

# Radboud Universiteit



**The state of OSS-Fuzz**

Assessing statefulness in OSS-Fuzz: performance and results

**Lars Esselink**

Supervisors:  
Erik Poll  
Bram Westerbaan

Second readers:  
Cristian Daniele  
Seyed Behnam Andarzian

Faculty of Science  
Computing Science  
Radboud University  
June 26, 2024

## Abstract

OSS-Fuzz is a project by Google that fuzzes open-source software on their data centers, to improve the reliability of open-source software. The underlying fuzzers of OSS-Fuzz are libFuzzer, AFL++, and Honggfuzz. We expected OSS-Fuzz to be ineffective because these fuzzers are not designed to deal with stateful software. Using ‘stateless’ fuzzers on stateful software might be faster in executions but it might also be worse at finding bugs because stateless fuzzers that use fuzzing targets effectively circumvent the state transitions and thus might not reach bugs that are unreachable. That is why we want to explore OSS-Fuzz to see if it gives results in stateful software effectively or is wasting resources by fuzzing stateful software, with a stateless fuzzer. We try to answer this question by comparing bugs between a stateful fuzzer and OSS-Fuzz, using open-source software that has known vulnerabilities.

We picked two open-source projects, ProFTPD, and open62541, which have versions with known vulnerabilities, and we fuzzed them with AFLNet, a stateful fuzzer, and OSS-Fuzz to see if OSS-Fuzz is wasting resources. Experimenting with ProFTPD showed us that AFLNet did find more bugs, but that was not because OSS-Fuzz was not good enough, but because the fuzzing targets created by the developers of ProFTPD were not good enough. AFLNet did not find any bugs in open62541 that we could reproduce. OSS-Fuzz, however, did find bugs. We do not know if this bug could be found by AFLNet because the malicious input would have been caught elsewhere in the program, which prevents the input from reaching the part that would have crashed (which might not happen when fuzzing using OSS-Fuzz).

We cannot conclude that OSS-Fuzz is wasting resources, but can conclude more research is needed. To get a more exact answer, we would require more time and more case studies to see if the results we found are flukes or a repeating pattern. It is not fair to compare AFLNet and OSS-Fuzz because OSS-Fuzz is heavily dependent on good fuzzing targets. When looking at ProFTPD we found that the lack of good fuzzing targets lowered the amount of bugs found while open62541 had good fuzzing targets and outperformed AFLNet.

OSS-Fuzz might find fewer bugs than AFLNet, but it does substantially more executions than AFLNet (30,000 vs 5) because AFLNet fuzzes the socket entry point which in turn uses the whole code while OSS-Fuzz only fuzzes parts of the code. This might make OSS-Fuzz better at finding bugs, but whenever a bug is not reachable by OSS-Fuzz (because the fuzzing targets do not cover enough code) they can be faster but they will never find the bug.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Fuzzing</b>	<b>5</b>
2.1	Stateless software & fuzzers . . . . .	5
2.1.1	AFL++ . . . . .	5
2.1.2	Honggfuzz and libFuzzer . . . . .	5
2.1.3	Reproducibility . . . . .	5
2.2	Stateful fuzzers . . . . .	6
2.2.1	AFLNet . . . . .	6
2.2.2	Socket vs Target . . . . .	7
2.3	OSS-Fuzz . . . . .	7
2.4	ASan and UBSan . . . . .	7
<b>3</b>	<b>Related Work</b>	<b>9</b>
3.1	Bugs found by OSS-Fuzz . . . . .	9
<b>4</b>	<b>Method</b>	<b>10</b>
4.1	Project selection . . . . .	10
4.2	Seed creation . . . . .	10
4.3	Fuzzing instructions . . . . .	12
4.3.1	AFLNet instructions . . . . .	12
4.3.2	OSS-Fuzz instructions . . . . .	13
4.4	Comparing results found by fuzzing . . . . .	13
<b>5</b>	<b>Fuzzing ProFTPD</b>	<b>15</b>
5.1	Fuzzing ProFTPD with AFLNet . . . . .	15
5.2	Fuzzing ProFTPD with OSS-Fuzz . . . . .	17
5.2.1	LibFuzzer . . . . .	17
5.2.2	AFL++ . . . . .	18
5.2.3	Honggfuzz . . . . .	18
5.3	Comparing AFLNet and OSS-Fuzz over ProFTPD . . . . .	19
<b>6</b>	<b>Fuzzing open62541</b>	<b>20</b>
6.1	Fuzzing open62541 with AFLNet . . . . .	20
6.2	Fuzzing open62541 with OSS-Fuzz . . . . .	22
6.2.1	LibFuzzer . . . . .	22
6.2.2	AFL++ . . . . .	24
6.2.3	Honggfuzz . . . . .	25
6.3	Comparing AFLNet and OSS-Fuzz over open62541 . . . . .	26
<b>7</b>	<b>Future work</b>	<b>27</b>
7.1	More stateful fuzzers . . . . .	27
7.2	More case studies . . . . .	27
7.3	Better fuzzing targets . . . . .	27

<b>8 Conclusion</b>	<b>28</b>
<b>Glossary</b>	<b>30</b>
<b>A Results</b>	<b>32</b>
A.1 ProFTPD in AFLNet . . . . .	32
A.2 ProFTPD in AFL++ . . . . .	32
A.3 ProFTPD in Honggfuzz . . . . .	33
A.4 Open62541 in AFLNet . . . . .	33
A.5 Open62541 in AFL++ . . . . .	34
<b>B Installation</b>	<b>36</b>
B.1 ProFTPD . . . . .	36
B.2 OSS-Fuzz . . . . .	36
B.2.1 Instructions for the OSS-Fuzz helper script . . . . .	36
B.3 AFL++ in OSS-Fuzz helper script . . . . .	37
B.4 Honggfuzz in OSS-Fuzz helper script . . . . .	37
<b>C Bugs</b>	<b>38</b>
C.1 ProFTPD . . . . .	38
C.2 Open62541 . . . . .	38
<b>D Source code</b>	<b>41</b>
D.1 fuzz_mdns_message.cc . . . . .	41
D.2 fuzz_binary_decode.cc . . . . .	42

# 1 Introduction

When we refer to stateless software, we mean software that does not maintain or track any state information. Let us take for example an implementation of FFmpeg which is stateless software that converts audio/video files to other types. It does not keep track of any states. But we also have stateful software, which is the opposite of stateless software. Stateful software does keep track of states, for example, an implementation of FTP (file transfer protocol). FTP has to be in specific states to do different commands. To download or upload files, you will need to be authenticated and thus be in the authenticated state.

Fuzzing is a popular technique used for testing software[14]. The advantage of fuzzing is that is almost fully automated. For normal test cases, you will need to write tests for each new code you write, while fuzzing only requires a one-time code creation (for example a fuzzing target) and thus saves time when trying to test your software. Another advantage is that it might catch edge cases that normal test cases would not catch. The downside of fuzzing is that you will need more computing power than normal test cases.

As long as you have computing power, you can test your software for every edge case and increase the reliability of your software. This is where Google OSS-Fuzz[8] comes in. Google created the OSS-Fuzz project to help improve open-source software. It takes open-source projects and fuzzes them using their data centers. Interestingly, OSS-Fuzz does not use any stateful fuzzer but fuzzes stateful software, and thus might not be as good at finding bugs as AFLNet at fuzzing stateful software.

In this thesis, we want to answer the following research question: is OSS-Fuzz wasting resources by fuzzing stateful software projects with stateless fuzzers? We investigate the matter by looking at a specific version of two stateful software projects with a known bug and fuzzing it with OSS-Fuzz's stateless and stateful fuzzers. If a stateful fuzzer can but OSS-Fuzz cannot find these bugs, it indicates that OSS-Fuzz is wasting resources. This is why we want to test it in an OSS-Fuzz environment to get the closest possible simulation of how OSS-Fuzz is doing it.

In Chapter 2, we discuss what fuzzing exactly is and the terms we are not on one line with. In Chapter 3, we look at related research. In Chapter 4, we discuss the method we use to get the results. Chapters 5 and 6 explain the results of fuzzing. In Chapter 7 we write down what we stumbled upon but was not in the scope of our research and can be researched in the future. In Chapter 8 we discuss the findings of the whole research, what we found, and explain the conclusions. In the last Chapter 8 we have a glossary for terms we use in this research.

## 2 Fuzzing

In this chapter, we discuss what fuzzing exactly is and the terms we are not on one line with. We also discuss the different fuzzing aspects we are not familiar with. Some abbreviations are not explained here but are written out in Chapter 8.

Fuzzing is a method to find vulnerabilities in software. This is done by sending many, (semi) automatically generated inputs to the software and seeing if it crashes or hangs. The mutation of messages is often guided based on multiple factors, like code coverage or the number of crashes. Fuzzers can be categorized into stateful and stateless groups, based on if they are aware of the states inside the software.

### 2.1 Stateless software & fuzzers

Software that does not have any internal states is called stateless, an example of stateless software is a browser that implements HTTP. The HTTP protocol does not keep track of any states and is therefore stateless. The fuzzers created for stateless software do not consider the states the software is in when making decisions.

#### 2.1.1 AFL++

AFL++[5], the daughter of AFL, is a coverage-guided fuzzer. The coverage guide works by looking at the coverage of the current mutation in the input and decides based on the coverage what mutation should be done next. The fuzzer also listens to the feedback, which is the part where it makes mutations based on the information it gets from the instrumentation (like crashes/hangs). The coverage guide and the feedback guide both get their information from the instrumentation, but the feedback guide looks at the program metrics (like memory usage), crashes, and hangs. You normally use standard command line input (stdin for short) to fuzz using AFL++, but in this research, we use a special C library to get it working on the fuzzing targets of libFuzzer.

#### 2.1.2 Honggfuzz and libFuzzer

Honggfuzz and libFuzzer[13] are coverage-guided fuzzers. They use so-called fuzzing targets, specific functions especially made for fuzzers inside the SUT, to execute specific functions in code (as if those functions are called by the program itself). They then check the code coverage and generate more mutations, maximizing code coverage. The difference between Honggfuzz and libFuzzer is minimal and both use the same principle but in different coding.

#### 2.1.3 Reproducibility

Fuzzing has a problem when it comes to reproducibility (also known as flakiness), especially fuzzing targets. Let us take a look at the simple state model

in Figure 1. Since fuzzing targets focus on specific parts of the code, the errors found work only in a specific state of the program. For example, when there is a bug found in fuzzing target 1, which operates in state 4, we can only reproduce that error when using that fuzzing target. When we want to reproduce the error on the original program, we need to construct a message that takes us from state 1 to state 4, using actions a, b, and c. This construction would be feasibly impossible to do and thus the error found with the fuzzing target would be impossible to reproduce outside the fuzzing target.




Figure 1: Simple state model

## 2.2 Stateful fuzzers

Software is generally stateful, meaning that during the execution of software, it can be in different states. Take the server side of FTP (file transfer protocol) for example, it has an idle state in which it can receive connections, but it also has a connected state where a client is connected and can send commands to the server. Stateless fuzzers do have some guiding mechanisms, like code coverage or instrumentation feedback, but none take states into account. With the example of FTP, a stateless fuzzer would fuzz FTP, but would not find the other state (It would be able to find other states when run for infinity). This is why we have stateful fuzzing. Stateful fuzzers also have state feedback that is taken into account when mutating input.

### 2.2.1 AFLNet

AFLNet[11] is a stateful fuzzer used to fuzz network protocols. It uses state feedback and measures code coverage. The state feedback works by receiving messages from the SUT and learning how the state model of the SUT is built. This feedback is then used to create a mutation of the input to trigger state transitions or stay in a specific state. This state feedback is also used to increase code coverage. It is used to fuzz stateful software (explained in Section 2.2) but

can also be used to fuzz stateless software (explained in Section 2.1) but it would run slower than stateless fuzzers.

### 2.2.2 Socket vs Target

One of the differences between AFLNet and OSS-Fuzz is the difference between socket vs target fuzzing (visually explained in Figure 2). Socket fuzzing, the method of AFLNet, fuzzes the program interacting with it via a socket. The program is fuzzed using normal operations and tries to find it as if were a normal user. Fuzzing target, the method of libFuzzer, AFL++ (in OSS-Fuzz), and Honggfuzz, use fuzzing targets that focus on pieces of code within the program. It does not fuzz the program as a whole but works more like unit testing, where different sections of code are tested. Both have pros and cons (as discussed in 2.1.3).

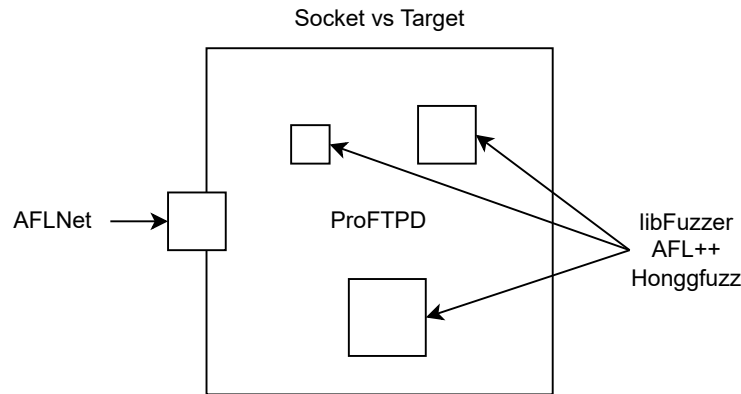


Figure 2: The difference between socket and target fuzzing

### 2.3 OSS-Fuzz

OSS-Fuzz[8] is a project by Google to fuzz open-source software using their data centers. The fuzzers included are libFuzzer[13], AFL++[5] and Honggfuzz[7]. None of these fuzzers specifically support stateful software. Google does not state why it chose to not include stateful fuzzers, but one assumption would be that using stateless fuzzers is less work for the OSS-Fuzz team. Stateless fuzzers only require fuzzing targets, while stateful fuzzers require a method of extracting the state, to use state feedback.

### 2.4 ASan and UBSan

ASan[9] (Address sanitizer) and UBSan[6] (Undefined Behavior sanitizer) are compiler flags to add checks in the compiled code for memory violations and undefined actions. ASan checks for common memory errors such as using memory



after freeing it, memory leaks, etc. For example, whenever we use the function `memcpy` in the C language, we copy a block of memory to another block of memory. But this operation does not check bounds, and thus whenever we go over bounds it is not detected. When run without ASan memory faults might go unnoticed since it might overflow in memory that is not used, but with ASan this is immediately noticed and the ASan crashes the program. UBSan checks at run time for undefined behavior, for example whenever we use a non-static variable in the C language before it is declared, will give undefined behavior because of the nature of the C language. Just like memory violation, this might go unnoticed as the effect is not catastrophic, but with UBSan this will be immediately noticed and UBSan will crash the program. ASan and UBSan, in our case, will help with finding bugs more easily since every memory violation and undefined behavior will immediately show up and count as a crash and thus be detected by the fuzzer.

### 3 Related Work

In Section 3.1, we present the results of Yu Ding and Le Goues[3], which classifies the most bugs found in OSS-Fuzz.

#### 3.1 Bugs found by OSS-Fuzz

OSS-Fuzz had its first launch in 2016. As of August 2023, it has fuzzed 1000 projects, found 36000 bugs, and identified over 10000 vulnerabilities [8]. Ding et al.[3] explored all the bugs and performed an empirical study. They analyzed over 4 years of data and 23,907 bugs. These bugs were categorized under 16 categories found in Figure 3. They were also categorized into fixed and unfixed and flakiness (reproducibility, whether a bug is reproducible or not). The biggest category is timeout. But they did not categorize based on stateless or stateful. This is what we planned on doing, albeit not that many projects, to see the difference between stateless and stateful fuzzers.




Figure 3: bugs found in OSS-Fuzz categorized in 15 categories[3]

## 4 Method

This chapter discusses an overview of the different steps to take to find out if OSS-Fuzz is wasting its resources with 2 case studies.

### 4.1 Project selection

First, we need to select a project to test our hypothesis on. We look at the list of already fuzzed projects<sup>1</sup> and manually filter it on if the project uses sockets. This list now contains projects that have sockets but are not necessarily stateful, for example, the project ffmpeg which does not have any states, and the socket use is optional. So now we manually apply a second filter, which filters on protocols that were specifically made for networking (as they are almost always stateful) and we get the following list as shown in Table 1. From this list, we can pick a suitable candidate. We chose ProFTPD and open62541 because both protocols (OPCUA and FTPD) already had an implementation in AFLNet which was easier for us to use and thus we did not have to create it ourselves.

bitcoin-core	libvnc
bluez	mosquitto
dropbear	msquic
fast-dds	open5gs
freeradius	open62541 (opcUA)
freerdp	opendnp3
gnutls	openssl
hiredis	openweave
irssi	pidgin
libgit2	pupnp
libmodbus	qpuid-proton
librdkafka	resiprocate
libssh	wolfssl
libssh2	wpantund
libtorrent	proftpd

Table 1: The list of projects after we applied 2 filters.

### 4.2 Seed creation

We are going to need seeds for fuzzers to work. This is not required but it might be better to give a starting point instead of randomly starting. These seeds are input strings that are used by the fuzzers and fed into the SUT. These seeds are also mutated to test different inputs. OSS-Fuzz does not require us to provide the seeds, because they are either provided by the project owner or not needed. For AFLNet however, we need to make the seed ourselves (as OSS-Fuzz uses seeds in another format). This is done using the normal operation of a SUT and tcpdump. Tcpdump records all traffic within the specified parameters, which lets us use that same traffic as input seed for AFLNet. For example, let us take

<sup>1</sup><https://github.com/google/oss-fuzz/tree/master/projects>

open62541. First, we need to use the SUT to create traffic, such that tcpdump can capture it. In this example, we use:

```
tcpdump -w traffic.pcap -i lo port 4840
```

This starts tcpdump and tells it to write to the file called “traffic.pcap” and listen specifically to the loopback network interface on port 4840 (the port used by open62541 tools). After which you perform the normal operation of the SUT and close tcpdump. Now we have a file called “traffic.pcap” which holds the traffic of the program. We can read this with Wireshark, which we can find in Figure 4.

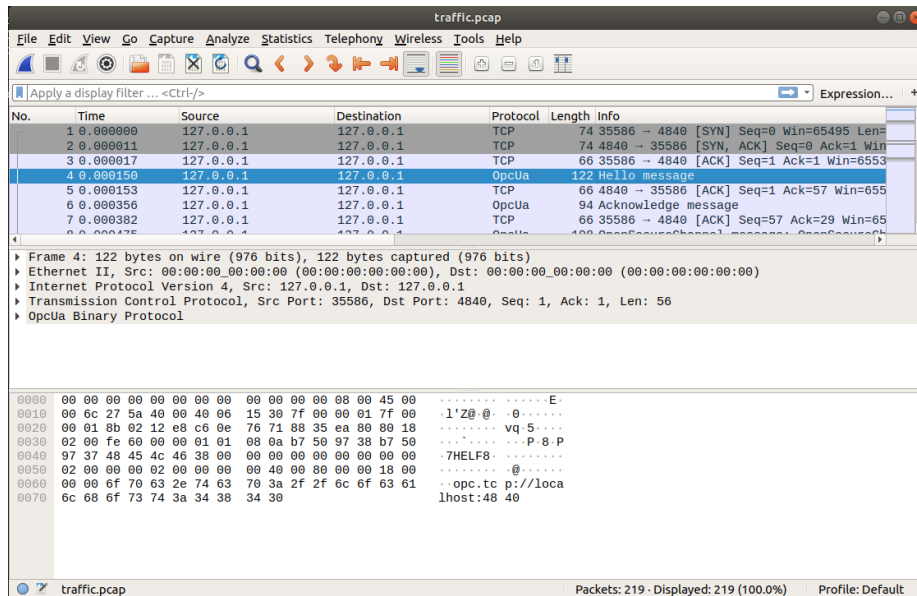


Figure 4: The content of traffic.pcap using Wireshark

In Figure 4 you see network packets. From this, we can view the TCP stream (traffic of a session) and view the traffic between the client and server of open62541. We only care about the traffic between the client and server (as seen in Figure 5), since that is what AFLNet is going to use for fuzzing (and not the server-to-client traffic). From this, we can save it as raw bytes (otherwise it will be sent as ASCII, which the server cannot interpret) and use it as a seed.

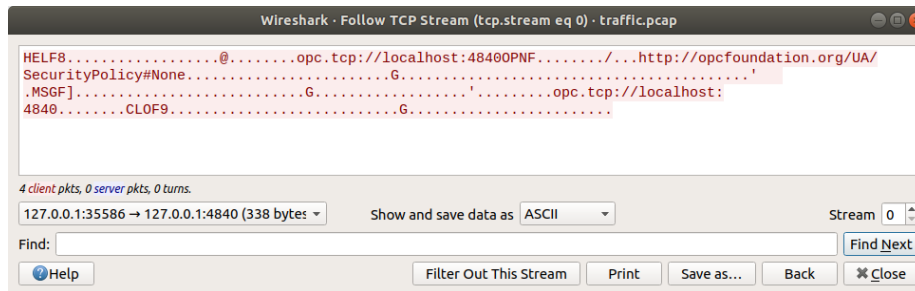


Figure 5: The specific TCP stream of a session between client and server

### 4.3 Fuzzing instructions

After we have the input seeds for AFLNet (and the supplied seeds for OSS-Fuzz), we can proceed to fuzz. AFLNet and OSS-Fuzz work a little bit differently. We use AFLNet in our environment, we do not run it in a pre-supplied way, while we run OSS-Fuzz using their helper script. This script does the setup for the SUT and the fuzzers for us using only commands. We chose to do this because we want to use OSS-Fuzz as close to their environment as possible. Doing it this way makes it easier for us to make conclusions as to why they chose some options (as discussed in Chapter 1) and makes it also easier to see if OSS-Fuzz is wasting resources.

We also need a SUT to test the difference. We pick a stateful SUT from the list of already fuzzed projects by OSS-Fuzz. Once we find a project, we search if the project has a version that has a known bug that can be found by fuzzers. If we find a project that is fuzzed by OSS-Fuzz and has a version with a known bug, we can start fuzzing that project version with AFLNet and OSS-Fuzz.

#### 4.3.1 AFLNet instructions

First, we start with AFLNet. For AFLNet we run it using:

```
afl-fuzz -i /path/to/seed/ -o /path/to/output -N
{tcp/udp}://{IP}/{PORT} -P {PROTOCOL} -q 3 -s 3 -E -K -R
-m none {path/to/executable}
```

You will need to supply the command with the following: the path to the seed, the path to where you want to save the output, does the SUT uses TCP or UDP, The IP and PORT of the SUT, the protocol of the SUT, and the path to the executable. The other parameters can be changed but do not make a big difference. Once it is fuzzing the program, we choose a time frame in which we fuzz the SUTs. We chose 24 hours as this is a reasonable amount of time to find a crash/hang, after which we manually stopped AFLNet. The results are stored in the output directory we defined.

### 4.3.2 OSS-Fuzz instructions

For OSS-Fuzz, we need to use the helper script. This script helps us build the SUT, build the fuzzers and run the fuzzers. First, building the SUT:

```
python3 /path/to/helper.py build_image {PROJECT}
```

The building of the SUT is reasonably straightforward, you just need to supply the project and it will build the SUT. After building the SUT, we need to build the fuzzer:

```
python3 /path/to/helper.py build_fuzzers --engine  
{libfuzzer,afl,honggfuzz} {PROJECT}
```

To build the fuzzers you need to supply the fuzzer you want to run and the project you want to fuzz. Once finished, we execute the last step, running the fuzzer:

```
python3 /path/to/helper.py run_fuzzer --corpus-dir  
/path/to/output/ --engine {libfuzzer,afl,honggfuzz}  
{PROJECT} {FUZZER TARGET}
```

To run the fuzzer you need to supply the script with an output directory, the fuzzer you want to fuzz with, the project, and one of the pre-supplied fuzzing targets. After executing this it will start fuzzing. AFL++ and Honggfuzz will run until stopped, while libFuzzer will stop once it finds a hang/crash. So we run AFL++ and Honggfuzz also 24 hours, just like AFLNet. We can then find the results in the output directory we defined.

## 4.4 Comparing results found by fuzzing

When comparing results we do not have a global method, since we do not yet know what we get for results. For a single bug found by a single fuzzer, we cannot compare it to anything so we do not have a specific method we use. When multiple fuzzers find a bug, we can use debug to find where the crash/hang came from.

Using a debugger like GDB, we can inspect the execution step by step to find out where the program crashes. GDB works by setting so-called breakpoints in the execution of the program. These breakpoints cause the execution of the code to pause at the breakpoint. For example, if you set a breakpoint at a function, the debugger loads the functions onto the stack but does not yet execute it and pauses the execution until further notice. Let us take for example the source code for the MDNS fuzzing target (found in Figure 19 on page 41).

1. First we analyze the source code and set breakpoints at important functions, in this case at `message_parse`, `mdnsd_new` and `mdnsd_in`.
2. After the breakpoints are set, we run the program with an input. The debugger stops at the first breakpoint (or crashes if there is a crash before the breakpoint). We then choose to “continue” which continues to the next crash/breakpoint, “step” which takes one step in execution, or “finish” which finishes the function it is currently at.
3. Repeat steps one and two until you find a crash. After which you can dive deeper into a function, set breakpoints at more functions, and analyze the results until you find the exact line of code that crashes the program

## 5 Fuzzing ProFTPD

For our experiment, we use a version of each SUT with a vulnerability that can be found using fuzzing. We compare AFLNet and OSS-Fuzz (with the underlying fuzzers) results.

For the first case, we use ProFTPD. ProFTPD is an open-source FTP implementation. FTP is used to transfer files from the client to the server and vice versa. ProFTPD is a highly configurable and secure implementation of FTP. One of the versions of ProFTPD[2] has an out-of-bounds vulnerability which can be found using AFLNet (commit 5e57d41)<sup>2</sup>, using ASan and UBSan. So we are going to test if for this version OSS-Fuzz can find the same crashes and hangs as AFLNet. In Section 5.1 we discuss the experiments we ran on ProFTPD using AFLNet and analyze those results. In Section 5.2 we talk about the experiments on ProFTPD using OSS-Fuzz and analyze the results. In Section 5.3 we take the results from Section 5.1 and 5.2, and we compare them to see if there is a difference in bugs found and execution speed.

### 5.1 Fuzzing ProFTPD with AFLNet

**Run 1** We first compile ProFTPD without ASan and UBSan. We then tried to fuzz commit 5e57d41, but we did not get any results. As you can see in Table 7 on page 32, there are no crashes or hangs found, and it is very slow. The low amount of executions is because AFLNet fuzzes the program as a whole using the socket. When the whole program needs to execute to fuzz a small input it results in few executions done per minute. This is different from running fuzzing targets, which only run the targeted code and not the whole program. The lack of crashes or hangs might be because the original issue was found with ASan and UBSan enabled.

**Run 2** We then run AFLNet with ASan and UBSan but AFLNet complained about memory usage when using ASan in combination with a 64-bit OS. It turns out that you need special arguments to run AFLNet with ASan, otherwise on a 64-bit OS you might use too much memory, namely the `m` flag. We managed to compile it with the instructions found in Figure B.1, which immediately gives us more results. After running AFLNet for 47 hours, we got 17 hangs where the last hang was found 24 hours before ending the fuzzer, meaning that it did not find anything new during the last 24 hours of fuzzing. Once we find a specific crash or hang, we can replay (using `afl-replay`) the exact test sequence used before it crashed or hung. We used `afl-replay` for our results and it immediately recognized the ‘flaw’ that caused these hangs. Namely, the ProFTPD configuration allowed for 3 login attempts and after that, it stopped the program. Which AFLNet classified as hang since the ProFTPD stopped. Thus all these hangs were false positives.

---

<sup>2</sup><https://github.com/ProFTPD/ProFTPD/issues/1683>



**Run 3** After changing the ProFTPD configuration to have unlimited login attempts, we ran it again, but this time for 24 hours since the last run did not have any new paths in the last 24 hours. The results (found in Table 2) showed 29 hangs. Unfortunately, the out-of-bound crash was not found.

ProFTPD	
Last new path	3 min
Last unique hang	2 hrs, 21 min
unique hangs	29
Last unique crash	None
unique crashes	0
Total execs	389k
Total paths	2808

Table 2: The third run of ProFTPD using AFLNet, using ASan and UBSan, with the correct configuration, extracted from table 8 on page 32.

Table 2 contains all the information of AFLNet. We can see how long it ran and how long it took until it found a new path/crash/hang. There are also statistics about the fuzzing process, for example, how many cycles it has done, the total paths found and the amount of (unique) crashes and hangs.

**Analyzing run 3** We use the tool that AFLNet has provided (called aflnet-replay) to replay the packet that causes the crash/hang. We start the example server, in this case, we start ProFTPD. Then we use aflnet-replay on all 29 of the unique replay packets. From the 29 hangs, 28 gave a segmentation fault on the aflnet-replay side and the last one did not give us an error. We focus on the 27th packet that caused a segmentation fault when using aflnet-replay. We do not get any information from the client or the server side. Using tcpdump, we can see if there is traffic at all. Traffic is sent, but only one way before crashing. The server side sends the welcome message to the client but nothing is sent by the client. It must be a problem with the first line of the input since nothing is sent. The first few lines are displayed in Figure 6.

```

00000000  0e 00 00 00 55 53 45 52 20 70 72 6f 66 74 70 64 |....USER proftpd|
00000010  0d 0a 0e 00 00 00 50 41 53 53 20 70 72 6f 66 74 |.....PASS prof|
00000020  70 64 0d 0a 06 00 00 00 51 55 49 54 0d 0a 0e 00 |pd.....QUIT....|
00000030  00 00 52 4e 46 52 20 74 65 73 74 2e 74 78 0d 0a |..RNFR test.tx..|
00000040  0e 00 00 00 50 41 53 53 20 70 72 6f 66 74 70 64 |....PASS proftpd|
00000050  0d 0a 0e 00 00 00 53 4d 4e 54 0d 70 72 6f 66 74 |.....SMNT.proft|
00000060  70 64 0d 0a 0e 00 00 00 50 41 53 53 20 70 72 6f |pd.....PASS pro|
00000070  66 74 70 64 0d 0a 08 00 00 00 54 59 50 45 20 41 |ftpd.....TYPE A|

```

Figure 6: The hex dump of the first few lines of the input that hang ProFTPD

The hex dump in Figure 6 does not say much and does not seem to be the problem of ProFTPD but is more likely Ubuntu or aflnet-replay. We do not

know why these segmentation faults are happening and to find out why, we would need to debug the aflnet-replay tool which is out of the scope of this research. We mark it as non-reproducible.

**Github input** We also tested ProFTPD using the input of the original GitHub issue<sup>3</sup>, which caused a heap buffer overflow (found in Appendix C.1). The problem was caused by sending 3722 ‘\’ and ProFTPD having no sanitation rule to sanitize that many wrong characters and thus overflowing in memory. Our AFLNet run, however, did not give us the same result, probably because fuzzing is nondeterministic meaning that it will not come up with the same mutations in the seed every time.

## 5.2 Fuzzing ProFTPD with OSS-Fuzz

We run the same experiments with OSS-Fuzz. We take the vulnerable version and run it with OSS-Fuzz. OSS-Fuzz does not use the whole program as SUT but uses fuzzing targets instead. Which focuses on a specific part of the program (explained in Section 2.2.2).

### 5.2.1 LibFuzzer

We fuzzed ProFTPd with libFuzzer (description found in Section 2.1.2) using the fuzzing targets in OSS-Fuzz. This is a relatively straightforward process. OSS-Fuzz contains a script to start the project to start fuzzing, which needs information about the fuzzer engine and fuzzing target. To submit a project to OSS-Fuzz its necessary to provide fuzzing targets for libFuzzer to use. We use these fuzzing targets to test how well libFuzzer can find bugs. We used the instructions in Appendix B.2.1 to run libFuzzer with ASan. Unfortunately, the OSS-Fuzz helper script does not allow multiple sanitizers at the same time, as we did with AFLNet. Because of this, we ran libFuzzer with address sensitization for 48 hours. Unfortunately, we did not find crashes or hangs. A closer look at the fuzzing target reveals that the target is not covering enough code. As you can see in Table 3, the coverage of the fuzzing target is high, while the coverage of the SUT is low. This means that the fuzzing target only focuses a small part of the code, even though the fuzzing target is fully (functionally) executed. The low code coverage on the SUT is a bad sign. The bug in AFLNet is impossible to find by libFuzzer because of the low code coverage and the fact that the bug we found was not inside the code reached by this fuzzing target. We can create a fuzzing target ourselves, but that defeats the purpose of testing the current OSS-Fuzz environment since we want to know if (without our help) libFuzzer can find the same bugs as AFLNet.

---

<sup>3</sup><https://github.com/ProFTPD/ProFTPD/issues/1683>

PATH	LINE COVERAGE	FUNCTION COVERAGE	REGION COVERAGE
proftpd/	1.26% (740/58550)	2.88% (55/1908)	1.30% (567/43665)
fuzzer.c	88.89% (16/18)	100.00% (1/1)	83.33% (5/6)
<b>TOTALS</b>	<b>1.29% (756/58568)</b>	<b>2.93% (56/1909)</b>	<b>1.31% (572/43671)</b>

Table 3: The code coverage, provided by the OSS-Fuzz helper script, on the ProFTPD fuzzing target

### 5.2.2 AFL++

AFL++ is simpler than libFuzzer. It does not require any fuzzing targets, just an input seed to start fuzzing. It also accepts a dictionary file that specifies the correct input AFL++ can use. The installation of ProFTPD is the same as AFLNet found in Appendix B.1. The only difference is that it uses other instrumentation tools to compile the binaries and after installation, it is run with another input.

**Run 1** The first run gave us no results. As you can see in Table 9 (on page 32), it says “odd, check syntax” after the last new path find. It means that AFL++ cannot find any path, this is a clear sign the fuzzer is not working correctly. It turns out to be because AFL++ does not support sockets in its normal mode, but AFL++ works with STDIN or AFL fuzzing targets. What we tried to do, however, was to run AFL++ directly on the executable. AFL++ did not do any fuzzing because the executable only used the socket entry point. OSS-Fuzz has a slightly different approach. OSS-Fuzz does not use STDIN mode or AFL fuzzing targets. What OSS-Fuzz does, is use the targets from libFuzzer and create some glue between the clang instrumentation and AFL++, using a C library, to get it to work with the non-AFL instrumentation<sup>4</sup>.

**Run 2** We can fuzz ProFTPD with AFL++ using the OSS-Fuzz helper script (instructions found in Appendix B.3). This time AFL++ worked as intended, but did not give us any errors or hangs (as found in Appendix A.1). It did find enough paths but it did not find any bugs. This is as expected since libFuzzer did not find anything, mainly because of the complexity of the fuzzing target (as explained in Section 5.2.1).

### 5.2.3 Honggfuzz

Just like AFL++ and libFuzzer, we compile Honggfuzz using the instructions for OSS-Fuzz (found in Appendix B.4) and let it run for 24 hours. Honggfuzz also did not give us any crashes or hangs (results found in Appendix A.3) for the same reason the other fuzzers did not find anything, namely a fuzzing target with low code coverage.

<sup>4</sup><https://github.com/google/oss-fuzz/issues/2094>

### 5.3 Comparing AFLNet and OSS-Fuzz over ProFTPD

When we look at ProFTPD, we can see a clear difference between AFLNet and OSS-Fuzz. Although we do not have a clear answer yet to the question “Is OSS-Fuzz wasting resources”, we can already see improvements when using AFLNet.

- AFLNet found a buffer overflow bug while OSS-Fuzz did not find any crashes or hangs.
- While fuzzing ProFTPD using AFLNet (using self-made seeds), we did not find any reproducible crashes or hangs. However, when we used the specific seed reported in a GitHub issue using AFLNet, we got ProFTPD to crash giving us a heap buffer overflow as a result.
- The fuzzing targets in ProFTPD do not provide good code coverage. There is only one fuzzing target which does not cover enough code to even find the heap buffer overflow found by AFLNet. This means that if OSS-Fuzz ran it, they would not find the same crash AFLNet found.

## 6 Fuzzing open62541

Open62541[12] is an open-source implementation of the protocol OPC Unified Architecture (OPC UA for short), which is used for data exchange between sensors and applications. Open62541 is a good candidate because its stateful and fuzzed by OSS-Fuzz. Fijneman[4] tested the execution speed difference between boofuzz<sup>5</sup> and AFLNet on open62541. Although our research’s main topic (comparing fuzzers) is the same, we will only use AFLNet. Fijneman’s results are summarized in Figure 7. As we can see there are 68 crashes and 14 hangs. We can use these results for a baseline of what we are supposed to find with AFLNet on open62541. For open62541 we will use commit 46d0395, just like Fijneman.

Implementation	Lang.	Crashes		Hangs		Exec/s	Runtime
		Total	Reproducible	Total	Reproducible		
legacy OPC UA ANSI-C Stack	C	8	8	0	0	$\pm 20 \cdot 8$	48 hours
open62541	C	68	0	14	0	$\pm 8 \cdot 8$	24 hours
FreeOpcUa	C++	43	43	9	0	$\pm 8 \cdot 8$	24 hours

Figure 7: The result of fuzzing OPC UA implementations in AFLNet by Fijneman[4]

In Section 6.1, we discuss the results of fuzzing open62541 in AFLNet. In Section 6.2 we discuss the the results of fuzzing open62541 with OSS-Fuzz. In the last Section, 6.3, we compare the results of AFLNet and OSS-Fuzz and see if it made any difference.

### 6.1 Fuzzing open62541 with AFLNet

We use the test application created by the developers of open62541, server\_ctt.c, and client.c . We modify the SUT as described in [4], such that server\_ctt and client can be used without authentication and be used by AFLNet. In detail, we:

1. removed **maxSecureChannels**, **maxSession** and **shutdownDelay** from server\_ctt.c . This causes the server\_ctt to be faster inside AFLNet.
2. changed the variable **enableAnon** to **true** in server\_ctt.c so that we can use it without authentication.
3. changed **UA\_Client\_connectUsername** to **UA\_Client\_connect** inside client.c. This causes the client to connect to server\_ctt anonymously.

<sup>5</sup><https://github.com/jtpereyda/boofuzz>

After this, we created a seed file for AFLNet to start fuzzing with. This can be done by using the instructions found in Section 4.2. After which, we can start fuzzing using the following line:

```
afl-fuzz -i seed/ -o 1run -N tcp://127.0.0.1/4840 -P OPCUA
-q 3 -s 3 -E -K -R -m none ./server_ctt
```

**Run 1** Running it with the instructions above, we did not get any problems (results found in Table 12 on page 33). There are 15 unique hangs with a good amount of executions.

**Analyzing run 1** However, after doing a test run, we stumbled upon a fault in our modifications of the SUT. AFLNet only reached 1 state, which can be found in the ipsm.dot file inside the output folder. In this file, AFLNet keeps track of the states it has reached. After talking to Fijneman[4] about the modifications of the SUT, he made a new improved version (since it was already known to Fijneman[4]) which handled the state transition correctly on the new version, but AFLNet could only reach 2 states (found in Figure 8), which is still incorrect since FTP has 4 states. The problem turns out to be the seed namely, this input only reaches 2 states and it would be more difficult for AFLNet to find other state transitions. As we needed to create a seed for AFLNet using tcpdump, we saved it in the wrong format (ASCII instead of raw). Saving in the right format, got us the correct state model (found in Figure 8).




Figure 8: The incorrect state model (left) and the correct state model (right) of open62541 found by AFLNet

**Run 2** AFLNet found 43 unique hangs with 143 total paths (found in Table 13, in page 34). This is still not the same as the crashes found by Fijneman[4]. But may give us some information about faulty coding.

**Analyzing run 2** As already mentioned, AFLNet found 43 hangs. We used aflnet-replay (as we did in Section 5.1). If we test all the hang packets, they all give a segmentation fault on the aflnet-replay side except the last hang. The last 2 give an error on the server side, BadInterError and BadSecureChannelTokenUnknown. They are not bugs, since the bad inputs are caught by the coding and are rejected.

**Run 3** We found another fault in our seed creation. When we create a seed for AFLNet to use (described in Section 4.2), we use tcpdump and Wireshark to create a seed from the TCP stream. We made the mistake of only looking at one TCP stream (as explained in Section 4.2) instead of all 4 TCP streams. Thus, this time we create 4 seeds from all the TCP streams and run AFLNet again. As we can see in Table 14 on page 34, There are no crashes or hangs found, enough paths found, and a good amount of total executions. Run 3 resulted in fewer hangs than run 2. This could be because fuzzing is a non-deterministic way of testing, meaning that it is not guaranteed to find the same bugs as the previous run with the same input.

## 6.2 Fuzzing open62541 with OSS-Fuzz

### 6.2.1 LibFuzzer

Just like ProFTPD, we fuzz open62541 using libFuzzer with the provided fuzzing targets. The open62541 developers provided multiple fuzzing targets, as shown in Table 4. We use these targets to test whether the libFuzzer can find the same crashes as AFLNet.

base 64 decode  
base 64 encode  
binary decode  
binary message  
JSON decode  
JSON decode encode  
MDNS message  
MDNS xht  
src ua util  
TCP message

Table 4: The list of fuzzing targets created by open62541 team

**Run 1** First we build the open62541 image, build the libFuzzer image and fuzzing targets for open62541, and then run it using the OSS-Fuzz helper script (instructions found in Appendix B.2.1) on each of the fuzzing target in the Table 4. The results can be found in Table 5.

Fuzzing open62541 results					
	crash	hang	execs / s	time ran	reproducible
Base64 decode			4096	24h	
Base64 encode			4096	24h	
Binary decode	x		0	$\leq 30s$	yes
Binary message	x		0	$\leq 10s$	no
JSON decode		x	4096	$\leq 5m$	no
JSON decode encode		x	4096	$\leq 5m$	no
MDNS message	x		0	$\leq 10s$	yes
MDNS xht			44k	24h	
src UA util			4096	24h	
TCP message			4	24h	

Table 5: The results of the different libFuzzer fuzzing targets

**Analyzing libFuzzer results** If we look at Table 5, we can see that 3 crashes were found by libFuzzer. Most of them were found right after startup, which explains why the run time was less than a few seconds. Let us first look at the crash (found in Figure 14, on page 39) of MDNS message. A heap buffer overflow caused the error. Rerunning the same tests, results in the same error but with different addresses. This rules out a coincidence. There also were 2 targets that did not respond (hang). The JSON decode and the JSON decode encode target both had a time out of more than 25 seconds, probably because the fuzzing target no longer responded to new inputs.

**Binary decode** First we take a look at the crash of the binary decode target (found in Figure 18 on page 40), which gives us an error because of an unknown address. When we want to debug and try to find out where it crashes, we use GDB (more information about the debugging process is found in Section 4.4). So we run GDB on binary decode with the seed that crashes the program. We set some breakpoints to see where it crashes. This gives us, that it crashes on the last assert (also seen in the crash itself). The assert in question:

```
UA_assert(ret == UA_STATUSCODE_GOOD);
```

Unfortunately, the debugger did not provide us with any information about local variables. This was because the dwarfdata (which is debugger data), is corrupted. Namely, if we start the debugger, we get “Cannot handle DW\_FORM\_unknown”. This fails because the variable ‘ret’ is not the same as 0x00000000.

**Binary message** The binary message crash (found in Figure 17 on page 40) was caused by a heap buffer overflow. It went 4 bytes over bounds. However, when we try to reproduce the crash using the seed directly on the seed without a



fuzzer, we cannot reproduce the crash because it does not crash. Which means this crash is not reproducible.

**JSON decode** The JSON decode hang (found in Figure 15 on page 39), was caused by an unknown bug. When we run it in GDB, and set some breakpoints around functions inside the fuzzing target, we find that it hangs around the execution of the “UA\_decodeJson” function. Whenever we try to dive deeper, we cannot get an answer, because it hangs immediately after going to the function somewhere else in the code. Analyzing this further, we found that after setting the memory limit for the program to unlimited, it succeeded. Thus, the timeout is the problem of not having enough memory given to the fuzzing target. Although the hang is reproducible, it turns out it is not a hang because it did not get enough memory.

**JSON decode encode** The JSON decode encode hang (found in Figure 16 on page 40), was caused by the same bug as the JSON decode fuzzing target. The JSON fuzzing target did not have enough memory, which was fixed by giving the fuzzing target more memory.

**MDNS message** The MDNS message crash (found in Figure 14 on page 39), was also caused by heap buffer overflow. It went immediately out of bounds with less than a byte. When we run GDB with the fuzzing target, we set a breakpoint at specific open62541 functions. And when we run the GDB with the seed that crashes the target, it crashes on the execution of the function “message\_parse”. When we dive deeper into the message parse function we see that it crashes on the execution of:

```
if (!_rrparse(m, m->an, m->ancount, &buf, m->_bufEnd))
    return false;
```

This function parses the RRS part of the MDNS message. When we dive even deeper, we find that it does a memcpy, which writes something in memory but does not check bounds and thus overflows in the allocated memory. There is a comment in the code that states “check if the message has the correct size, i.e. the count matches the number of bytes” (found in 22 on page 44), which the code below should check the size and stop our crash from happening, but does not.

## 6.2.2 AFL++

As discussed in Section 5.2.2, OSS-Fuzz uses the fuzzing targets from libFuzzer to fuzz using AFL++, using a C library created by the OSS-Fuzz team. So we run AFL++ on the existing fuzzing targets with the helper script (described in Appendix B.3).

**Run 1** As shown in Table 6, AFL++ found crashes in the same piece of code but fewer crashes/hangs than libFuzzer. Although AFL++ found a hang in Base64 and MDNS xht, taking a quick peek at the results, it seems that the PC where these fuzzers were run on, had a hang itself since the Base64 encode and the MDNS xht hang happened at the same time. The Binary decode and the MDNS message crashes were no coincidence, they were also found by libFuzzer.

Fuzzing open62541 results				
	crash	hang	execs / s	time ran
Base64 decode			56k	24h
Base64 encode		1	30k	24h
Binary decode	97		10k	24h
Binary message			5	24h
JSON decode			30k	24h
JSON decode encode			30k	24h
MDNS message	15		30k	24h
MDNS xht		1	30k	24h
src UA utitl			30k	24h
TCP message			5	24h

Table 6: The results of the different libFuzzer fuzzing targets using AFL++

**Binary decode** As shown in Table 6, there are 61 crashes. However, when we tried to analyze the crashes, we found out that the OSS-Fuzz helper script did not save the seeds that caused the crashes. So we had to rerun AFL++ on the MDNS fuzzing target, but on the second run, it found 97 crashes which is more crashes. When we try to recreate the crashes we get the same error we found using libFuzzer in Section 6.2.1, binary decode.

**MDNS message** As we can see in Table 6, there are 15 crashes. After running the target seeds with the input seeds that caused these crashes, we can see they all cause a buffer overflow (as we have seen in Section 6.2.1, MDNS message). When we use the debugger to find what causes the crash and try the same method we used in Section 6.2.1 (with the same breakpoints), we can see that it crashes in the same location as found with libFuzzer on the MDNS message fuzzing target.

### 6.2.3 Honggfuzz

Upon trying to fuzz the fuzzing targets with Honggfuzz, we found out that the developers of open62541 do not specify that they want to be fuzzed with Honggfuzz, as every project that is fuzzed by OSS-Fuzz needs to provide specific information, for example, what fuzzer to use. Knowing that they did not specify Honggfuzz to be used, we will not try to get Honggfuzz working on their fuzzing targets (as we also found out that Honggfuzz does not work on open62541).

### 6.3 Comparing AFLNet and OSS-Fuzz over open62541

Observing the bugs found in open62541 (just like in Chapter 5) we noticed a clear difference in the bugs found. This time it is in favor of OSS-Fuzz. The results of open62541 make answering the question “Is OSS-Fuzz wasting resources” more difficult, since this experiment gave opposite results than in Chapter 5.

- OSS-Fuzz found a buffer overflow bug while AFLNet did not find any crashes or hangs.
- Unlike ProFTPD, the fuzzing targets were good enough to find a bug in programming. AFLNet however, did not find anything using the socket. This could probably be because the bad input would be filtered out when reaching the part that crashes on OSS-Fuzz.
- The underlying fuzzers of OSS-Fuzz did not have many differences since AFL++ and libFuzzer found the same bug in the same place.

As we discussed in Section 2.1.3, it is not easy to reproduce fuzzing target crashes on the original program (generally speaking) because it is nearly impossible to reproduce a trace with the part that caused the fuzzing target to crash (as discussed in Section 2.1.3). That is also why it is not easy to compare these results as we cannot test the bad inputs of OSS-Fuzz on the program as a whole but we can test the bad inputs of AFLNet. This was no problem when using ProFTPD as SUT because the fuzzing target in ProFTPD did not find any crashes or hangs. This would have been a problem if the fuzzing target did find crashes or hangs.

## 7 Future work

### 7.1 More stateful fuzzers

In the future, it might be an idea to test different stateful fuzzers, like SGFuzzer[1], to see if there is a better fuzzer than AFLNet, that would work better for OSS-Fuzz in both bug finding and getting it to work with OSS-Fuzz. This is because AFLNet requires the developer to create a packet parser in code into AFLNet and there might be an even better stateful fuzzer at finding bugs.

### 7.2 More case studies

It is also an idea to test using more case studies because 2 case studies might not be enough to exactly tell the difference between AFLNet and OSS-Fuzz. It would also be good to test protocols with encryption. We did not test protocols with encryption because the nature of encryption makes it more difficult to fuzz those protocols than protocols without.

### 7.3 Better fuzzing targets

In this project by Google, called OSS-Fuzz-gen[10], Google uses LLMs (Large Language Models), like OpenAI ChatGPT or Google Gemini, to create fuzzing targets to increase the code coverage of the fuzzing for stateless and stateful projects. If you have better code coverage, you have more chance of finding faults in the code. OSS-Fuzz depends heavily on good fuzzing targets for all their fuzzers (libFuzzer, AFL++, Honggfuzz), this is why it would be an idea to try to use the LLM OSS-Fuzz-gen, to improve code coverage and thus increase OSS-Fuzz bug finding vs AFLNet.

## 8 Conclusion

**Speed** OSS-Fuzz does substantially more executions than AFLNet (30,000 vs 5) because AFLNet fuzzes the socket entry point which in turn uses a large part of the code while OSS-Fuzz only fuzzes smaller parts of the code at once. This might make OSS-Fuzz better at finding bugs, but whenever a bug is not reachable by OSS-Fuzz (because of state transitions) they can be faster but they will never find the bug.

**ProFTPD** AFLNet improved the bug finding on ProFTPD compared to OSS-Fuzz. Although OSS-Fuzz is not wasting a lot of resources since ProFTPD has only one fuzzing target, OSS-Fuzz might not even run the fuzzing target for long as it might notice that some statistics (like code coverage and crashes) are not changing, and thus terminating early. Points to note on ProFTPD (for more analysis of ProFTPD, see Section 5.3),

- ProFTPD has a lot to improve on when talking about the fuzzing targets. Since the fuzzing targets only reached 2.8% ~ coverage (see analysis in Section 5.3), ProFTPD can add more targets like open62541 and improve the ones they have.
- AFLNet could improve OSS-Fuzz when fuzzing ProFTPD. Even though we saw that the bug finding is dependent on the creation of good fuzzing targets, when using AFLNet it would decrease the burden of the creation of good fuzzing targets to have results (as we see in Section 5.3). Implementing AFLNet does, however, require developers who request fuzzing by AFLNet, to provide an entry point (unified socket) and an implementation of the protocol into AFLNet.

**Open62541** The results for open62541, however, show us something different. While we expected the same results for open62541 as ProFTPD, the results for open62541 are the opposite. Points to note on open62541 (for more analysis of open62541, see Section 6.3),

- Open62541 does not have a lot to improve on when talking about the fuzzing target, because while AFLNet did not find any reproducible bugs, OSS-Fuzz did find bugs.
- AFLNet would not improve OSS-Fuzz when fuzzing open62541, as open62541 had good enough fuzzing targets to find bugs.
- OSS-Fuzz did find crashes, more specifically AFL++ and libFuzzer did. Both fuzzers found the same bugs, which caused a buffer overflow. This bug was created by a memcpy function that did not check any bounds.
- OSS-Fuzz also found more unreachable bugs, than reachable bugs. Namely, when executing the normal path of the program, the bugs found by OSS-Fuzz might be not possible since the input might have already been filtered and checked.

**AFLNet vs OSS-Fuzz** We do not have a clear answer to the question “Is OSS-Fuzz wasting resources?”. ProFTPD shows us that OSS-Fuzz is wasting resources since it is not fuzzing with many results. Open62541 however, shows us a different picture. AFL++ and libFuzzer found bugs in open62541 while AFLNet did not. AFLNet would be better at finding bugs against bad fuzzing targets. There would be no improvement when there are good fuzzing targets. Answering this question would require more time, more specifically a bigger sample rate.

Another way of looking at it would be that OSS-Fuzz is not wasting time but the developers are, since OSS-Fuzz depends heavily on good fuzzing targets. Creating good fuzzing targets would be wasting time instead of using AFLNet which only requires a one-time creation of a code parser inside AFLNet and takes less time.

**OSS-Fuzz** One of the assumptions we made is that the developers know the effect of the quality of their fuzzing targets on the quality of the results of OSS-Fuzz. This assumption might be wrong since ProFTPD has bad fuzzing targets and we do not know the exact reason why. The reason for bad fuzzing targets could also have been that the developers submitted their project for the money since OSS-Fuzz gives out money for submitting your project.

**Recommendations** Although there is no clear answer to the question “Is OSS-Fuzz wasting resources?”, we do have a recommendation for developers who submit their projects. As OSS-Fuzz is heavily dependent on good fuzzing targets, making good fuzzing targets would increase the bug finding with OSS-Fuzz. This would also be a recommendation to OSS-Fuzz as they might want to recommend creating fuzzing targets with high code coverage to ensure good fuzzing targets.

## Glossary

**ASan** Address Sanitizer

**UBSan** Undefined Behaviour Sanitizer

**OPCUA** OPC Unified Architecture

**FTP** File Transfer Protocol

**SUT** System Under Test

**TCP** Transmission Control Protocol

**UDP** User Datagram Protocol

## References

- [1] Jinsheng Ba, Marcel Böhme, Zahra Mirzamomen, and Abhik Roychoudhury. Stateful greybox fuzzing. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 3255–3272, Boston, MA, August 2022. USENIX Association.
- [2] Castaglia. Proftpd. <https://github.com/proftpd/proftpd>, 2014.
- [3] Zhen Yu Ding and Claire Le Goues. An empirical study of OSS-Fuzz bugs. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*, pages 131–142, 2021.
- [4] Mark Fijneman. *Fuzzing open source OPC UA implementations*. bachelor thesis, Radboud universiteit, Nijmegen, Gelderland, August 2023.
- [5] Andrea Fioraldi, Dominik Maier, Heiko Eißfeldt, and Marc Heuse. AFL++ : Combining incremental steps of fuzzing research. In *14th USENIX Workshop on Offensive Technologies (WOOT 20)*. USENIX Association, August 2020.
- [6] Google. UndefinedBehaviorSanitizer. <https://clang.llvm.org/docs/UndefinedBehaviorSanitizer.html>.
- [7] Google. Honggfuzz. <https://github.com/google/honggfuzz>, 2015.
- [8] Google. OSS-Fuzz. <https://github.com/google/oss-fuzz>, 2016.
- [9] Google. AddressSanitizer. <https://github.com/google/sanitizers/wiki/AddressSanitizer>, 2019.
- [10] Dongge Liu, Jonathan Metzman, and Oliver Chang. AI-powered fuzzing: Breaking the bug hunting barrier. <https://security.googleblog.com/2023/08/ai-powered-fuzzing-breaking-bug-hunting.html>, Aug 2023.
- [11] Van-Thuan Pham, Marcel Böhme, and Abhik Roychoudhury. Afnet: A greybox fuzzer for network protocols. In *2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST)*, pages 460–465, 2020.
- [12] Stefan Profanter, Ulius Pfrommer, and Noel Graf. open62541. <https://github.com/open62541/open62541>, 2013.
- [13] LLVM project team. libFuzzer. <https://llvm.org/docs/LibFuzzer.html>.
- [14] Michael Sutton, Adam Greene, and Pedram Amini. *Fuzzing: Brute Force Vulnerability Discovery*. Addison-Wesley Professional, Boston, MA, 2007.



## A Results

### A.1 ProFTPD in AFLNet

The results of fuzzing ProFTPD commit 5e57d41 with AFLNet without any sanitation over multiple runs.

Run time	1 day, 20 hours, 2 minutes, 1 seconds
Last unique crash found	None
Last unique hang found	None
Total unique hangs	None
Total unique crashes	None
Total paths	$800 \times 6$
Exec speed	$\sim 3 / \text{sec}$

Table 7: The summary of the first run of AFLNet on ProFTPD, summing the statistics of 6 fuzzers

Run time	1 day, 6 hours, 42 minutes, 11 seconds
Last unique crash found	None
Last unique hang found	2 hours, 21 minutes, 3 seconds
Total unique hangs	29
Total unique crashes	None
Total paths	2808
Exec speed	1 / sec

Table 8: The summary of the third run of ProFTPD using AFLNet, using ASan and UBSan, with the correct configuration

### A.2 ProFTPD in AFL++

The results of fuzzing ProFTPD commit 5e57d41 with AFL++ with sanitation over multiple runs.

Run time	0 days, 20 hours, 42 minutes, 7 seconds
Last new path found	None
Last unique crash found	None
Last unique hang found	None
Total unique hangs	0
Total unique crashes	None
Total execs	33.7 million
Exec speed	507 / sec

Table 9: The summary of the first run of AFL++ on ProFTPD

Run time	1 days, 0 hours, 22 minutes, 36 seconds
Last new path found	0 days, 3 hours, 54 minutes, 31 seconds
Last unique crash found	None
Last unique hang found	None
Total unique hangs	0
Total unique crashes	None
Total execs	4.32 billion
Exec speed	48.4k / sec

Table 10: The summary of the second run of AFL++ on ProFTPD

### A.3 ProFTPD in Honggfuzz

The results of fuzzing ProFTPD commit 5e57d41 with AFL++ with sanitation over a single run.

Summary iterations	634679822
time	85000
speed	7466
crashes_count	0
timeout_count	0
new_units_added	661
slowest_unit_nm	714
guard_nb	19394
branch_coverage_percent	1
peak_rss_mb	0

Table 11: The results of ProFTPD in Honggfuzz

### A.4 Open62541 in AFLNet

The results of fuzzing open62541 commit 5e57d41 with AFLNet with sanitation over multiple runs.

Run time	1 day, 0 hours, 21 minutes, 49 seconds
Last unique crash found	None
Last unique hang found	0 days, 17 hours, 53 minutes, 36 seconds
Total unique hangs	15
Total unique crashes	None
Total paths	32
Total execs	79.7k
Exec speed	~ 5 / sec

Table 12: The summary of the first run of open62541 in AFLnet

Run time	1 day, 0 hours, 36 minutes, 55 seconds
Last unique crash found	None
Last unique hang found	0 days, 17 hours, 53 minutes, 36 seconds
Total unique hangs	43
Total unique crashes	None
Total paths	32
Total execs	89.5k
Exec speed	~ 5 / sec

Table 13: The summary of the second run of open62541 in AFLNet, now with fixed AFLNet implementation

Run time	1 day, 0 hours, 33 minutes, 45 seconds
Last unique crash found	None
Last unique hang found	None
Total unique hangs	0
Total unique crashes	None
Total paths	219
Total execs	381k
Exec speed	~ 5 / sec

Table 14: The summary of the third run of open62541 in AFLNet, now with 4 seeds

## A.5 Open62541 in AFL++

The results of fuzzing open62541 commit 5e57d41 with AFL++ with sanitation over a single run.

	base64 encode	base64 decode	binary decode	binary message	JSON decode
Run time	24 hours	24 hours	24 hours	24 hours	24 hours
Last unique crash	None	None	23 minutes	None	None
Last unique hang	1 hour	None	None	None	None
Total unique hangs	1	0	0	0	0
Total unique crashes	0	0	61	0	0
Total execs	1.61G	5.57G	698M	797K	2.83G
Exec speed	17.9k/sec	56.8k/sec	13.0k/sec	14.0k/sec	37.8k/sec

Table 15: The first part of the summary of the afl AFL++ results on open62541

	JSON encode	MDNS message	MDNS xht	src ua util	TCP message
Run time	24 hours	24 hours	24 hours	24 hours	24 hours
Last unique crash	None	11 hours	2 hours	None	None
Last unique hang	None	None	None	None	None
Total unique hangs	0	0	1	0	0
Total unique crashes	0	15	0	0	0
Total execs	2.72G	1.6G	234M	3.44G	349k
Exec speed	28.3k/sec	33.8k/sec	277/sec	45.2k/sec	8k/sec

Table 16: The second part of the summary of the afl AFL++ results on open62541

## B Installation

The way the different fuzzers and SUTs are compiled and run.

### B.1 ProFTPD

```
$ CFLAGS='-g -fsanitize=address, undefined'
  LDFLAGS='-fsanitize=address, undefined'
  CXXFLAGS='-g -fsanitize=address, undefined' CC=
  "afl-clang-fast" CXX="afl-clang-fast++" ./
  configure --enable-devel=nodaemon:nofork
$ make
$ make install
$ afl-fuzz -d -i in-ftp/ -o fourthrun/ -x /home/
  ubuntu/aflnet/tutorials/lightftp/ftp.dict -N
  tcp://127.0.0.1/2021 -R -E -q 3 -s 3 -K -D
  10000 -m none -P FTP proftpd -n -c /tmp/PFTEST
  /PFTEST.conf
```

Figure 9: Instructions to compile and fuzz ProFTPD

### B.2 OSS-Fuzz

The instructions to run libFuzzer and AFL++ with the help of the OSS-Fuzz helper script.

#### B.2.1 Instructions for the OSS-Fuzz helper script

```
$ cd path/to/oss-fuzz
$ python3 infra/helper.py build_image
  $PROJECT_NAME
$ python3 infra/helper.py build_fuzzers --engine
  libfuzzer $PROJECT_NAME
$ python3 infra/helper.py run_fuzzer --corpus-dir
  =<path-to-temp-corpus-dir> --engine afl
  $PROJECT_NAME <fuzz_target>
```

Figure 10: Instructions to fuzz projects with libFuzzer using OSS-Fuzz helper scripts

### B.3 AFL++ in OSS-Fuzz helper script

```
$ cd path/to/oss-fuzz
$ python3 infra/helper.py build_image
  $PROJECT_NAME
$ python3 infra/helper.py build_fuzzers --engine
  afl $PROJECT_NAME
$ python3 infra/helper.py run_fuzzer --corpus-dir
  =<path-to-temp-corpus-dir> --engine afl
  $PROJECT_NAME <fuzz_target>
```

Figure 11: Instructions to fuzz projects with AFL++ using OSS-Fuzz helper scripts

### B.4 Honggfuzz in OSS-Fuzz helper script

```
$ cd path/to/oss-fuzz
$ python3 infra/helper.py build_image
  $PROJECT_NAME
$ python3 infra/helper.py build_fuzzers --engine
  honggfuzz $PROJECT_NAME
$ python3 infra/helper.py run_fuzzer --corpus-dir
  =<path-to-temp-corpus-dir> --engine honggfuzz
  $PROJECT_NAME <fuzz_target>
```

Figure 12: Instructions to fuzz projects with Honggfuzz using OSS-Fuzz helper scripts

## C Bugs

The bugs we found during the fuzzing of software.

### C.1 ProFTPD

The crash when trying to reproduce the crash from AFLNet on ProFTPD using the seeds from the original GitHub issue.

```
==22008==ERROR: AddressSanitizer: heap-buffer-overflow on address 0x621000015320 at pc 0x00000052e099 bp 0x7fffffff7b10 sp 0x7fffffff7b10
READ of size 1 at 0x621000015320 thread T0
#0 0x52e098 (/usr/local/sbin/proftpd+0x52e098)
#1 0x52ac66 (/usr/local/sbin/proftpd+0x52ac66)
#2 0x5298be (/usr/local/sbin/proftpd+0x5298be)
#3 0x55ac8c (/usr/local/sbin/proftpd+0x55ac8c)
#4 0x550b76 (/usr/local/sbin/proftpd+0x550b76)
#5 0x54b100 (/usr/local/sbin/proftpd+0x54b100)
#6 0x547bba (/usr/local/sbin/proftpd+0x547bba)
#7 0x7ffff9e4c86 (/lib/x86_64-linux-gnu/libc.so.6+0x21c86)
#8 0x42b109 (/usr/local/sbin/proftpd+0x42b109)

0x621000015320 is located 0 bytes to the right of 4640-byte region [0x621000014100,0x621000015320)
allocated by thread T0 here:
#0 0x4eafc0 (/usr/local/sbin/proftpd+0x4eafc0)
#1 0x58f453 (/usr/local/sbin/proftpd+0x58f453)
#2 0x58e869 (/usr/local/sbin/proftpd+0x58e869)
#3 0x57097d (/usr/local/sbin/proftpd+0x57097d)
#4 0x582736 (/usr/local/sbin/proftpd+0x582736)
#5 0x5816bf (/usr/local/sbin/proftpd+0x5816bf)
#6 0x5837e6 (/usr/local/sbin/proftpd+0x5837e6)
#7 0x527ff8 (/usr/local/sbin/proftpd+0x527ff8)
#8 0x5298be (/usr/local/sbin/proftpd+0x5298be)
#9 0x55ac8c (/usr/local/sbin/proftpd+0x55ac8c)
#10 0x550b76 (/usr/local/sbin/proftpd+0x550b76)
#11 0x54b100 (/usr/local/sbin/proftpd+0x54b100)
#12 0x547bba (/usr/local/sbin/proftpd+0x547bba)
#13 0x7ffff9e4c86 (/lib/x86_64-linux-gnu/libc.so.6+0x21c86)

SUMMARY: AddressSanitizer: heap-buffer-overflow (/usr/local/sbin/proftpd+0x52e098)
Shadow bytes around the buggy address:
 0x0c427fffa10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x0c427fffa20: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x0c427fffa30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x0c427fffa40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x0c427fffa50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
=>0x0c427fffa60: 00 00 00 [fa]fa fa fa fa fa fa fa fa fa fa fa
 0x0c427fffa70: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c427fffa80: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c427fffa90: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c427fffaa0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c427fffab0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
Shadow byte legend (one shadow byte represents 8 application bytes):
Addressable: 00
Partially addressable: 01 02 03 04 05 06 07
Heap left redzone: fa
Freed heap region: fd
Stack left redzone: f1
Stack mid redzone: f2
Stack right redzone: f3
Stack after return: f5
Stack use after scope: f8
Global redzone: f9
Global init order: f6
Poisoned by user: f7
Container overflow: fc
Array cookie: ac
Intra object redzone: bb
ASan internal: fe
Left alloca redzone: ca
Right alloca redzone: cb
==22008==ABORTING
```

Figure 13: Heap buffer overflow found by AFLNet

### C.2 Open62541

These are the bugs found by fuzzing different targets of open62541 using libFuzzer. The fuzzing targets include MDNS message, JSON decode, JSON decode encode, binary message, binary decode.







## D Source code

In this chapter, we show different source codes of the SUTs and the fuzzing targets.

### D.1 fuzz\_mdns\_message.cc

```
0 #include <cstdint>
1 #include <cstdio>
2 #include <libmdnsd/mdnsd.h>
3 #include <cstring>
4 #include <cstdlib>
5
6 /*
7 ** Main entry point. The fuzzer invokes this function with each
8 ** fuzzed input.
9 **/
10 extern "C" int
11 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
12
13     struct message m;
14     memset(&m, 0, sizeof(struct message));
15
16     unsigned char *dataCopy = (unsigned char *)malloc(size);
17     if (!dataCopy) {
18         return 0;
19     }
20
21     memcpy(dataCopy, data, size);
22
23     bool success = message_parse(&m, dataCopy, size);
24
25     free(dataCopy);
26
27     if (!success)
28         return 0;
29
30     mdns_daemon_t *d = mdnsd_new(QCLASS.IN, 1000);
31
32     struct sockaddr_storage addr;
33     mdnsd_in(d, &m, (struct sockaddr *)&addr, 2000);
34
35     mdnsd_free(d);
36
37     return 0;
38 }
39 }
```

Figure 19: The source code of the MDNS message fuzzing target

## D.2 fuzz\_binary\_decode.cc

```
0 #include "custom_memory_manager.h"
1 #include <open62541/plugin/log_stdout.h>
2 #include <open62541/server_config_default.h>
3 #include <open62541/types.h>
4
5 /*
6 ** Main entry point. The fuzzer invokes this function with each
7 ** fuzzed input.
8 **/
9 extern "C" int LLVMFuzzerTestOneInput(const uint8_t *data, size_t
    size) {
10     if(size <= 6)
11         return 0;
12
13     // set the available memory
14     if(!UA_memoryManager_setLimitFromLast4Bytes(data, size))
15         return 0;
16
17     data += 4;
18     size -= 4;
19
20     // get some random type
21     uint16_t typeIndex = (uint16_t)(data[0] | data[1] << 8);
22     data += 2;
23     size -= 2;
24
25     if(typeIndex >= UA_TYPES_COUNT)
26         return UA_FALSE;
27
28     void *dst = UA_new(&UA_TYPES[typeIndex]);
29     if(!dst)
30         return 0;
31
32     const UA_ByteString binary = {
33         size, //length
34         (UA_Byte *) (void *) data
35     };
36
37     UA_StatusCode ret = UA_decodeBinary(&binary, dst, &UA_TYPES[
        typeIndex], NULL);
38     if(ret != UA_STATUSCODE_GOOD) {
39         UA_delete(dst, &UA_TYPES[typeIndex]);
40         return 0;
41     }
```

Figure 20: The first part of the source code of the binary decode fuzzing target

```

42 // copy the datatype to test
43 void *dstCopy = UA_new(&UA_TYPES[typeIndex]);
44 if(!dstCopy) {
45     UA_delete(dst, &UA_TYPES[typeIndex]);
46     return 0;
47 }
48 ret = UA_copy(dst, dstCopy, &UA_TYPES[typeIndex]);
49 if(ret != UA_STATUSCODE_GOOD) {
50     UA_delete(dst, &UA_TYPES[typeIndex]);
51     UA_delete(dstCopy, &UA_TYPES[typeIndex]);
52     return 0;
53 }
54
55 // compare with copy
56 UA_assert(UA_order(dst, dstCopy, &UA_TYPES[typeIndex]) ==
57           UA_ORDER_EQ);
58 UA_delete(dstCopy, &UA_TYPES[typeIndex]);
59
60 // now also test encoding
61 size_t encSize = UA_calcSizeBinary(dst, &UA_TYPES[typeIndex]);
62 UA_ByteString encoded;
63 ret = UA_ByteString_allocBuffer(&encoded, encSize);
64 if(ret != UA_STATUSCODE_GOOD) {
65     UA_delete(dst, &UA_TYPES[typeIndex]);
66     return 0;
67 }
68
69 ret = UA_encodeBinary(dst, &UA_TYPES[typeIndex], &encoded);
70 UA_assert(ret == UA_STATUSCODE_GOOD);
71
72 UA_ByteString_clear(&encoded);
73 UA_delete(dst, &UA_TYPES[typeIndex]);
74 return 0;
75 }

```

Figure 21: The second part of the source code of the binary decode fuzzing target

```

504 // check if the message has the correct size, i.e. the count
505 // matches the number of bytes
506
507 /* Process questions */
508 my(m->qd, (sizeof(struct question) * m->qdcount), struct
509 question *);
510 for (i = 0; i < m->qdcount; i++) {
511     if (!_label(m, &buf, m->_bufEnd, &(m->qd[i].name)))
512     {
513         return false;
514     }
515     if (buf + 4 > m->_bufEnd) {
516         return false;
517     }
518     m->qd[i].type = net2short(&buf);
519     m->qd[i].clazz = net2short(&buf);
520 }
521 if (buf > m->_bufEnd) {
522     return false;
523 }
524
525 /* Process rrs */
526 my(m->an, (sizeof(struct resource) * m->ancount), struct
527 resource *);
528 my(m->ns, (sizeof(struct resource) * m->nscount), struct
529 resource *);
530 my(m->ar, (sizeof(struct resource) * m->arcount), struct
531 resource *);
532 if (!_rrparse(m, m->an, m->ancount, &buf, m->_bufEnd))
533     return false;
534 if (!_rrparse(m, m->ns, m->nscount, &buf, m->_bufEnd))
535     return false;
536 if (!_rrparse(m, m->ar, m->arcount, &buf, m->_bufEnd))
537     return false;
538 return true;
539 }

```

Figure 22: The part of the source code where the rrpars function is called with the wrong size