A webserver's nightmare – Serving files that let me pwn you

BerlinSides 0x7E2

@gehaxelt

June 23, 2018

Agenda

1. Intro & something about webservers
2. Interesting files
3. Scanning for files
4. Feedback || Answers && Questions

Attention!

Intro & something about webservers

$>whoami

- Mr. @gehaxelt / `0day.work`
- Co-Founder of `Internetwache.org`
- MSc CS student at TU Berlin
- <3 CTFs @ ENOFLAG
  - Join us for the FAUST-CTF
  - Or sponsor our Defcon trip ;-)

Webservers...

- How do we identify webservers?

[1] https://news.netcraft.com/archives/2018/04/26/april-2018-web-server-survey.html

Webservers...

- How do we identify webservers?

```
PORT    STATE SERVICE
80/tcp  open  http
443/tcp open  https
```

- Who operates a webserver?

[1] https://news.netcraft.com/archives/2018/04/26/april-2018-web-server-survey.html

Webservers...

- How do we identify webservers?

```
PORT    STATE SERVICE
80/tcp  open  http
443/tcp open  https
```

- Who operates a webserver?
- Who shut off his server because of GDPR? ;-)

[1] https://news.netcraft.com/archives/2018/04/26/april-2018-web-server-survey.html

Webservers...

- How do we identify webservers?

```
PORT    STATE SERVICE
80/tcp  open  http
443/tcp open  https
```

- Who operates a webserver?
- Who shut off his server because of GDPR? ;-)
- What's the most used webserver software?

[1] https://news.netcraft.com/archives/2018/04/26/april-2018-web-server-survey.html

Webservers...

- How do we identify webservers?

```
PORT     STATE SERVICE
80/tcp   open  http
443/tcp  open  https
```

- Who operates a webserver?
- Who shut off his server because of GDPR? ;-)
- What's the most used webserver software?

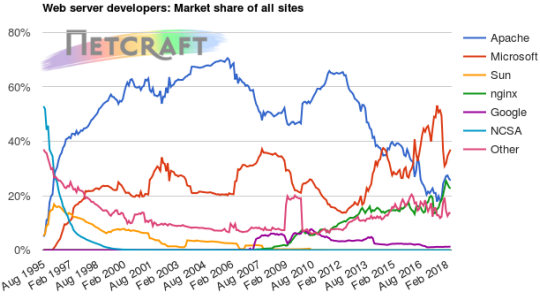

1

Developer	March 2018	Percent	April 2018	Percent	Change
Microsoft	633,719,941	35.80%	658,800,756	36.94%	1.15
Apache	464,340,535	26.23%	456,169,336	25.58%	-0.65
nginx	409,124,174	23.11%	403,381,961	22.62%	-0.49
Google	21,802,670	1.23%	22,460,562	1.26%	0.03

1

[1] https://news.netcraft.com/archives/2018/04/26/april-2018-web-server-survey.html

Who should listen? Someone who...

- ... develops websites using Git/SVN/Mercurial ?
  - ... deploys them on the server using these tools (e.g. git pull)?
- ... has a MacOS based system?
  - ... deploys using rsync/scp/(s)ftp ?
- ... develops using Sublime Text and the 'SFTP'-Plugin?

Who should listen? Someone who...

- ... develops websites using Git/SVN/Mercurial ?
  - ... deploys them on the server using these tools (e.g. git pull)?
- ... has a MacOS based system?
  - ... deploys using rsync/scp/(s)ftp ?
- ... develops using Sublime Text and the 'SFTP'-Plugin?

- ... or just wants to pwn those people's servers?

Attention!

Interesting files - Part I

- ... develops websites using Git/SVN/Mercurial ?
  - ... deploys them on the server using these tools (e.g. git pull)?

.git directories (1)

- VCS developed by Linus Torvalds
- Commands: git init / add / commit / push / pull / ...
- Data is stored in the .git directory

```
gehaxelt@LagTop /b/g/h/r/p/berlinsides (master)> ls -lha .git
total 20K
drwxr-xr-x 1 gehaxelt gehaxelt  144 May 24 22:24 .
drwxr-xr-x 1 gehaxelt gehaxelt  320 May 24 14:40 ..
drwxr-xr-x 1 gehaxelt gehaxelt    0 May 19 17:58 branches
-rw-r--r-- 1 gehaxelt gehaxelt   15 May 24 14:40 COMMIT_EDITMSG
-rw-r--r-- 1 gehaxelt gehaxelt   92 May 19 17:58 config
-rw-r--r-- 1 gehaxelt gehaxelt   73 May 19 17:58 description
-rw-r--r-- 1 gehaxelt gehaxelt   23 May 19 17:58 HEAD
drwxr-xr-x 1 gehaxelt gehaxelt  414 May 19 17:58 hooks
-rw-r--r-- 1 gehaxelt gehaxelt 1.7K May 24 14:40 index
drwxr-xr-x 1 gehaxelt gehaxelt   14 May 19 17:58 info
drwxr-xr-x 1 gehaxelt gehaxelt   16 May 24 14:40 logs
drwxr-xr-x 1 gehaxelt gehaxelt  100 May 24 14:40 objects
drwxr-xr-x 1 gehaxelt gehaxelt   18 May 19 17:58 refs
```

.git directories (2)

- Objects can be commits, trees and blobs.

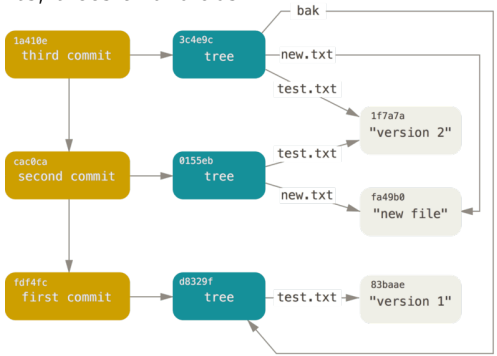

[1]

```
gehaxelt@LagTop /b/g/h/r/p/berlinsides (master)> find .git/objects/ -type f
.git/objects/df/947a00232ca1e0488d32318f81438f602dcf68
.git/objects/ef/ef9e9b3982954994a90ef0b5dc0078a85b1206
.git/objects/7a/228c9c75501e4201acdbf6cc18b07994392aa3
```

[1] Figure https://git-scm.com/book/en/v2/Git-Internals-Git-Objects

.git directories (3)

What if the deployment process is ‘cd /var/www/html && git pull‘?

[1] https://github.com/internetwache/GitTools

.git directories (3)

What if the deployment process is 'cd /var/www/html && git pull'?

The /.git/ folder might be accessible!

[1] https://github.com/internetwache/GitTools

.git directories (3)

What if the deployment process is 'cd /var/www/html && git pull'?
The /.git/ folder might be accessible!

Directory listing enabled

- It's trivial to download all object files and restore the repository.
- wget –mirror –include-directories=/.git http://domain.tld/.git/

[1] https://github.com/internetwache/GitTools

.git directories (3)

What if the deployment process is 'cd /var/www/html && git pull'?
The /.git/ folder might be accessible!

Directory listing enabled

- It's trivial to download all object files and restore the repository.
- wget –mirror –include-directories=/.git http://domain.tld/.git/

Directory listing disabled

- Obtain first hash (.git/HEAD, .git/refs/heads/master)
- Download object file and get new object hashes
- Repeat until nothing new is found!
- Automation: GitTools[1]

[1] https://github.com/internetwache/GitTools

.git directories (4)

Demo!

.git directories (5)

Consequences

- Source code disclosure
  - Get the source and find other vulns ;-)
  - Find committed credentials and escalate privileges.

[1] https://en.internetwache.org/dont-publicly-expose-git-or-how-we-downloaded-your-websites-sourcecode-an-analysis-of-alexas-

.git directories (5)

Consequences

- Source code disclosure
  - Get the source and find other vulns ;-)
  - Find committed credentials and escalate privileges.
- In some cases .git/config contains HTTP-BasicAuth credentials
  - Instant access to company's repositories (e.g. GitLab / GitHub / ...)
  - Access to the CI (e.g. GitLabCI): Build scripts and auto-deployment may lead to server pwnage

[1] https://en.internetwache.org/dont-publicly-expose-git-or-how-we-downloaded-your-websites-sourcecode-an-analysis-of-alexas-

.git directories (5)

Consequences

- Source code disclosure
  - Get the source and find other vulns ;-)
  - Find committed credentials and escalate privileges.
- In some cases .git/config contains HTTP-BasicAuth credentials
  - Instant access to company's repositories (e.g. GitLab / GitHub / ...)
  - Access to the CI (e.g. GitLabCI): Build scripts and auto-deployment may lead to server pwnage
- A scan[1] showed: ~10k out of Alexa's Top 1M are affected.
  - ~250 had HTTP-BasicAuth

[1] `https://en.internetwache.org/dont-publicly-expose-git-or-how-we-downloaded-your-websites-sourcecode-an-analysis-of-alexas-`

Other VCS

Other VCS can be affected, too!

- Subversion
- Mercurial
- ...

Attention!

Interesting files - Part II

- ... has a MacOS based system?
  - ... deploys using rsync/scp/(s)ftp ?

.DS_Store files (1)

```
gehaxelt@LagTop /b/g/h/r/p/b/d/ds_store (master)> file samples/.DS_Store.ctf
samples/.DS_Store.ctf: Apple Desktop Services Store
```

- Apple's proprietary Desktop Service Store format[1] on MacOS.
- Holds meta information (e.g. icons, file name, attributes) about files in a directory.
- Hidden and automatically created when entering a directory with 'Finder'.

[1] https://en.wikipedia.org/wiki/.DS_Store

.DS_Store files (2)

- Header contains magic byte, ‘checksum‘, location of ‘root block‘
- Root block holds structural information
  - Offsets to leaf nodes
  - Tables of content
  - Free lists

Header:

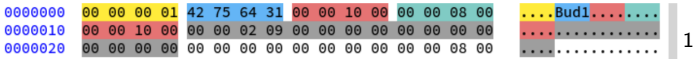

[1]

[1] https://0day.work/parsing-the-ds_store-file-format/

.DS_Store files (2)

- Header contains magic byte, ‘checksum‘, location of ‘root block‘
- Root block holds structural information
  - Offsets to leaf nodes
  - Tables of content
  - Free lists
- Offsets
  - Encoded address and size of a data block

Header:

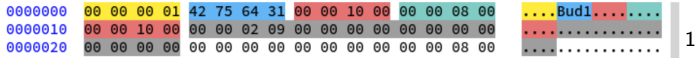

Offsets:

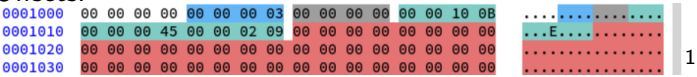

[1] https://0day.work/parsing-the-ds_store-file-format/

.DS_Store files (2)

- Header contains magic byte, ‘checksum‘, location of ‘root block‘
- Root block holds structural information
  - Offsets to leaf nodes
  - Tables of content
  - Free lists
- Offsets
  - Encoded address and size of a data block

Header:

```
0000000  00 00 00 01 42 75 64 31 00 00 10 00 00 00 08 00  ....Bud1........
0000010  00 00 10 00 00 00 02 09 00 00 00 00 00 00 00 00  ................
0000020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 08 00  ................
```

Offsets:

```
0001000  00 00 00 00 00 00 00 03 00 00 00 00 00 00 10 0B  ................
0001010  00 00 00 45 00 00 02 09 00 00 00 00 00 00 00 00  ...E............
0001020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
0001030  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
```

ToC:

```
0001400  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 01  ................
0001410  04 44 53 44 42 00 00 00 01 00 00 00 00 00 00 00  .DSDB...........
```

- Tables of content
  - Usually ‘DSDB‘
  - Block IDs as the index for the Offset list

[1] https://0day.work/parsing-the-ds_store-file-format/

.DS_Store files (3)

- Tree root
  - First data block ID
  - # of internal blocks
  - # of records
  - # of blocks

Tree root:

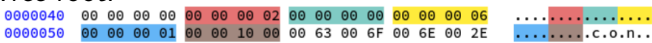

[1]

[1] https://0day.work/parsing-the-ds_store-file-format/

.DS_Store files (3)

- Tree root
  - First data block ID
  - # of internal blocks
  - # of records
  - # of blocks
- Data block
  - Block mode
  - Number of Records
- Record
  - File name's length
  - UTF-16 file name
  - Structure ID
  - Structure type
  - # bytes to skip

Tree root:

```
0000040  00 00 00 00 00 00 00 02 00 00 00 00 00 00 00 06  ....|....|....|....
0000050  00 00 00 01 00 00 10 00 00 63 00 6F 00 6E 00 2E  ....|...|.c.o.n..
```

1

Data block:

```
0000200  00 00 00 00 00 00 00 00 00 00 00 06 00 00 00 0B  ....|....|....|....
0000210  00 66 00 61 00 76 00 69 00 63 00 6F 00 6E 00 2E  .f.a.v.i.c.o.n..
0000220  00 69 00 63 00 6F 49 6C 6F 63 62 6C 6F 62 00 00  .i.c.oIlocblob..
0000230  00 10 00 00 00 46 00 00 00 28 FF FF FF FF FF FF  .....F...(......
0000240  00 00 00 00 00 04 00 66 00 6C 00 61 00 67 49 6C  .......f.l.a.gIl
0000250  6F 63 62 6C 6F 62 00 00 00 10 00 00 02 E4 00 00  ocblob..........
0000260  00 28 FF FF FF FF FF FF 00 00 00 00 00 06 00 73  .(.............s
0000270  00 74 00 61 00 74 00 69 00 63 49 6C 6F 63 62 6C  .t.a.t.i.cIlocbl
0000280  6F 62 00 00 00 10 00 00 00 CC 00 00 00 28 FF FF  ob.............(..
0000290  FF FF FF FF 00 00 00 00 00 09 00 74 00 65 00 6D  ...........t.e.m
00002A0  00 70 00 6C 00 61 00 74 00 65 00 73 49 6C 6F 63  .p.l.a.t.e.sIloc
00002B0  62 6C 6F 62 00 00 00 10 00 00 01 52 00 00 00 28  blob.....R...(
00002C0  FF FF FF FF FF FF 00 00 00 00 00 0D 00 76 00 75  .............v.u
00002D0  00 6C 00 6E 00 65 00 72 00 61 00 62 00 6C 00 65  .l.n.e.r.a.b.l.e
00002E0  00 2E 00 70 00 79 49 6C 6F 63 62 6C 6F 62 00 00  ...p.yIlocblob..
00002F0  00 10 00 00 01 D8 00 00 00 28 FF FF FF FF FF FF  .........(......
0000300  00 00 00 00 00 0F 00 76 00 75 00 6C 00 6E 00 65  .......v.u.l.n.e
0000310  00 72 00 61 00 62 00 6C 00 65 00 2E 00 77 00 73  .r.a.b.l.e...w.s
0000320  00 67 00 69 49 6C 6F 63 62 6C 6F 62 00 00 00 10  .g.iIlocblob....
0000330  00 00 02 5E 00 00 00 28 FF FF FF FF FF FF 00 00  ...^...(........
0000340  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
```

1

[1] https://0day.work/parsing-the-ds_store-file-format/

.DS_Store files (4)

What if the deployment process is ‘scp / rsync / ftp ./code/ server:/var/www/html/‘?

[1] https://github.com/internetwache/Python-dsstore
[2] http://github.com/internetwache/ds_storescanner

.DS_Store files (4)

What if the deployment process is 'scp / rsync / ftp ./code/ server:/var/www/html/'?

All files, including .DS_Store, are transfered and exposed!

[1] https://github.com/internetwache/Python-dsstore

[2] http://github.com/internetwache/ds_storescanner

.DS_Store files (4)

What if the deployment process is 'scp / rsync / ftp ./code/ server:/var/www/html/'?

All files, including .DS_Store, are transfered and exposed!

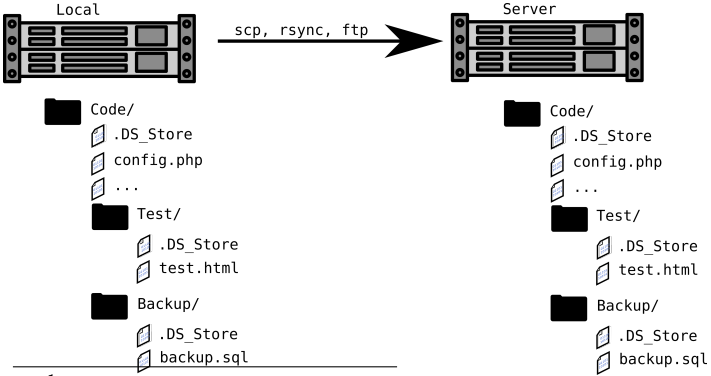

- I've developed tools for parsing[1] and recursively enumerating/checking[2] referenced files.

[1] https://github.com/internetwache/Python-dsstore

[2] http://github.com/internetwache/ds_storescanner

.DS_Store files (5)

Demo!

.DS_Store files (6)

Consequences

- Directory listing ‘bypass‘

[1] https://en.internetwache.org/scanning-the-alexa-top-1m-for-ds-store-files-12-03-2018/

.DS_Store files (6)

Consequences

- Directory listing ‘bypass‘
- Disclosure of (probably) accessible files on the server.
  - Backup files
  - Database files
  - Temporary files
  - Key files

[1] https://en.internetwache.org/scanning-the-alexa-top-1m-for-ds-store-files-12-03-2018/

.DS_Store files (6)

Consequences

- Directory listing ‘bypass‘
- Disclosure of (probably) accessible files on the server.
  - Backup files
  - Database files
  - Temporary files
  - Key files
- A scan[1] showed: ~10k from Alexa's Top 1M are affected
  - >850k URLs with response HTTP 200/OK
  - Files: .bak, .gz, .db, .eml, .old, .inc, .config, .sql, .pem, ...

[1] https://en.internetwache.org/scanning-the-alexa-top-1m-for-ds-store-files-12-03-2018/

Attention!

Interesting files - Part III

- ... develops using Sublime Text and the ‘SFTP‘-Plugin?

sftp-config.json (1)

- The ‘Sublime SFTP‘[1] plugin manages file uploads to a server
- Configuration file contains (S)FTP and/or SSH credentials
- Maintainer: ‘Plugin excludes it from file transfers.‘

[1] `https://wbond.net/sublime_packages/sftp`

sftp-config.json (1)

- The 'Sublime SFTP'[1] plugin manages file uploads to a server
- Configuration file contains (S)FTP and/or SSH credentials
- Maintainer: 'Plugin excludes it from file transfers.'

- What if the files are uploaded using another channel?

[1] `https://wbond.net/sublime_packages/sftp`

sftp-config.json (1)

- The 'Sublime SFTP'[1] plugin manages file uploads to a server
- Configuration file contains (S)FTP and/or SSH credentials
- Maintainer: 'Plugin excludes it from file transfers.'

- What if the files are uploaded using another channel?

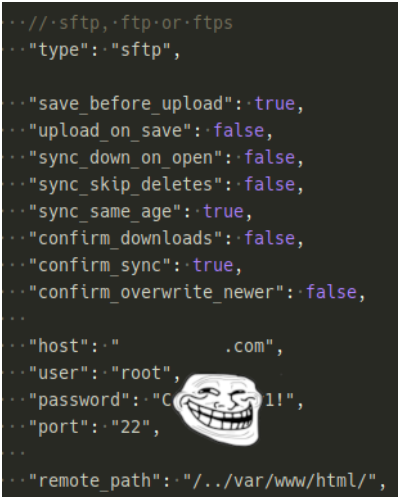

[1] https://wbond.net/sublime_packages/sftp

sftp-config.json (2)

No Demo :(

sftp-config.json (3)

Consequences

- Instant FTP or shell access to a server.

sftp-config.json (3)

Consequences

- Instant FTP or shell access to a server.
- Sometimes even as root!

sftp-config.json (3)

Consequences

- Instant FTP or shell access to a server.
- Sometimes even as root!
- A scan showed: ~400 hits on the Alexa Top 1M
  - ~300 have username + password
  - 20 'root's

Attention!

Interesting files - Part IV

- ... or just wants to pwn those people's servers?

Other interesting files

- There are a LOT more files with sensitive content!
- Be motivated to ~~exploit~~ explore them!

Examples

- .svn/
- .idea/
- .swp
- .old
- .htpasswd
- coredump
- wsftp.ini
- winscp.ini
- filezilla.xml
- domain.tld.key

Tools

- Snallygaster[1]
- Bfac[2]
- GitTools[3]
- DS_StoreScanner[4]

[1] https://github.com/hannob/snallygaster/
[2] https://github.com/mazen160/bfac
[3] https://github.com/internetwache/GitTools
[4] https://github.com/internetwache/ds_storescanner

Attention!

Scanning for files

General tipps

- Dataset: Alexa Top 1 Million[1]
- User-Agents: Firefox/Chromium **not** curl or wget or python-requests!
- Limits: DNS-Requests, Bandwidth, Abuse-Mails

[1] `http://s3-us-west-1.amazonaws.com/umbrella-static/top-1m.csv.zip`

General tipps

- Dataset: Alexa Top 1 Million[1]
- User-Agents: Firefox/Chromium **not** curl or wget or python-requests!
- Limits: DNS-Requests, Bandwidth, Abuse-Mails

```
This warning is from the Financial Security Institute(FSI) of Korea.

Our job is to protect Korean financial organizations from illegal intrusion attacks.
We have received a report of unauthorized access trial originating from your site as shown below.

--------------------------------------------------------------------------------
Date/Time(GMT+9): 2017/03/31 07:00:49 ~ 2017/03/31 08:46:40
Source IP : 207.154.202.22
Destination IP :
124.243.61.13,175.126.235.30,203.229.168.2,203.229.168.79,203.229.175.89,203.233.85.2
Attack Type : F-SCN-WEB-170305-GitRepository_scan_attempt
--------------------------------------------------------------------------------
```

[1] http://s3-us-west-1.amazonaws.com/umbrella-static/top-1m.csv.zip

A distributed tool

- Celery Task Queue[1]
  - Backend: Redis
  - Message broker: RabbitMQ
- 1 master node & *n* slave nodes
- Not finished and released yet, but hopefully soon on GitHub[2].

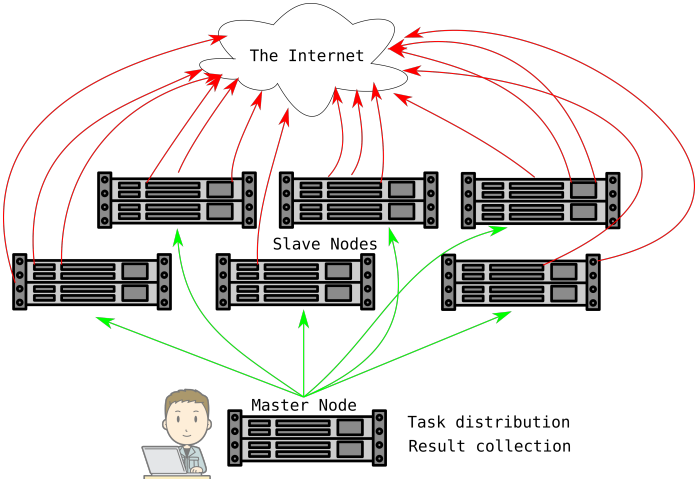

[1] http://www.celeryproject.org/
[2] http://github.com/gehaxelt/

F||A&Q

Feedback || Answers && Questions?

@gehaxelt
contact@0day.work
... or talk to me :-)