

# ProFuzzBench: A Benchmark for Stateful Protocol Fuzzing

Roberto Natella

roberto.natella@unina.it

Università degli Studi di Napoli Federico II, Italy

Van-Thuan Pham

thuan.pham@unimelb.edu.au

University of Melbourne, Australia

## ABSTRACT

We present a new benchmark (*ProFuzzBench*) for stateful fuzzing of network protocols. The benchmark includes a suite of representative open-source network servers for popular protocols, and tools to automate experimentation. We discuss challenges and potential directions for future research based on this benchmark.

## KEYWORDS

Fuzzing; Benchmarking; Network Protocols

## 1 INTRODUCTION

Protocols are pervasive in computer systems, as they enable communication among parties over a local network or the internet, such as, for instant messaging, email, file serving and sharing, multimedia streaming, and several other applications. Therefore, protocol *implementations* are an appealing target for malicious actors, as a vulnerability in an implementation may be remotely exploited. However, it is challenging to test the security of protocol implementations, since protocols are most often *stateful* in nature. This means that, compared to stateless programs, the input space for testing is not limited to the format of individual messages (which can be very large by its own), but is further enlarged by the potential combinations of several messages. Therefore, protocol security testing is in need of techniques that can efficiently test such input space, by taking into account the protocol states.

Much of recent research on security testing has been on *fuzzing* [13, 15], which automatically exercises a target system with large volumes of inputs, either *generated* using a model of the input space (e.g., a protocol state machine), or *mutated* from an initial set of regular inputs (seeds). Research on fuzzing is of experimental nature, as this technique lies on heuristics, such as on seed selection and mutation operators. Heuristics leverage empirical insights about typical vulnerabilities, such as, knowing that out-of-bounds integers are often not well managed, mutation operators perform integer operations on the seeds.

Therefore, one of the key factors for progress in fuzzing research is the availability of **benchmarks**, that is, representative and well-defined experimental targets for applying heuristic techniques, and for obtaining reproducible and quantitative evaluations. However, there is still a debate about how to benchmark fuzzers, e.g., how to measure progress and depth of fuzzing, how to avoid biases, and whether the targets are representative of real vulnerable systems. Existing proposals for benchmarking fuzzers include *LAVA-M* [7], a set of open-source program injected with synthetically-generated bugs; *Google FuzzBench* [14], a cloud service for running a fuzzer on a set of target programs from the OSS-Fuzz project; the DARPA *Cyber Grand Challenge* corpus [5]; and *Magma* [11], a benchmark of targets with manually-curated vulnerabilities. However, existing benchmarks focus on *stateless* libraries and programs (e.g., parsers

for multimedia file formats, data compression tools, OS utilities, etc.), but have limited or no support for *stateful* software, which would be useful to make progress on fuzzing techniques suitable for protocol testing.

In this paper, we present a new benchmark (*ProFuzzBench*) for stateful fuzzing of network protocols. This benchmark provides a suite of open-source programs that implement popular network protocols, and tools for fully automating the execution of fuzzing experiments. We selected programs representative of the ones adopted in previous research on stateful fuzzing, and of network protocols targeted by commercial and open-source fuzzers. Based on our experience with these programs, we discuss some open research issues for stateful network protocol fuzzing, and potential future directions. We release the benchmark as open-source, inviting the community to use it for evaluation of new fuzzing techniques, and for further extend it with more targets.

In the following of this paper, Section 2 presents the benchmark targets; Section 3 elaborates on fuzzing automation; Section 4 discusses challenges for applying fuzzers on stateful network protocol implementations; Section 5 concludes the paper.

## 2 BENCHMARK TARGETS

We selected a set of network protocols that we wanted to cover with the benchmark. We looked for protocols that are mature and widely used, both by enterprises and by individual users. Moreover, the selection of the protocols took into account which protocols were analyzed by previous fuzzing studies, as reported in academic papers, books, technical blogs, or supported by non-academic fuzzing tools (e.g., commercial products).

For every selected protocol, we looked for open-source software that implements that protocol, to be subject to fuzz testing. In this case too, we looked for mature software suitable for use in real-world applications (thus, worth of being tested for security vulnerabilities), and actively maintained and popular among users.

Table 1 lists the protocols, and the related software implementations, that were selected for the benchmark. In total, we selected 10 protocols, and one software implementation for each of them. FTP is the only exception, for which we included two implementations in the benchmark. Since FTP is often adopted as running example in books and papers to present fuzzing techniques and tools, we found several FTP implementations that were subjected to fuzz testing. Thus, we included more than one software for this protocol.

For each protocol, the table provides a brief description of the objective of the protocol, and of the stateful aspects of the protocol. In many cases, the state of the protocol keeps track of the progress of a communication session between a client and a server. For example, the FTP protocol prescribes initial states for authentication, with messages for providing the username and password; then, authentication is followed by one or more commands for file transfer; and finally, the session ends with a log-out. Similarly, in the case of

**Table 1: Benchmark targets**

Protocol	Implementation	Description	Protocol states	Related papers	Non-academic fuzzing tools
RTSP	Live555	Real-time media streaming	Progress of ongoing streaming	[9, 16]	[4, 19]
SMTP	Exim	Email transmission	Message queue, progress of the session	[20]	[4, 17, 19]
FTP	ProFTPD, LightFTP	File transfer	CWD, session flags, progress of the session	[10, 16, 20]	[4, 17, 19]
SSH	OpenSSH	Secure remote shell	User authentication progress, session configuration	[24]	[4, 17, 19]
TLS	OpenSSL	Secure socket connection	User authentication progress, session configuration	[6, 21]	[4, 19]
DTLS	TinyDTLS	Secure datagram communication	User authentication progress, session configuration	[8]	[19]
DNS	dnsmasq	Network domain names	Cached DNS records	[20, 24]	[4, 17, 19]
SIP	Kamailio	Signaling protocol for real-time sessions	User registrations, progress of the session	[1-3]	[4, 19]
DAAP	forked-daapd	HTTP-based audio library streaming	Progress of ongoing streaming, playlist	[18]	[4, 17, 19]
DICOM	demqrsep	Image retrieval	Progress of the session	[22]	[19]

the secure communication protocols (SSH, TLS, DTLS), the protocol state tracks the order of the steps in the authentication process, including the exchange of public key and other information for encryption. In some protocols, the protocol state reflects the current state of the server process, and the history of its interactions with the client: for example, in SMTP and DAAP, the server enqueues emails to be transmitted, and songs to be played; while they are processed, the server continues to interact with the client through further messages. The benchmark also includes the DNS protocol, despite it is in principle a stateless one. We still opted to include this protocol in the benchmark, since it is a popular fuzzing target, and since its implementations can still provide stateful behavior, such as the caching of records from past queries.

The table provides a non-exhaustive list of papers and books related to each protocol. In these references, the protocols were adopted as case studies for evaluating the effectiveness of new techniques for stateful protocol fuzzing. In the case of some references, the protocols were adopted to provide context for examples, such as to show how to customize and to configure a fuzzing tool for a new protocol. For example, in generational fuzzing, the user needs to provide a model for the protocol under test, either by writing it from scratch, or by reusing an existing protocol model (e.g., based on protocol specifications).

Finally, the table lists non-academic fuzzing tools (commercial and open-source) that provide support for each protocol. The commercial Defensics tool from Synopsis (originally developed by Codenomicon) and beSTORM from Beyond Security, support an extensive set of network protocols, including the ones that are part of the benchmark. The Metasploit Framework (best known for its exploit database) also provides a small set of fuzzers for few basic network protocols, which were included in the benchmark. The DAAP protocol is not directly supported by these fuzzers, but can be fuzzed using these fuzzing tools using their support for HTTP.

### 3 AUTOMATION

The protocol fuzzing benchmark has been published as a set of utilities for automating the compilation, configuration, execution, and analysis of target software (Figure 1). These utilities are based on Docker to automate the deployment of the target software and of fuzzing tools, in order to achieve reproducible experiments, and to support the comparative analysis of different fuzzing techniques under controlled conditions.

The first stage is the build of a container image. The target software, and a set of protocol fuzzing tools, are first cloned from their open-source repositories. The tools currently included in the

benchmark are AFLNWE (a basic variant of AFL, to support fuzzing over network sockets) and AFLNET (a protocol-aware fuzzer, also based on AFL) [16]. The benchmark comes with a set of patches for each target software, which are applied during this state. For most target software, the patches de-randomize the program, such as, by embedding a fixed seed in the source code. De-randomization is important to have reproducible behavior, i.e., if the program is executed again with the same input, then the same execution path is covered. This reproducible behavior is an implicit assumption for the correct application of coverage-driven fuzzing techniques. Moreover, de-randomization supports the comparison of different fuzzing techniques, as it makes it less likely that a fuzzer perform better than other by chance, rather than for its merits. In some cases, the patches decrease or remove delays in the program (e.g., uses of `sleep()`), which are adopted for synchronization in network settings, but which unnecessarily slow down fuzzing since it can be performed locally with smaller delays. Other patches make the program compatible with fuzzers. For example, AFL uses UNIX signals to manage the target process and to detect its failures. Thus, any custom handler for these signals should be disabled to avoid interfering with AFL-based fuzzing tools.

Some patches of the benchmark support the experimentation of tools tailored for protocol fuzzing. For example, AFLNET [16] takes in input a corpus of network traffic traces, where an individual fuzz input is represented by one traces, which consists of a sequence of messages exchanged with the target software. In order to test more states of the protocol, a tool may need to focus fuzzing on a specific message in the sequence. Thus, AFLNET performs a parsing of the traffic traces, by identifying the position of the individual messages. AFLNET includes several parsers for the supported protocols. Where available, the parsers look for marker characters to split the messages, such as, the characters `0x0D0A` in FTP. Some protocol implementations do not provide such markers (e.g., a message ends when the connection is closed, without an explicit marker). In these cases, the patches introduce such characters to simplify the implementation of message parsers.

Moreover, the patches support protocol fuzzing tools that keep track of protocol states in order to maximize state coverage. Some protocols, such as FTP, provide distinctive status codes in response messages from the server to the client (e.g., the number 230 when the login succeeded and the server awaits for commands or number 257 when a new folder has been successfully created). However, it might not be the case for all protocols. For instance, RTSP protocol uses the same response code 200 for all successful requests, regardless of the current server state. To make stateful fuzzing

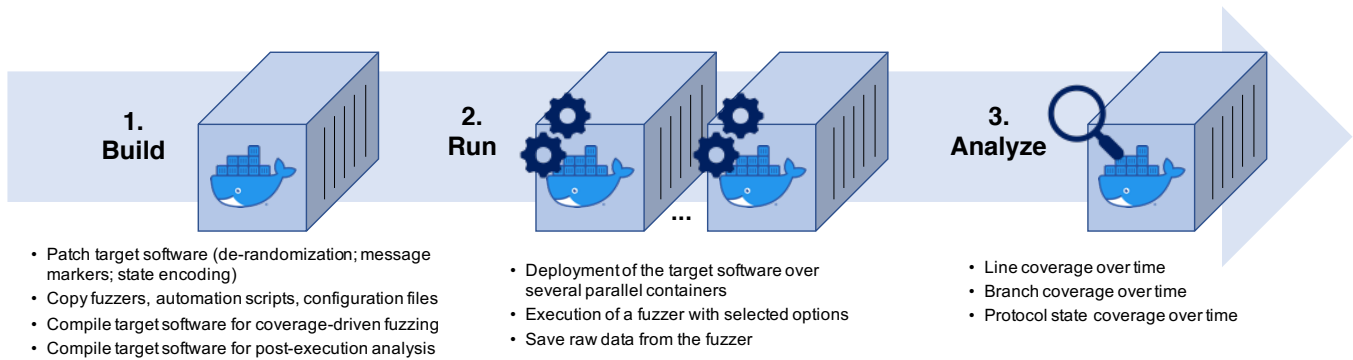


Figure 1: Workflow of benchmark automation.

more effective, our benchmark patches the target (Live555 server) to send responses with different status codes e.g., 201 when a PLAY command has been processed and 203 for a PAUSE command.

The (patched) target software is then compiled within the container image. It is compiled a first time using the AFL compiler wrapper (i.e., `afl-clang-fast/afl-clang-fast++`), in order to include the runtime instrumentation needed by AFL-based coverage-driven fuzzing tools. This binary will be the target for fuzzing. The target software is compiled a second time with support for `gcov`: this binary will be used in the last stage of the benchmark, in order to compute coverage metrics after fuzzing. The build stage of the benchmark finalizes the container image by adding target-specific files (e.g., configuration file, media sources). Finally, the fuzzing tools and automation scripts are copied in the container image.

In the second stage, the benchmark runs the selected fuzzer on the target software, using the container image. In order to compare different fuzzers in a rigorous way, the same fuzzing experiment should be repeated several times. In this way, two fuzzers can be compared with respect to their results over several repetitions, using statistical techniques analysis to achieve confidence that one of the fuzzers is indeed superior to the other [12]. Therefore, the utilities provided with the benchmark allow the user to run multiple instances of the same container in parallel, by taking advantage of multicore architectures. Finally, raw data from the fuzzer is copied from the container to the host machine. The raw data include inputs generated by the fuzzer that covered a new path in the program (e.g., “interesting inputs” in AFL); inputs that discovered new state transitions in the protocol state machine (e.g., as in AFLNET); and inputs that caused a crash or hang of the target software.

The final stage analyzes the raw data to generate statistics about the fuzzing experiments. It runs again the inputs using the target software compiled with `gcov`, in order to obtain line- and branch-level information about the coverage of each input. Then, it generates time series in CSV format for the line and branch coverage, based on the timestamps at which the inputs were generated, to be used for plotting and for comparing the fuzzers.

## 4 DISCUSSION

The availability of this benchmark is an opportunity for research in new techniques for protocol fuzzing. In selecting the target software,

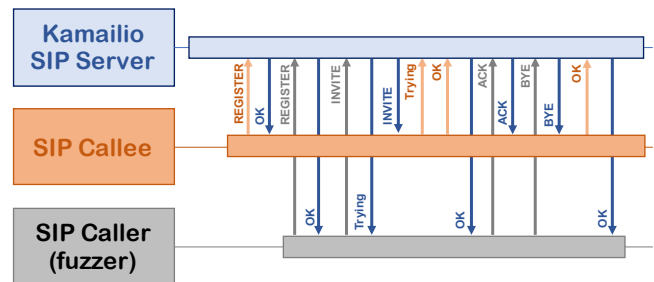


Figure 2: SIP protocol fuzzing.

and in integrating them with a fuzzer, we found the following open research issues.

**Configuration and multi-party protocols.** The fuzzing of network protocols requires the configuration of their software implementations. Selecting a specific configuration determines which parts of a network server can be tested by fuzzing. The configuration space can be quite complex for real-world network servers, but the problem has not been studied yet. Previous research has been mostly focused on stateless programs, such as libraries and command-line utilities, where the main challenge comes from the complexity of the inputs and of their constraints for reaching deep execution paths. As the configuration affects the attack surface, it indirectly becomes a concern for fuzzing.

In particular, protocols can require to run more than two parties. Existing fuzzing techniques only adopt a client-server scheme, where the fuzzer acts as client. However, in the case of SIP protocol, at least three parties are required: the server under test, a “caller” and a “callee”. One or more “callee” and “caller” must first register on the server; then, a caller can contact the server, by indicating the name of a callee; the server will act as proxy between the callee and the caller, in order to establish a communication (e.g., a VoIP call) between them. In the process, the server will handle the “ringing”, the possibility that the callee is busy, etc.

Figure 2 shows an example of interactions between the fuzzing tool and the Kamailio SIP target software. To fuzz a SIP session, which establishes a call between a caller and callee, the fuzzer should take one of these roles (in our benchmark, it is the caller). In addition to the fuzzer, it is necessary to have another process to

take the other role. Therefore, the benchmark comes with scripts to automate the execution both of the SIP server and of a callee using a SIP client; after the callee has registered, the fuzzers can replay and fuzz the sequences of message from the corpus. The corpus has been prepared to match the configuration of the processes, i.e., the same port numbers and client names are both referenced in corpus, and used by the SIP client and by the fuzzer.

In more complex protocols, such as for IP routing and network management, even more parties may be required for a meaningful testing. When setting up multiple parties, identifying a configuration for the target becomes even more challenging. Moreover, these protocols are beyond the support of current fuzzers, as they focus on fuzzing an individual stream of data to the server.

**Deterministic execution.** Another challenge towards fuzzing a complete network server is non-deterministic behavior, due to the extensive use of threads and event-based I/O. When the same input is applied to the target software, its non-determinism causes variations of the coverage of program paths, which in turn hinders coverage-driven fuzzing. For this reason, the AFL fuzzer performs a “calibration” on every new input that expands the coverage, in order to assess whether the execution is deterministic, and to report the percentage of anomalous new inputs (the *stability* metrics). In the benchmark, we followed the best practices suggested for AFL [23], by disabling threads where possible; or, by replacing kernel-level threads (as in *pthread*s) with user-level ones (as in *GNU Pth*) and related I/O APIs, as we did in the *forked-daapd* target. Despite these countermeasures, the reported stability is still lower than 50% in several benchmarks, due to the use of threads and I/O in libraries called by the target software. Moreover, user-level threads increase the complexity of the testing setup (e.g., in *forked-daapd*, it hinders the saving of coverage data).

**Test execution speed.** The execution speed of fuzz testing for network servers can be quite slow. Compared to fuzzing libraries and command-line utilities (which can reach up thousands of executions per second in AFL), fuzzing network servers is slowed down by the complex initialization of the process (e.g., to initialize a cache or a database), the networking overhead (due to communication and context switches between processes), and the long sequences of messages in a test case. For example, in fuzzing SIP, the initialization was slowed down by the synchronization between the fuzzer and the other processes (the server and the SIP callee), with less than 5 executions per second. Notably, in fuzzing DAAP, the server under test takes about 2 seconds to reliably complete its initialization hence the fuzzing speed significantly drops.

**State identification** The goal of protocol fuzzing is to generate tests that cover protocol states. One important open area of research is to perform fuzzing when a formal specification of the protocol state machine is not available (e.g., using model inference techniques, which reverse a protocol by trial-and-error). In our benchmark, we also faced the problem of identifying the current state of a protocol at run-time, when the protocol does not explicitly encode the current state in the messages (e.g., RTSP protocol).

## 5 CONCLUSION

In this paper, we present a new benchmark (*ProFuzzBench*) for stateful fuzzing of network protocols. The benchmark currently

supports 10 protocols and it can be extended. Based on our preliminary experiments, we also discuss open challenges for future research. As future work we plan to add more targets and integrate more fuzzers into the benchmark. Moreover, we will run large-scale experiments on the benchmark and analyze the results to gain more insights into the identified challenges.

## BENCHMARK AVAILABILITY

The source code of *ProFuzzBench* is available online on GitHub at: <https://github.com/profuzzbench/profuzzbench>

## REFERENCES

- [1] Humberto J Abdelnur, Radu State, and Olivier Festor. 2007. KiF: a stateful SIP fuzzer. In *Intl. Conf. on Princ., Sys. and Apps. of IP Telecom.* 47–56.
- [2] Thoulfekar Alrahem, Alex Chen, N DiGiuseppe, Jefferey Gee, Shang-Pin Hsiao, Sean Mattox, Taejoon Park, IG Harris, et al. 2007. Interstate: A stateful protocol fuzzer for SIP. *Defcon 15* (2007), 1–5.
- [3] Greg Banks, Marco Cova, Viktoria Felmetzger, Kevin Almeroth, Richard Kemmerer, and Giovanni Vigna. 2006. SNOOZE: toward a Stateful NetwOrk prOtoCol fuzZER. In *Intl. Conf. on Information Security.* 343–358.
- [4] Beyond Security. 2020. beSTORM Black Box Testing. <https://beyondsecurity.com/solutions/bestorm.html>. [Online; accessed 12-10-2020].
- [5] Brian Caswell. 2020. Cyber Grand Challenge Corpus. <http://www.lungetech.com/cgc-corpus/>. [Online; accessed 12-01-2021].
- [6] Joeri De Ruiter and Erik Poll. 2015. Protocol State Fuzzing of TLS Implementations. In *24th USENIX Security Symp.* 193–206.
- [7] Brendan Dolan-Gavitt, Patrick Hulin, Engin Kirda, Tim Leek, Andrea Mambretti, Wil Robertson, Frederick Ulrich, and Ryan Whelan. 2016. LAVA: Large-scale automated vulnerability addition. In *Symp. on Security and Privacy (SP)*. 110–121.
- [8] Paul Fiterau-Brosteau, Bengt Jonsson, Robert Merget, Joeri de Ruiter, Konstantinos Sagonas, and Juraj Somorovsky. 2020. Analysis of DTLS Implementations Using Protocol State Fuzzing. In *29th USENIX Security Symp.*
- [9] Zicong Gao, Weiyeu Dong, Rui Chang, and Yisen Wang. 2020. Fw-fuzz: A code coverage-guided fuzzing framework for network protocols on firmware. *Concurrency and Computation: Practice and Experience* (2020).
- [10] Hugo Gascon, Christian Wressnegger, Fabian Yamaguchi, Daniel Arp, and Konrad Rieck. 2015. Pulsar: Stateful black-box fuzzing of proprietary network protocols. In *Intl. Conf. on Sec. and Priv. in Comm. Sys.* 330–347.
- [11] Ahmad Hazimeh, Adrian Herrera, and Mathias Payer. 2020. Magma: A Ground-Truth Fuzzing Benchmark. *ACM Meas. Anal. Comput. Syst.* 4, 3 (2020).
- [12] George Klees, Andrew Ruef, Benji Cooper, Shiyi Wei, and Michael Hicks. 2018. Evaluating fuzz testing. In *ACM Conf. on Comp. and Comm. Security.* 2123–2138.
- [13] V. J. M. Manès, H. Han, C. Han, S. K. Cha, M. Egele, E. J. Schwartz, and M. Woo. 2019. The Art, Science, and Engineering of Fuzzing: A Survey. *IEEE Transactions on Software Engineering* (2019).
- [14] László Szekeres Jonathan Metzman, Abhishek Arya, and L Szekeres. 2020. FuzzBench: Fuzzer benchmarking as a service. *Google Security Blog* (2020).
- [15] Barton P. Miller, Louis Fredriksen, and Bryan So. 1990. An Empirical Study of the Reliability of UNIX Utilities. *Commun. ACM* (1990).
- [16] Van-Thuan Pham, Marcel Böhme, and Abhik Roychoudhury. 2020. AFLNET: A Greybox Fuzzer for Network Protocols. In *Intl. Conf. on Software Testing, Verification and Validation (Testing Tools Track)*.
- [17] Rapid7. 2020. Metasploit Vulnerability & Exploit Database. <https://www.rapid7.com/db/?q=fuzzer&type=metasploit>. [Online; accessed 12-10-2020].
- [18] Rebecca Shapiro, Sergey Bratus, Edmond Rogers, and Sean Smith. 2011. Identifying vulnerabilities in SCADA systems via fuzz-testing. In *Intl. Conf. on Critical Infrastructure Protection.* 57–72.
- [19] Synopsis, Inc. 2020. Defensics Fuzz Testing. <https://www.synopsys.com/software-integrity/security-testing/fuzz-testing.html>. [Online; accessed 12-10-2020].
- [20] Ari Takanen, Jared D Demott, Charles Miller, and Atte Kettunen. 2018. *Fuzzing for software security testing and quality assurance*. Artech House.
- [21] Andreas Walz and Axel Sikora. 2017. Exploiting dissent: Towards fuzzing-based differential black box testing of TLS implementations. *IEEE Trans. on Dep. and Sec. Comp.* (2017).
- [22] Zhiqiang Wang, Quanqi Li, Yazhe Wang, Biao Liu, Jianyi Zhang, and Qixu Liu. 2019. Medical Protocol Security: DICOM Vulnerability Mining Based on Fuzzing Technology. In *ACM Conf. on Comp. and Comm. Security.* 2549–2551.
- [23] Michal Zalewski. 2020. AFL Documentation - Understanding the status screen. [https://github.com/mirrorer/afl/blob/master/docs/status\\_screen.txt](https://github.com/mirrorer/afl/blob/master/docs/status_screen.txt). [Online; accessed 12-01-2021].
- [24] Yaowen Zheng, Ali Davanian, Heng Yin, Chengyu Song, Hongsong Zhu, and Limin Sun. 2019. FIRM-AFL: High-throughput greybox fuzzing of IoT firmware via augmented process emulation. In *28th USENIX Security Symp.* 1099–1114.