-2936	6.2	20	Postfix smtpd 6.0p	- -

table 7: Top ten vulnerabilities found while scanning SMTP using Shodan

SSH CVE frequencies			Versions, CVE's and max CVSS		
Count	CVE	CVSS	Count	Version	#CVE's Max CVSS
6	CVE-2000-0999	10	191	OpenSSH 5.3	7 7.5
77	CVE-2006-5051	9.3	77	OpenSSH 4.3	15 9.3
77	CVE-2006-4924	7.8	55	OpenSSH 6.0p	- -
288	CVE-2010-4478	7.5	51	OpenSSH 5.9p	- -
83	CVE-2007-4752	7.5	27	OpenSSH 6.6.1	- -
6	CVE-2001-0572	7.5	26	OpenSSH 5.5p	- -
301	CVE-2014-1692	7.5	15	OpenSSH 6.7p	- -
77	CVE-2009-2904	6.9	13	OpenSSH 5.3p	- -
8	CVE-2008-1657	6.5	10	OpenSSH 6.9p	- -
6	CVE-2013-4548	6	9	OpenSSH 6.1p	- -

table 8: Top ten vulnerabilities found while scanning SSH using Shodan

## 5.2.1 Differences Between Scans.io and Shodan

This Shodan scan confirms some of the results of the Scans.io results in the previous section, but deviates in several places. There are several reasons for this that we discovered during our research:

- Shodan gave some false positives with clearly incorrect results. For example, it reported Exim version numbers for SSH, and OpenSSH as FTP servers.. These have been left out of our overall analysis.
- When comparing scan results from a single host, it seemed that Shodan reports contain a relatively high number of incorrect versions. In our research, several HTTP hosts were reported by Shodan as "Apache 2", while the Scans.io raw data showed that they were not running this version, or were more accurate in the actual version number.
- Shodan information is more up-to-date (probably less than a few weeks old), while the Scans.io database used was about three months old when the final passive scan was performed.
- Considering the differences in version numbers, minor differences in version numbers can result in very different CVE distribution and maximum CVSS levels being reported. This is because a single software version can be prone to several different CVE's.

These differences make it hard to directly compare the results. On the other hand, both of them show that there are many vulnerable hosts in the Netherlands.

### 5.3 *Heartbleed*

This is a serious vulnerability in a specific version of the commonly-used OpenSSL cryptographic software library, which provides TLS (and SSL) capabilities for applications and services. Heartbleed allows an attacker to read parts of the service process' memory, including the keys used for the TLS connection. It can go undetected for long because TLS and even old SSL implementations don't often get updated, which in turn makes the impact much worse. The publicity around the Heartbleed Bug helped with resolving this, but it remains a problem with many "legacy" (i.e., outdated) systems without an active maintainer. Web, email, instant messaging, and some VPNs can all be affected by this malicious bug.

We found that 11 hosts out of the 1380 hosts scanned positive for the Heartbleed vulnerability. This amounts to 0.80% of hosts scanned. Compared to the number of other vulnerabilities found, this is a relatively low number. This is probably due to all the publicity around Heartbleed, as mentioned above.

### 5.4 *HTTP Headers*

The previous chapters showed information about vulnerabilities that were found. This chapter deals with mitigations: Certain HTTP headers can protect users from certain types of attacks. We evaluated how many websites implemented these headers.

X-Frame-Options
-----------------

Found by 22 out of 1380 hosts (1.6%)

This header provides Clickjacking protection. It can have the following values:

- deny – No rendering of the content within a frame
- sameorigin – No rendering within a frame if it isn't the same site as the content
- allow-from: DOMAIN – Allow rendering if framed by specified domain

Clickjacking is a type of attack where malicious actors present content of a legitimate site into a rogue frame. This enables attackers to 'siphon off' user clicks, that e.g. trigger malicious actions, like the installation of a virus. 98.4% of HTTP Headers scanned do not include this protection against clickjacking. It is recommended that website developers include this header using one of the values listed above.

Content-Security-Policy
-------------------------

Found by 0 of the 1380 hosts (0%)

If enabled correctly, a Content Security Policy protects users from potential cross-site scripting attacks using vulnerabilities in webpages. It impacts the way a browser renders pages, e.g., inline JavaScript disabled by default and must be explicitly allowed in the policy. None of the scanned headers have this protection. A Content Security Policy (header) requires careful tuning and precise definition of the policy.

X-Content-Type-Options
------------------------

Found by 24 out of 1380 hosts (1.7%)

When set with the only allowed value ("nosniff"), this header prevents Internet Explorer and Google Chrome from MIME-sniffing the response, as opposed to respecting the declared content-type. This also applies to Google Chrome when downloading extensions. This reduces exposure to drive-by download attacks and sites serving user-uploaded content that, by clever naming, could be treated by MSIE as executable or dynamic HTML files. 98.3% of scanned websites are potentially vulnerable to this type of attack.

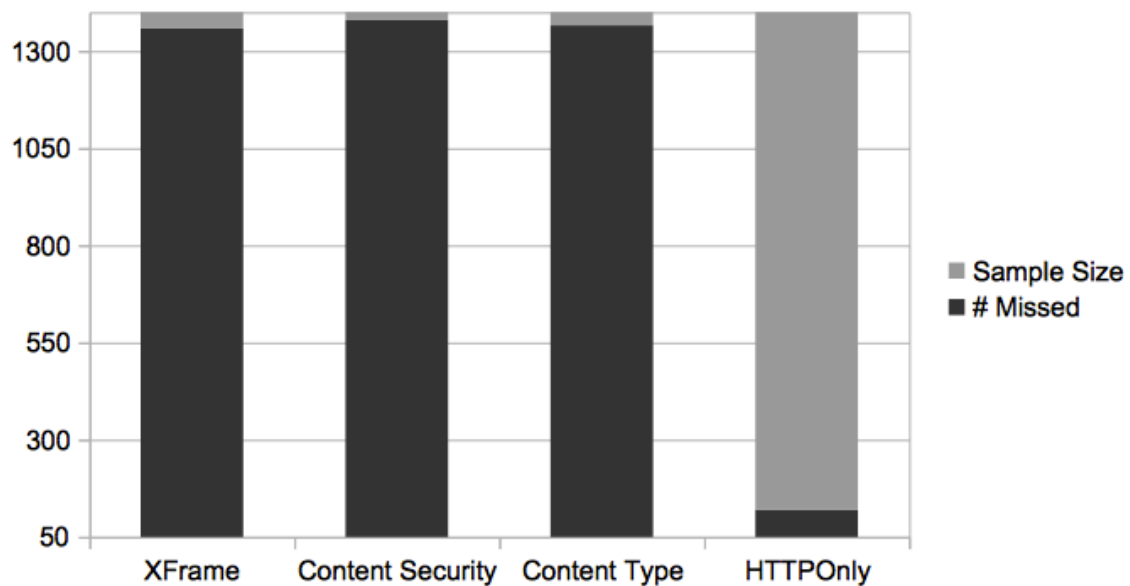


Figure 5: HTTP header sample size and missing HTTP headers

HttpOnly

Found by 1260 out of 1380 hosts (91.3%)

HttpOnly is an additional flag included in a Set-Cookie HTTP response header. Using the HttpOnly flag when setting a cookie makes sure that the cookie value will only be sent over the HTTP protocol. It mitigates against the risk of attackers trying to access the cookie using Javascript. 8.69% of scanned headers are vulnerable to this.

If the HttpOnly flag is included in the HTTP Set-Cookie response header, the cookie cannot be accessed through client side scripting. As a result, even if a cross-site scripting (XSS) flaw exists and a user accidentally accesses a link that exploits this flaw, the browser will not reveal the cookie to a third party.

If a browser does not support HttpOnly and a website attempts to set an HttpOnly cookie, the HttpOnly flag will be ignored by the browser, thus creating a traditional, script accessible cookie. As a result, the cookie (typically your session cookie) becomes vulnerable to theft or modification by malicious script. This is not a thorough fix, but with more browsers supporting it, the coverage it provides is becoming better and better.

## 5.5 SSL Research

We pulled all the passive SSL data from the CIRCL.lu database (<https://circl.lu/services/passive-ssl/>) and looked for issues with certificates. Some certificates for example are blacklisted:

certificates whose corresponding private keys have been leaked, or certificates that are known to be (ab)used by malware.

Abuse.ch was used to pull IP addresses and other important information that can help us uncover issues with certificates. We found certificates flagged by Abuse.ch by consulting the Abuse.ch SSL Blacklist provided here: <https://sslbl.abuse.ch/blacklist/>. We checked every SSL/TLS certificate hash against the SSL blacklist CSV.

If certificates are found in this blacklist this means they are known to be used by malicious attackers for spreading malicious software (malware). They could either have once been valid certificates that have been stolen or valid certificates that have been created by the malicious actors.

#### Findings:

A total of 6 blacklisted SSL/TLS certificate hashes were found in the dataset. The most common bad certificate was ee14e4ab0b243b397315e094935f5b74a67a1bc7 <https://sslbl.abuse.ch/intel/ee14e4ab0b243b397315e094935f5b74a67a1bc7>

The various blacklisted certificates were all associated with different malware command and control centers; an infrastructure of servers used to control malware like viruses, worms, spyware, adware, and trojan horses. There are several different 'brands' identified. "Malware Command & Control (C&C)" is a generic brand of malware. The others (URLzone, Gozi, KINS, and Dridex) are known brands of malware. This usually means that the servers have been taken over by attackers and are being used as assets by someone else.

The full list is as follows:

- 415a586a121158602392d56394a5903dfe222a0c - URLzone C&C
- 8ae76c20358ef9b66c8f547f5aecd15a7fdb169e - Gozi MITM
- 9cc93e0add5efdf43bc566836e5666d7049701ed - Gozi C&C
- a8fb10cb9a222372646f0b7b3a6e4d228ef254d3 - KINS C&C
- ceb8616d116b0bcfeefc3a1b15dde6e004f8cc1a - Dridex C&C
- ee14e4ab0b243b397315e094935f5b74a67a1bc7 - Malware C&C



## 6 Future Work

We think this tool is the starting point for a lot of interesting research. We currently see the following new avenues, applications and extensions for the future. Examples of this include:

- Trend spotting: regularly running the scanner over the same address pool and compare the differences between the runs.
- Research differences between geographical or economical areas: run the scanner on IP spaces from different geographical areas (countries, continents) or economical areas (western/developed world, BRICS, developing world).
- Extend to scanner to "pivot" on other services than just web services: the scanner currently select web servers and scans those for CVE's on common services. The same kind of scan can be performed on e.g. mail servers, SCADA devices, NTP servers, DNS servers etc.
- Extend to scanner to perform a more generic scan: this sort of scan would be comparable to what many existing tools do, by scanning an IP address range for vulnerabilities over a range of different services. The added value of our scanner is that it is passive, and because it samples the population, it is faster.
- Extend the scanner to scan more kinds of services.
- Extend the scanner to work with IPv6. Obviously, IPv6 is becoming more adopted and important in the next few years. As a personal observation of one of the authors (falling outside the scope of this research), his IPv6-enabled host with public services has attracted much more IPv6 scans in 2015 than in 2014. IPv6 seems to be around the inflection point for large-scale adoption.
- Attackers often use large scale scanning and compromising of servers for various attacks. SIDN could benefit from offering their customers additional active scans or email them the results of passive scans.
- We think it may be worth reviewing our sample size estimation method; does sampling for a binary choice (yes/no) really have the same minimum size estimate as sampling for a many-options choice (banners; we found over > 100 different banners)?
- Matching banners and versions in the CVE database is not always easy to do, and is therefore not 100% reliable. It may be worth investigating this and improving banner-CVE matching.

## 7 Conclusion

We found that it's possible to assess the state of security using only existing, passive data sets in a non-obtrusive way. We developed software to perform such a scan, which is publicly available on GitHub (see also the [Compilation appendix](#) (page 39)).

Subsequently we used the tool to scan a random subset of Dutch IP addresses running webservers. Based on our scan results, we found that a large number of the scanned Dutch IP addresses contained vulnerabilities having high and medium impact.

The top 10 of most encountered vulnerabilities per protocol shows the following:

- MySQL vulnerabilities:
  - Oracle MySQL features in every spot of the top 10
  - Vulnerabilities are either "Unspecified" or a Denial of Service
- SMTP vulnerabilities:
  - Exim and Postfix have the most vulnerabilities assigned in the Top 10 list
  - Remote code execution is the most prevalent vulnerability
- SSH vulnerabilities:
  - Most of the vulnerabilities involve either Denial of Service or information disclosure
- HTTP vulnerabilities:
  - Microsoft IIS and Nginx feature in the top 10
  - Remote code execution is the most prevalent vulnerability
- FTP vulnerabilities:
  - The most-used software is clearly ProFTPD. Many ProFTPD services were found vulnerable.
  - Denial of Service, user enumeration and information disclosure are the primary vulnerabilities found in FTP

The majority of the vulnerabilities found are due to outdated, vulnerable software being used. We found that about a quarter all dutch webservers are vulnerable. This indicates that a large number of Dutch webservers are not frequently updated to new, more secure versions.

HTTP headers that improve security for users of the webservice, like 'X-Frame-Options', 'Content-Security-Policy,' and 'X-Content-Type-Options', were rarely used.

Scans.io proved to be a very useful resource from scanning information. We found Shodan to be useful but somewhat inaccurate. After some initial research, Archive.org has not been considered any further because the added value was less than the error margin of the statistical test. This made the results from Archive.org statistically insignificant and thus not worth the effort.

Services like IMAP, POP3 were also not considered in the analysis because these services rarely announced their version number.

## Appendix 1 CVE Vulnerability Entries

The following information explains vulnerabilities having a CVSSv2 base score of 5.0 or higher.

### FTP PROTOCOLS

#### **CVE-2013-4359:**

Software: ProFTPD FTP Server

CVSS2 Score: 5.0

Vulnerability: Denial of Service

Summary: An "Integer Overflow" in kbdint.c of the "mod\_sftp" module of ProFTPD 1.3.4d and 1.3.5r3 could cause a Denial of Service vulnerability.

Impact: Low despite access being network based with no authentication required. This is due to only the "availability" being affected upon successful exploitation.

#### **CVE-2015-3306**

Software: ProFTPD FTP Server

CVSS2 Score: 10

Vulnerability: Information Disclosure

Summary: The "mod\_copy" module in version 1.3.5 allows remote attackers to read and write to arbitrary files via the built in "CPFR" and "CPTO" commands.

Impact: High impact as attackers could use this vulnerability to read or write sensitive files on the remote system. There are also publicly available exploits which make it much easier to attackers to exploit the vulnerability.

#### **CVE-2006-4924**

Software: OpenSSH

CVSS2 Score: 7.9

Vulnerability: Denial of Service

Summary: The SSHD binary in OpenSSH before version 4.4, when using the version 1 SSH protocol allowed remote attackers to cause a Denial of Service attack (CPU Consumption).

Impact: A proof of concept exploit is available and as such the exploitability of this vulnerability is high.

#### **CVE-2006-5051**

Software: OpenSSH

CVSS2 Score: 9.3

Vulnerability: Denial of Service

Summary: There is a signal handler race condition in OpenSSH prior to version 4.4 that could allow a remote attacker to crash the service. There is also the possibility of remote code execution should GSSAPI authentication be enabled in the daemon configuration.

Impact: If the denial of service/crash is exploited, valid users of the service will not be able to access the server. If the remote code execution vulnerability is exploited, an attacker could execute code on the server.

#### **CVE-2006-5052**

Software: OpenSSH

CVSS2 Score: 5.0

Vulnerability: User enumeration

Summary: An "Unspecified" vulnerability in portable OpenSSH prior to version 4.4 running on some platforms could allow username enumeration.

Impact: Low impact as only the usernames can be enumerated. This would allow an attacker to launch brute force attacks against the server; however, there should be measures in place to defend against this.

#### **CVE-2007-2243**

Software: OpenSSH

CVSS2 Score: 5.0

Vulnerability: User enumeration

Summary: In OpenSSH 4.6 and earlier with the "ChallengeResponseAuthentication" option enabled allows remote attackers to determine valid users accounts by attempting to authenticate via the S/KEY authentication method.

Impact: Low impact as this will only give valid users to the attacker. These could be used later in brute force attacks, however these are easy to detect and defend against.

#### **CVE-2007-4752**

Software: OpenSSH

CVSS2 Score: 7.5

Vulnerability: Privilege escalation

Summary: The "ssh" client binary prior to version 4.7 doesn't handle authentication X11 cookies properly.

Impact: High: an attacker could use untrusted X11 cookies to violate configuration policy and cause an X client to be treated as trusted.

#### **CVE-2008-3259**

Software: OpenSSH

CVSS2 Score: 1.2

Vulnerability: Information Exposure

Summary: OpenSSH prior to 5.1 sets a socket option which could be hijacked by attackers when using the X11 port forwarding option.

Impact: Low impact due to the specific requirements and vulnerable platforms of the affected service.

#### **CVE-2009-2904**

Software: OpenSSH

Vulnerability: Privilege Escalation

CVSS2 Score: 6.9

Summary: A certain Red Hat modification to the ChrootDirectory feature in OpenSSH 4.8, as used in SSHD in OpenSSH 4.3 in Red Hat Enterprise Linux (RHEL) 5.4 and Fedora 11, allows local users to gain privileges via hard links to setuid programs that use configuration files within the chroot directory.

Impact: Low.

### **HTTP PROTOCOLS**

#### **CVE-2010-1256**

Software: Microsoft IIS web server

CVSS2 Score: 8.5

Vulnerability: Remote code execution

Summary: Unspecified vulnerability in Microsoft IIS 6.0, 7.0, and 7.5, when Extended Protection for Authentication is enabled, allows remote authenticated users to execute arbitrary code via unknown vectors related to \"token checking\" that trigger memory corruption, aka \"IIS Authentication Memory Corruption Vulnerability.

Impact: Attackers could execute code of their choosing on the affected server.

#### **CVE-2010-1899**

Software: Microsoft IIS web server

CVSS2 Score: 4.3

Vulnerability: Denial of Service

Summary: Stack consumption vulnerability in the ASP implementation in Microsoft Internet Information Services (IIS) 5.1, 6.0, 7.0, and 7.5 allows remote attackers to cause a denial of service (daemon outage) via a crafted request, related to asp.dll, aka \"IIS Repeated Parameter Request Denial of Service Vulnerability.\"

Impact: Valid users of the service will not be able to access the service.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS10-065.msp>

#### **CVE-2010-2730**

Software: Microsoft IIS web server

CVSS2 Score: 9.3

Vulnerability: Remote code execution

Summary: When the FastCGI option is enabled on version 7.5 of IIS remote attackers could execute arbitrary code on the server via a specially crafted header in a request.

Impact: High impact due to attackers being able to execute code of their choosing on the server.

This is somewhat mitigated by the complexity of the attack.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS10-065.msp>

#### **CVE-2010-3972**

Software: Microsoft IIS FTP service

CVSS2 Score: 10

Vulnerability: Remote code execution / Denial of Service

Summary: Remote attackers could execute arbitrary code or cause a denial of service condition via a specially crafted FTP command.

Impact: High. Exploits are publicly available.

Additional Information: <http://www.microsoft.com/technet/security/Bulletin/MS11-004.msp>

#### **CVE-2012-2531**

Software: Microsoft IIS Web service

CVSS2 Score: 2.1

Vulnerability: Sensitive information disclosure

Summary: Weak permissions on the Operational log could allow local users to discover credentials by reading the file.

Impact: Medium, Sensitive information disclosure based on specific application configuration.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS12-073.msp>

#### **CVE-2014-3616**

Software: Nginx

CVSS2 Score: 4.3

Vulnerability: Improper access control

Summary: Version 0.5.6 through 1.7.4 when using the same specific shared SSL configuration options for multiple servers can reuse a caches SSL session in an unrelated context which could allow a remote attacker with certain privileges to conduct "virtual host confusion attacks".

Impact: Low.

#### **CVE-2013-4547**

Software: Nginx

CVSS2 Score: 7.5

Vulnerability: Access control bypass / sensitive information disclosure

Summary: Certain configurations in certain versions of Nginx could allow remote attackers to bypass intended restrictions via an unescaped space character in the URI.

Impact: Medium: Under certain conditions an attacker could bypass access restrictions to privileged files.

#### **CVE-2013-0337**

Software: Nginx

CVSS2 Score: 7.5

Vulnerability: Sensitive Information Disclosure

Summary: Insecure permissions on the access.log and error.log files allow local users to obtain sensitive information by reading the files.

Impact: Medium. Sensitive information could be disclosed should the application log this kind of information to file.

#### **CVE-2013-2070**

Software: Nginx

CVSS2 Score: 5.8

Vulnerability: Denial of Service / Information disclosure

Summary: A denial of service condition could be caused by a remote attacker when the proxy\_pass option in the ngx\_htt\_proxy\_module of Nginx is used by an untrusted HTTP server. The attacker could then obtain sensitive information from the Nginx worker process memory via a specially crafted proxy response.

Impact: Medium. Specific configuration is required along with attack complexity.

#### **CVE-2014-2957**

Software: Exim Mail software - Not relevant in this section

### **HTTP PROTOCOLS**

#### **CVE-2010-1256**

Software: Microsoft IIS web server

CVSS2 Score: 8.5

Vulnerability: Remote code execution

Summary: Unspecified vulnerability in Microsoft IIS 6.0, 7.0, and 7.5, when Extended Protection for Authentication is enabled, allows remote authenticated users to execute arbitrary code via unknown vectors related to "\"token checking\" that trigger memory corruption, aka "\"IIS Authentication Memory Corruption Vulnerability.

Impact: High impact to confidentiality, integrity and availability.

#### **CVE-2010-1899**

Software: Microsoft IIS web server

CVSS2 Score: 4.3

Vulnerability: Denial of Service

Summary: Stack consumption vulnerability in the ASP implementation in Microsoft Internet Information Services (IIS) 5.1, 6.0, 7.0, and 7.5 allows remote attackers to cause a denial of service (daemon outage) via a crafted request, related to asp.dll, aka "IIS Repeated Parameter Request Denial of Service Vulnerability".

Impact: Medium impact to availability.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS10-065.msp>

#### **CVE-2010-2730**

Software: Microsoft IIS web server

CVSS2 Score: 9.3

Vulnerability: Remote code execution

Summary: When the FastCGI option is enabled on version 7.5 of IIS remote attackers could execute arbitrary code on the server via a specially crafted header in a request.

Impact: High impact to confidentiality, integrity and availability.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS10-065.msp>

#### **CVE-2010-3972**

Software: Microsoft IIS FTP service

CVSS2 Score: 10

Vulnerability: Remote code execution / Denial of Service

Summary: Remote attackers could execute arbitrary code or cause a denial of service condition via a specially crafted FTP command.

Impact: High impact to confidentiality, integrity and availability. Exploits are publicly available.

Additional Information: <http://www.microsoft.com/technet/security/Bulletin/MS11-004.msp>

#### **CVE-2012-2531**

Software: Microsoft IIS Web service

CVSS2 Score: 2.1

Vulnerability: Sensitive information disclosure

Summary: Weak permissions on the Operational log could allow local users to discover credentials by reading the file.

Impact: Medium impact to confidentiality.

Additional information: <http://www.microsoft.com/technet/security/Bulletin/MS12-073.msp>

#### **CVE-2014-3616**

Software: Nginx

CVSS2 Score: 4.3

Vulnerability: Improper access control

Summary: Version 0.5.6 through 1.7.4 when using the same specific shared SSL configuration options for multiple servers can reuse a caches SSL session in an unrelated context which could allow a remote attacker with certain privileges to conduct "virtual host confusion attacks".

Impact: Medium impact to integrity.

#### **CVE-2013-4547**

Software: Nginx

CVSS2 Score: 7.5

Vulnerability: Access control bypass / information disclosure

Summary: Certain configurations in certain versions of Nginx could allow remote attackers to bypass intended restrictions via an unescaped space character in the URI.



Impact: Medium impact to confidentiality, integrity and availability under certain conditions an attacker could bypass access restrictions to privileged files.

#### **CVE-2013-0337**

Software: Nginx

CVSS2 Score: 7.5

Vulnerability: Sensitive Information Disclosure

Summary: Insecure permissions on the access.log and error.log files allow local users to obtain sensitive information by reading the files.

Impact: Medium impact to confidentiality, integrity and availability. Sensitive information could be disclosed should the application log this kind of information to file.

#### **CVE-2013-2070**

Software: Nginx

CVSS2 Score: 5.8

Vulnerability: Denial of Service / Information disclosure

Summary: A denial of service condition could be caused by a remote attacker when the proxy\_pass option in the ngx\_htt\_proxy\_module of Nginx is used by an untrusted HTTP server. The attacker could then obtain sensitive information from the Nginx worker process memory via a specially crafted proxy response.

Impact: Medium impact to confidentiality and availability. Specific configuration is required along with attack complexity.

### **MYSQL SERVICES**

#### **CVE-2015-0499**

Software: Oracle MySQL 5.5.42 and earlier / 5.6.23 and earlier

CVSS2 Score: 3.5

Vulnerability: Denial of Service

Summary: An unspecified vulnerability allows remote authenticated users to affect availability.

Impact: Low impact to availability due to authentication being required for exploitation.

#### **CVE-2015-0501**

Software: Oracle MySQL 5.5.42 and earlier / 5.6.23 and earlier

CVSS2 Score: 5.7

Vulnerability: Denial of Service

Summary: An unspecified vulnerability allows remote authenticated users to affect availability.

Impact: Medium impact availability due to authentication being required for exploitation.

#### **CVE-2015-0505**

Software: Oracle MySQL 5.5.42 and earlier / 5.6.23 and earlier

CVSS2 Score: 3.5

Vulnerability: Denial of Service

Summary: An unspecified vulnerability allows remote authenticated users to affect availability.

Impact: Low impact to availability due to authentication being required for exploitation.

#### **CVE-2014-6464**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 4.0

Vulnerability: Denial of Service

Summary: An unspecified vulnerability allows remote authenticated users to affect availability.

Impact: Low impact to availability due to authentication being required for exploitation.

**CVE-2014-6469**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 6.8

Vulnerability: Denial of Service

Summary: An unspecified vulnerability allows remote authenticated users to affect availability.

Impact: Medium impact to availability due to authentication being required for exploitation.

**CVE-2014-6491**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 7.5

Vulnerability: Unspecified vulnerability

Summary: Unspecified vulnerability allows remote attackers to affect confidentiality, integrity, and availability via vectors related to SERVER:SSL:yaSSL.

Impact: Medium impact to confidentiality, integrity and availability.

**CVE-2014-6494**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 4.3

Vulnerability: Denial of Service

Summary: Unspecified vulnerability allows remote attackers to affect confidentiality, integrity, and availability via vectors related to SERVER:SSL:yaSSL.

Impact: Low impact to availability.

**CVE-2014-6496**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 4.3

Vulnerability: Denial of Service

Summary: Unspecified vulnerability allows remote attackers to affect confidentiality, integrity, and availability via vectors related to CLIENT:SSL:yaSSL.

Impact: Low impact to availability.

**CVE-2014-6500**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 7.5

Vulnerability: Unspecified vulnerability

Summary: Unspecified vulnerability allows remote attackers to affect confidentiality, integrity, and availability via vectors related to SERVER:SSL:yaSSL

Impact: Medium impact to confidentiality, integrity and availability

**CVE-2014-6507**

Software: Oracle MySQL 5.5.39 and earlier / 5.6.20 and earlier

CVSS2 Score: 8.0

Vulnerability: Unspecified vulnerability

Summary: Unspecified vulnerability allows remote authenticated users to affect confidentiality, integrity, and availability via vectors related to SERVER:DML

Impact: Medium impact to confidentiality/integrity + high impact to availability. This is somewhat mitigated by the need for authentication.

**SMTP SERVICES**

**CVE-2014-2957**

Software: Exim before 4.82.1

CVSS2 Score: 6.8

Vulnerability: Remote code execution

Summary: With the EXPERIMENTAL\_DMARC option enabled, remote attackers could execute arbitrary code via the From header in an email, which is passed to the expand\_string function.

Impact: Medium impact to confidentiality, integrity and availability.

**CVE-2014-2972**

Software: Exim before 4.83

CVSS2 Score: 4.6

Vulnerability: Privilege escalation / arbitrary command execution

Summary: expand.c in Exim expands mathematical comparisons twice which could allow local users to gain privileges and execute arbitrary commands via a specially crafted lookup value.

Impact: Medium impact to confidentiality, integrity and availability but local access is required.

**CVE-2012-5671**

Software: Exim 4.70 - 4.80

CVSS2 Score: 5.4

Vulnerability: Remote code execution

Summary: Remote attackers could execute arbitrary code when DKIM support is enabled and configuration options acl\_smtp\_connect and acl\_smtp\_rcpt are not set to "warn control = dkim\_disable\_verify.

Impact: Medium impact to confidentiality, integrity and availability. However attack complexity and specific configuration requirements mitigate the risk to a certain degree.

**CVE-2010-4345**

Software: Exim 4.72 and earlier

CVSS2 Score: 6.9

Vulnerability: Privilege escalation

Summary: Local users can leverage the ability of the Exim user account to specify an alternative configuration file to gain privileges.

Impact: High impact to confidentiality, integrity and availability. This is compounded by publically available exploits. However local access is required.

**CVE-2011-0017**

Software: Exim 4.72 and earlier

CVSS2 Score: 6.9

Vulnerability: Summary: The open\_log function in log.c does not check the return value from setuid or setgid system calls. This could allow local users to append log data to arbitrary files via a symlink attack.

Impact: Medium impact to confidentiality, integrity and availability mitigated somewhat by the local access requirement.

**CVE-2011-1407**

Software: Exim 4.76 and earlier

CVSS2 Score: 7.5

Vulnerability: Remote code execution/Information disclosure

Summary: The DKIM implementation permits matching for DKIM identities to apply to lookup items instead of only strings which could allow remote attackers to execute arbitrary code or access the filesystem via a specially crafted identity.

Impact: Medium impact to confidentiality, integrity and availability.

#### **CVE-2011-1764**

Software: Exim 4.76

CVSS2 Score: 7.5

Vulnerability: Format string vulnerability

Summary: Format string vulnerability in the `dkim_exim_verify_finish` function in `src/dkim.c` in Exim before 4.76 might allow remote attackers to execute arbitrary code or cause a denial of service.

Impact: Medium impact to confidentiality, integrity and availability based off configuration requirements.

#### **CVE-2011-1720**

Software: Postfix before 2.5.13, 2.6.10, 2.7.4 and 2.8.3

CVSS2 Score: 6.8

Vulnerability: Buffer overflow

Summary: The SMTP server, when certain Cyrus SASL authentication methods are enabled, does not create a new server handle after client authentication fails, which allows remote attackers to cause a denial of service or possibly execute arbitrary code via an invalid AUTH command with one method followed by an AUTH command with a different method.

Impact: Medium impact to confidentiality, integrity and availability.

#### **CVE-2014-3956**

Software: Sendmail before 8.14.9

CVSS2 Score: 1.9

Vulnerability: Information disclosure

Summary: The `sm_close_on_exec` function in `conf.c` in sendmail before 8.14.9 has arguments in the wrong order, and consequently skips setting expected `FD_CLOEXEC` flags, which allows local users to access unintended high-numbered file descriptors via a custom mail-delivery program.

Impact: Low impact to confidentiality due to local access requirement and access complexity.

#### **CVE-2012-0811**

Software: Postfix Admin before 2.3.5

CVSS2 Score: 6.5

Vulnerability: SQL injection

Summary: SQL injection vulnerabilities in Postfix Admin could allow remote authenticated users to execute arbitrary SQL commands.

Impact: Medium impact of confidentiality, integrity and availability mitigated somewhat by the requirement for authenticated users.

### **SSH SERVICES**

#### **CVE-2010-4478**

Software: OpenSSH 5.6 and earlier

CVSS2 Score: 7.5

Vulnerability: Authentication Bypass

Summary: With J-PAKE enabled, OpenSSH does not properly validate the public parameters in the protocol which could allow remote attackers to bypass the need for knowledge of the shared secret and successfully authenticate.

Impact: Medium impact but with specific configuration requirements.

#### **CVE-2010-4755**

Software: OpenSSH 5.8 and earlier

CVSS2 Score: 4.0

Vulnerability: Denial of Service

Summary: The remote\_blob function in sftp-glob.c and process\_put in sftp.c allow remote authenticated users to cause a denial of service condition via crafted glob expressions that do not match any path names.

Impact: Medium, however authentication is required for the exploit to work.

#### **CVE-2010-5107**

Software: OpenSSH through 6.1

CVSS2 Score: 5.0

Vulnerability: Denial of Service

Summary: The default configuration of OpenSSH enforces a fixed time limit between establishing a TCP connection and completing a login. This could make it easier for remote attackers to cause a denial of service condition by periodically making many new TCP connections.

Impact: Medium. No authentication required.

#### **CVE-2011-4327**

Software: OpenSSH before 5.8p2

CVSS2 Score: 2.1

Vulnerability: Information Disclosure

Summary: ssh-keygen.c in ssh-keygen on certain platforms executes ssh-rand-helper with unintended open file descriptors, which allows local users to obtain sensitive key information via the ptrace system call.

Impact: Medium impact to confidentiality. Local access is required.

#### **CVE-2011-5000**

Software: OpenSSH 5.8 and earlier

CVSS2 Score: 3.5

Vulnerability: Denial of Service

Summary: The ssh\_gssapi\_parse\_ename function in gss-serv.c allows remote authenticated users to cause a denial of service condition via a large value in a certain length field.

Impact: Medium impact to availability, mitigated by the requirement for authentication.

#### **CVE-2012-0814**

Software: OpenSSH before 5.7

CVSS2 Score: 3.5

Vulnerability: information disclosure

Summary: The auth\_parse\_options function in auth-options.c in sshd provides debug messages containing authorized\_keys command options which allows remote authenticated attackers to obtain potentially sensitive information by reading these messages.

Impact: Medium impact to confidentiality, mitigated by the requirement for authentication.

#### **CVE-2014-1692**

Software: OpenSSH through 6.4

CVSS2 Score: 7.5

Vulnerability: Denial of Service

Summary: The hash\_buffer function in schnorr.c, when Makefile.inc is modified to enable the J-PAKE protocol, does not initialize certain data structures which could allow remote attackers to cause a denial of service condition.

Impact: Medium impact to confidentiality, integrity and availability.

#### **CVE-2006-4924**

Software: OpenSSH before 4.4

CVSS2 Score: 7.9

Vulnerability: Denial of Service

Summary: sshd, when using version 1 SSH protocol, allows remote attackers to cause a denial of service condition via an SSH packet that contains duplicate blocks, which is not properly handled by the CRC compensation attack detector.

Impact: Medium impact to availability.

#### **CVE-2006-5051**

Software: OpenSSH before 4.4

CVSS2 Score: 9.3

Vulnerability: Denial of Service/Remote Code Execution

Summary: A signal handler race condition could allow remote attackers to cause a denial of service condition and possibly execute arbitrary code if GSSAPI authentication is enabled.

Impact: High impact to confidentiality, integrity and availability.

#### **CVE-2006-5052**

Software: OpenSSH before 4.4

CVSS2 Score: 5.0

Vulnerability: Username enumeration

Summary: An unspecified vulnerability in the portable version of OpenSSH prior to 4.4 when running on some platforms, could allow remote attackers to determine the validity of usernames via known vectors involving a GSSAPI "authentication abort".

Impact: Medium impact to confidentiality.

## Appendix 2 PassiveScanning Tool

### App 2.1 Compilation

<https://github.com/radicallyopensecurity/PassiveNLSurvey> contains a compilation guide and a formatted usage.

radicallyopensecurity / **PassiveScanningTool**

22 commits 1 branch 0 releases 1 contributor

Branch: master **PassiveScanningTool** / +

File	Commit Message	Time Ago
koenj2 Create LICENSE.md	Latest commit 3af7457	5 hours ago
Cve	Added support for Shodan.	15 days ago
Properties	First commit.	3 months ago
Results	Added archive.org.	22 hours ago
Scanslo	Fixed a small mistake where the port was not output correctly.	a month ago
Shodan	Added support for Shodan.	15 days ago
FindServiceDescriptor.cs	Added the possibility of using Rapid7 scan results.	a month ago
Host.cs	Added archive.org.	22 hours ago
HostList.cs	Removed a call from the host list.	2 months ago
LICENSE.md	Create LICENSE.md	5 hours ago
Makefile	First commit.	3 months ago
Newtonsoft.Json.dll	First commit.	3 months ago

Code

Issues 4

Pull requests 0

Pulse

Graphs

HTTPS clone URL

<https://github.com/radicallyopensecurity/PassiveScanningTool>

You can clone with HTTPS or Subversion.

Clone in Desktop

Download ZIP

### App 2.2 Usage

To be able to use the application, the following ZMAP data files have to be downloaded from the repository. The files should be stored in the path /data.

ZMAP results and ZGRAB results:

```
imap-starttls
ftp-banner
pop3s-tls
https-heartbleed
smtp-starttls
imaps-tls
https-tls
pop3-starttls
Rapid7 results:
http
```

Additionally, all files from the NVD data repository (<https://nvd.nist.gov/download.cfm>) have to be downloaded in order to build a CVE database (2002 to Recent). The files should be stored in the folder, "nvdcve." The CVE database allows CVEs to be correlated to version numbers. The Shodan data source is automatically loaded during run-time only for the sampled hosts. Archive.org is also automatically loaded.

Finally, a list of IP ranges should be supplied in a format similar to that of all Dutch IPs (<http://www.nirsoft.net/countryip/nl.html>) found on <http://www.nirsoft.net/countryip/nl.csv>.

The program can now run, but it will generate a few intermediate outputs in the directory, '/data/output'. These intermediate files (services-\*) are output to a CSV format 'host;name;port;jsondata'. The intermediate results are then analyzed, which results in more output files. The first set of output files (software-frequency-\*) shows the frequency of specific software versions used, which follows the format 'software;frequency'. The second set of output files (cve-frequency-\*) shows the frequency of specific CVEs found, which also follows the format 'software;score;frequency'.

In order to speed up analysis after the first run, it generated intermediate files that stored the more time-intensive results. These were stored in 'UniqueAddressList', 'HostInformation', 'ShodanHostInformation', and 'CveDatabase'. In case the investigated host had changed, the 'UniqueAddressList', 'HostInformation', and 'ShodanHostInformation' should be removed before restarting the application.

### PassiveNLSurvey Example Output

The following is an example of output from our PassiveNLSurvey tool:

```
Melanies-MacBook-Pro:PassiveScanning melanie$ mono PassiveScanning.exe 85.17.171.95
Loading CVE database...
CVE database loaded.  PassiveNLSurvey
Found 1 unique HTTP servers.
1 hosts fetched.
Generating host information...
Generated host information.
Searching for software banners/versions and CVE's...
Found missing header X-Frame-Options for host 85.17.171.95.
Found missing header Content-Security-Policy for host 85.17.171.95.
Found missing header X-Content-Type-Options for host 85.17.171.95.
Done!
Loading Shodan host list...
Total hosts: 1
Loaded 1 hosts from Shodan.
Found 1 Shodan host.
Searching for Shodan software banners/versions and CVE's...
Found software OpenSSH 6.7p.
Searching Shodan for comments on websites...
Total hosts: 1.
Processing host 1.
Processing hostname webmail.dubbele.com.
Processing hostname 85.17.171.95.
```



## Appendix 3 Researcher Biographies

Melanie Rieback	Melanie Rieback is a former Asst. Prof. of Computer Science from the VU, who is also the co-founder/CEO of Radically Open Security.
Koen Jeukendrup	Koen is a lifelong hardware and software enthusiast, employed as a security consultant for the past four years. He's currently working on a Master's in Electrical Engineering.
Matt Erasmus	Matt is an Information Security professional with interests in network forensics, malware analysis and penetration testing. He enjoys splitting his time between building and breaking in equal measure.
Eireann Leverett	Eireann has studied psychology, philosophy, artificial intelligence, software engineering, and computer security at various times in his life. He holds a BEng from Edinburgh Univesity and an MPhil from the University of Cambridge in Advanced Computer Science. He still enjoys punting at Darwin College when he has the time.