

**Methodology for Characterizing
Network Behavior of Internet of Things
Devices**

Paul Watrobski
Murugiah Souppaya
Joshua Klosterman
William Barker

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8349-draft>

Methodology for Characterizing Network Behavior of Internet of Things Devices

Paul Watrobski
*Applied Cybersecurity Division
Information Technology Laboratory*

Murugiah Souppaya
*Computer Security Division
Information Technology Laboratory*

Joshua Klosterman
*The MITRE Corporation
McLean, VA*

William Barker
*Dakota Consulting
Gaithersburg, MD*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8349-draft>

January 2022



U.S. Department of Commerce
Gina Raimondo, Secretary

National Institute of Standards and Technology
*James K. Olthoff, Performing the Non-Exclusive Functions and Duties of the Under Secretary of Commerce
for Standards and Technology & Director, National Institute of Standards and Technology*

53 National Institute of Standards and Technology Interagency or Internal Report 8349
54 43 pages (January 2022)

55 This publication is available free of charge from:
56 <https://doi.org/10.6028/NIST.IR.8349-draft>

57 Certain commercial entities, equipment, or materials may be identified in this document in order to describe an
58 experimental procedure or concept adequately. Such identification is not intended to imply recommendation or
59 endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best
60 available for the purpose.

61 There may be references in this publication to other publications currently under development by NIST in accordance
62 with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies,
63 may be used by federal agencies even before the completion of such companion publications. Thus, until each
64 publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For
65 planning and transition purposes, federal agencies may wish to closely follow the development of these new
66 publications by NIST.

67 Organizations are encouraged to review all draft publications during public comment periods and provide feedback to
68 NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at
69 <https://csrc.nist.gov/publications>.

70 **Public comment period:** January 11, 2022 – February 11, 2022

71 **Submit comments on this publication to:** iot-ddos-nccoe@nist.gov

72 National Institute of Standards and Technology
73 Attn: Applied Cybersecurity Division, Information Technology Laboratory
74 100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-2000

75 All comments are subject to release under the Freedom of Information Act (FOIA).

76

Reports on Computer Systems Technology

77 The Information Technology Laboratory (ITL) at the National Institute of Standards and
78 Technology (NIST) promotes the U.S. economy and public welfare by providing technical
79 leadership for the Nation’s measurement and standards infrastructure. ITL develops tests, test
80 methods, reference data, proof of concept implementations, and technical analyses to advance
81 the development and productive use of information technology. ITL’s responsibilities include the
82 development of management, administrative, technical, and physical standards and guidelines for
83 the cost-effective security and privacy of other than national security-related information in
84 federal information systems.

85

Abstract

86 This report describes an approach to capturing and documenting the network communication
87 behavior of Internet of Things (IoT) devices. From this information, manufacturers, network
88 administrators, and others can create and use files based on the Manufacturer Usage Description
89 (MUD) specification to manage access to and from those IoT devices. The report also describes
90 the current state of implementation of the approach and proposals for future development.

91

Keywords

92 access control; device characterization; Internet of Things (IoT); Manufacturer Usage
93 Description (MUD); network communications.

94

Disclaimer

95 Any mention of commercial products or reference to commercial organizations is for information
96 only; it does not imply recommendation or endorsement by the National Institute of Standards
97 and Technology (NIST), nor does it imply that the products mentioned are necessarily the best
98 available for the purpose.

99

Additional Information

100 For additional information on NIST’s cybersecurity programs, projects, and publications, visit
101 the [National Cybersecurity Center of Excellence](#) (NCCoE) and the [Computer Security Resource](#)
102 [Center](#). Information on other efforts at [NIST](#) and in the [Information Technology Laboratory](#)
103 (ITL) is also available.

104

Supplemental Content

105 NIST’s NCCoE created MUD-PD, a tool to assist in developing MUD files. The tool is
106 described in greater detail in Section 3.2. See GitHub: [https://www.github.com/usnistgov/MUD-](https://www.github.com/usnistgov/MUD-PD)
107 [PD](#).

108

Acknowledgments

109 The authors wish to thank the following individuals for their generous contributions of expertise
110 and time: Luca Deri, Institute for Informatics and Telematics; Donna Dodson and William Polk,
111 NIST; Karen Scarfone, Scarfone Cybersecurity; Parisa Grayeli, Blaine Mulugeta, and Susan
112 Symington, The MITRE Corporation; Hassan Habibi Gharakheili, University of New South
113 Wales; and Russ Housley, Vigil Security, LLC.

114

Audience

115 This report is written for those who would like to build, create, or utilize MUD files, including:

- 116 • IoT device manufacturers and developers;
- 117 • network administrators;
- 118 • IoT device vulnerability researchers and analysts;
- 119 • network equipment developers and manufacturers; and
- 120 • service providers that develop and utilize components based on the MUD specification.

121

Document Conventions

122 This report utilizes several terms for which contradictory or generic definitions exist in literature.
123 For purposes of this paper, the following definitions have been coined or adopted:

124 **Characterizing** is the act of collecting, analyzing, and/or storing information intended to be used
125 in describing behavior and/or characteristics pertaining to a device.

126 **Fingerprinting** is the act of collecting information intended to help uniquely identify a device
127 type.

128 A **MUD file** contains information that describes an IoT device and its network behavior, as
129 described in the MUD specification [1]. The term “MUD profile” is used throughout existing
130 literature and is synonymous with “MUD file.” This paper adheres to the use of “MUD file” as
131 defined in the MUD specification.

132 **MUD file accuracy** describes how precisely a MUD file captures the full communication
133 requirements of an IoT device—in particular, the extent to which it lists all potential
134 communications that the device may need to perform its intended functions (comprehensiveness)
135 and the extent to which it avoids listing communications that the device does not need
136 (correctness). Note that it may be impossible to ensure complete accuracy of a MUD file even if
137 the file is created by the manufacturer of the device. For some devices, it may be impractical or
138 even impossible to test every possible situation or network configuration capable of altering
139 device behavior. In addition, potential communication requirements that would be revealed by
140 those situations may remain unknown.

141

Trademark Information

142 All registered trademarks or trademarks belong to their respective organizations.

143

Call for Patent Claims

144 This public review includes a call for information on essential patent claims (claims whose use
145 would be required for compliance with the guidance or requirements in this Information
146 Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be
147 directly stated in this ITL Publication or by reference to another publication. This call also
148 includes disclosure, where known, of the existence of pending U.S. or foreign patent applications
149 relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

150

151 ITL may require from the patent holder, or a party authorized to make assurances on its behalf,
152 in written or electronic form, either:

153

154 a) assurance in the form of a general disclaimer to the effect that such party does not hold
155 and does not currently intend holding any essential patent claim(s); or

156

157 b) assurance that a license to such essential patent claim(s) will be made available to
158 applicants desiring to utilize the license for the purpose of complying with the guidance
159 or requirements in this ITL draft publication either:

160

161 i. under reasonable terms and conditions that are demonstrably free of any unfair
162 discrimination; or

163

164 ii. without compensation and under reasonable terms and conditions that are
165 demonstrably free of any unfair discrimination.

165

166 Such assurance shall indicate that the patent holder (or third party authorized to make assurances
167 on its behalf) will include in any documents transferring ownership of patents subject to the
168 assurance, provisions sufficient to ensure that the commitments in the assurance are binding on
169 the transferee, and that the transferee will similarly include appropriate provisions in the event of
170 future transfers with the goal of binding each successor-in-interest.

171

172 The assurance shall also indicate that it is intended to be binding on successors-in-interest
173 regardless of whether such provisions are included in the relevant transfer documents.

174

175 Such statements should be addressed to: iot-ddos-nccoe@nist.gov

176 Executive Summary

177 Characterizing and understanding the expected network behavior of Internet of Things (IoT)
178 devices is essential for cybersecurity purposes. It enables the implementation of appropriate
179 network access controls (e.g., firewall rules or access control lists) to protect the devices and the
180 networks on which they are deployed. This may include limiting a device's communication to
181 only that which is deemed necessary. It also enables identifying when a device may be
182 misbehaving, a potential sign of compromise. The ability to restrict network communications for
183 IoT devices is critically important, especially given the increased number of these devices.

184 Network behavior for most IoT devices is situation-dependent. For example, many IoT devices
185 have multiple mechanisms for interaction and control, such as voice commands, physical
186 interaction with a person, other devices (e.g., a smartphone or IoT hub), and services (e.g., cloud-
187 based). Any given action may result in different network behavior, depending on the mechanism
188 through which it was performed. Additionally, certain patterns of network behavior may only
189 occur in specific stages of a device's lifecycle (i.e., setup, normal operation, and
190 decommissioning). Also, network behavior may change over time as device software is updated.
191 For these reasons, the expected network behavior of a device needs to be characterized and
192 understood for all intended scenarios and during each stage of its lifecycle. Otherwise, necessary
193 steps for device setup, operation, or decommissioning may be blocked by network access
194 controls, preventing them from being performed fully or at all.

195 This publication describes recommended techniques for IoT device manufacturers and
196 developers, network administrators, and researchers to accurately capture, document, and
197 characterize the entire range of a device's network behavior in MUD (Manufacturer Usage
198 Description) files. MUD provides a standard way to specify the network communications that an
199 IoT device requires to perform its intended functions. MUD files tell the organizations using IoT
200 devices what access control rules should apply to each IoT device, and MUD files can be
201 automatically consumed and used by various security technologies.

202 This publication also presents a National Cybersecurity Center of Excellence (NCCoE)
203 developed open-source tool, MUD-PD, that can be used to catalog and analyze the collected
204 data, as well as generate both reports about the device and deployable MUD files. This tool is
205 intended to aid IoT device manufacturers and developers, network administrators, and
206 researchers who want to create or edit MUD files.

207 **Table of Contents**

208 **Executive Summary v**

209 **1 Introduction 1**

210 1.1 Challenges 1

211 1.2 Purpose and Scope 2

212 1.3 Report Structure 3

213 **2 Network Traffic Capture Methodology 4**

214 2.1 Capture Strategy 4

215 2.1.1 IoT Device Life-Cycle Phases 4

216 2.1.2 Environmental Variables 5

217 2.1.3 Activity-Based and Time-Based Capture Approaches 7

218 2.1.4 Network Architecture and Capture Approach 7

219 2.1.5 Capture Tools 8

220 2.2 Capture Procedure 9

221 2.2.1 Device Setup Capture 9

222 2.2.2 Normal Operation Capture 9

223 2.2.3 Decommissioning/Removal Capture 10

224 2.3 Documentation Strategy 10

225 **3 Analysis Use Cases and Tools 12**

226 3.1 Manual MUD File Generation 12

227 3.1.1 Wireshark 12

228 3.1.2 NetworkMiner 12

229 3.1.3 Overview of Manual MUD File Generation Process 12

230 3.2 MUD-PD 13

231 3.2.1 Current Feature Set 14

232 3.2.2 GUI Overview 15

233 3.2.3 MUD-PD Uses 24

234 3.3 MUD-PD Support for Privacy Analysis 25

235 **4 Future Work 26**

236 4.1 Extending MUD-PD Features 26

237 4.2 Developing a MUD Pipeline 26

238 4.3 Open Problems for the Community 29

239 **References 31**

240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268

List of Appendices

Appendix A— Example Capture Environment..... 33
Appendix B— Acronyms 34

List of Figures

Figure 1: MUD-PD main window with buttons and list boxes labeled..... 16
 Figure 2: Prompt for providing Fingerbank API key..... 16
 Figure 3: Prompt for creating a new database 16
 Figure 4: Prompt for connecting to an existing database 17
 Figure 5: Prompt for importing packet captures into database 18
 Figure 6: Window listing devices imported and to import during the packet capture
 import process..... 19
 Figure 7: Window prompt for importing a device 20
 Figure 8: Window prompting to update the firmware version logged in the database ... 20
 Figure 9: Prompt for generating a human-readable device report..... 21
 Figure 10: Example device report showing the details of a single packet capture 21
 Figure 11: Prompt for selecting a device for which the MUD file will be generated 22
 Figure 12: Prompt for providing device details including the support and document
 URLs 22
 Figure 13: Prompt for providing internet communication rules 23
 Figure 14: Prompt for providing local communication rules..... 23
 Figure 15: Preview of the MUD file to be generated..... 24
 Figure 16: MUD pipeline for the device manufacturer or developer use case 27
 Figure 17: MUD pipeline for the network administrator use case 28
 Figure 18: The overarching MUD pipeline, particularly as it may be used for research
 and development..... 28
 Figure 19: Example capture architecture..... 33

269 **1 Introduction**

270 The National Institute of Standards and Technology's (NIST's) National Cybersecurity Center of
271 Excellence (NCCoE) is working to improve the ability of network administrators and operators
272 of Internet of Things (IoT) networks to identify, understand, and document network
273 communication requirements of IoT devices. Documenting the types of devices and
274 communication behaviors of those devices can allow creation of files based on the Manufacturer
275 Usage Description (MUD) specification, which can be used by network administrators to
276 manage access to and from those devices [1].

277 The [Document Conventions](#) section earlier in this report defines several terms used throughout
278 the report. Readers should review those definitions before proceeding.

279 **1.1 Challenges**

280 For network administrators to properly secure networks, they need to understand what devices
281 are on the network in question and what network communication each device requires to perform
282 its intended functions. In the case of networks that include IoT devices, it is often difficult to
283 identify each individual device, much less know what access is required by each device to other
284 network components, and what access other network components need to each device. To
285 address this challenge, many organizations are implementing IoT device fingerprinting and
286 characterization methods to identify the types of devices on a network.

287 Once the IoT device type is known for each device, the network administrator can begin to
288 manage security and access control for the devices [2]. This involves collecting information
289 regarding the devices' characteristics and behavior. Approaches like those of the Princeton IoT
290 Inspector [3] and ProfilIoT's use of machine learning [4] are being used to characterize and
291 identify IoT devices, which can provide insight into security and privacy issues associated with
292 each device. However, not all fingerprinting and characterization schemes are equivalent. These
293 schemes are often created based on a limited set of data derived from network traffic that allows
294 them to accurately identify just the device type. The network traffic information used to develop
295 these schemes include packet headers, network ports, packet timing, handshakes, and other
296 information that might be unique to a particular IoT device [5], [6]. Given the limited set of data
297 used to develop the fingerprints, the fingerprints do not contain the information necessary to
298 determine a device's full range of potential behaviors.

299 Comprehensively describing the characteristics of IoT devices is made difficult by several
300 factors. For example, IoT devices are often subject to internal changes that may affect their
301 behavior. These changes can be caused by software updates, firmware updates, new or
302 supplemental hardware, and so on. External changes can also occur with hardware replacements,
303 integrations with other IoT devices, connections to new networks, and more. These changes can
304 increase the complexity involved in tracking an IoT device's behavior and, by extension,
305 increase the difficulty of accurately characterizing an IoT device. User activities can also
306 significantly affect an IoT device's behavior. For example, two cameras created by the same
307 manufacturer may display drastically different behaviors if they are used for different purposes.
308 Additionally, behaviors may be distinct for different firmware or hardware revisions of the same
309 device. Many IoT devices are also created as variants based on the design of an existing IoT

310 device, which can make their behaviors appear similar, even if the IoT devices are technically
311 distinct from one another.

312 The goal of the MUD specification [1] is to provide a standard method for IoT devices to “signal
313 to the network the access and network functionality they require to properly function.” This is
314 accomplished by using a MUD file, which can allow a network administrator to know what
315 access control rules should apply to the IoT device. If a network administrator enforces an
316 inaccurate MUD file, the functionality of the device can be severely impaired or potentially lead
317 to vulnerabilities. Therefore, it is imperative that any MUD file be as accurate as possible.

318 A MUD file’s accuracy is based on two concepts: *comprehensiveness*—the extent to which it
319 lists all potential network communications that the device may need to perform its intended
320 functions, and *correctness*—the extent to which it avoids listing network communications that
321 the device does not need. However, because a manufacturer may not be able to predict all
322 operational environments in which a device is used, there is no guarantee that all manufacturer-
323 provided MUD files are comprehensive. The final decision of what actions a device may perform
324 is ultimately up to the local network administrator [1] tasked with implementing the device; they
325 may decide that the device’s MUD file should be more or less restrictive than the MUD file
326 provided by the manufacturer. Additionally, a network administrator may wish to create a MUD
327 file for a device without a manufacturer-provided MUD file.

328 1.2 Purpose and Scope

329 This report describes a way to build an accurate MUD file based on network traffic data that
330 reveals information about the IoT device’s potential network behavior. Developing MUD files
331 consists of two major steps: traffic capture and traffic analysis. The methodology described in
332 the report is designed to create an accurate set of network traffic data, capturing as much of the
333 IoT device’s potential behavior as possible. The methodology seeks to allow for analysis of the
334 full range of IoT device network traffic behaviors that can reasonably be expected. This includes
335 examining a variety of factors that could potentially alter an IoT device’s behavior at each stage
336 of the device’s life cycle.

337 Developers, network administrators, and researchers can take advantage of the methodology to
338 develop a comprehensive data set that can be used for generating MUD files, investigating
339 security and privacy concerns, developing machine learning algorithms, and more. The
340 methodology described has been developed on Internet Protocol (IP)-based networks, but it can
341 potentially be utilized with other types of networks as well. It is important to note that this type
342 of analysis assumes that

- 343 • the IoT devices have not been tampered with or compromised by a malicious actor at any
344 point in the analysis process, and
- 345 • the IoT devices are operating as their manufacturers intended.

346 In addition to prescribing a methodology for capturing an IoT device’s behavior on a network,
347 this report also explores how the NCCoE-developed MUD-PD tool can leverage this behavior
348 information to create MUD files. MUD-PD requires a diverse set of network traffic captures to
349 generate accurate MUD files. The tool extracts and aggregates pertinent information that allows
350 creation of accurate MUD files without manually parsing a large set of network traffic data. This

351 tool can drastically reduce the time and effort required to generate MUD files compared with
352 manually creating MUD files.

353 Enforcement of rules generated from the MUD file is outside the scope of this report, but several
354 different approaches are described in the NCCoE preliminary draft Practice Guide, Special
355 Publication 1800-15, *Securing Small Business and Home Internet of Things (IoT) Devices:
356 Mitigating Network-Based Attacks Using Manufacturer Usage Description (MUD)* [7].

357 IoT device detection and identification is also out of scope, other than a description in Section
358 2.1.5 of two tools that support manual device identification and analysis.

359 **1.3 Report Structure**

360 The rest of this report is organized into the following sections and appendices:

- 361 • Section 2 discusses traffic capture strategy, tools, example procedures, and
362 documentation.
- 363 • Section 3 discusses analysis of network communications, privacy implications, and MUD
364 file generation using the MUD-PD tool.
- 365 • Section 4 explores possible future work, such as developing enhancements and additional
366 features of MUD-PD, and continuing research in the area of device characterization.
- 367 • The References section defines the references cited throughout the publication.
- 368 • Appendix A presents an example capture environment that supports analysis of both
369 wired and wireless IoT devices.
- 370 • Appendix B contains an acronym list.

371 **2 Network Traffic Capture Methodology**

372 Properly generating an accurate MUD file requires a comprehensive data set that reflects the
373 greatest possible range of intended device behaviors for each networked device. In the case of
374 MUD files that can and will be used for network security and access control, it is imperative that
375 each generated file be sufficiently accurate to prevent false reporting of legitimate network
376 activity and placing restrictions on devices that may prevent them from functioning properly.
377 The methodology described in this section is designed to support capture of the information
378 needed for IoT device analysis and MUD file generation.

379 This methodology is based on network traffic and does not account for device behavior that
380 cannot be observed from network traffic. Observed device behaviors outside the scope of this
381 methodology should be documented through other means.

382 **2.1 Capture Strategy**

383 Capturing a wide range of intended device behaviors requires that communications to and from
384 the IoT device be captured under a wide range of states and environmental conditions throughout
385 the device life cycle. This section describes network traffic capture approaches, strategies, and
386 tools. The information listed in this section should be documented for each capture activity for
387 each IoT device to support analysis of the device's behavior.

388 **2.1.1 IoT Device Life-Cycle Phases**

389 Various taxonomies are used to describe IoT device life cycles, but this report organizes device
390 life-cycle components into three broad phases for IoT device traffic analysis: setup, normal
391 operation, and decommissioning/removal.

392 **2.1.1.1 Setup**

393 The setup phase includes everything needed to initially connect an IoT device to a network and
394 to take configuration actions necessary for the device to be fully functional and ready to begin
395 normal operations. Setup typically begins with a wired or wireless connection of the device to
396 the network. Once the device is connected, setup processes can include firmware updates;
397 connections to smart hubs, smartphones, and other devices; and other processes that must be
398 completed. While following the manufacturer's instructions may be adequate for most situations
399 involving setup behaviors, deviation from those instructions may be necessary to capture the
400 device's behavior under some circumstances (e.g., not connecting an IoT device to an associated
401 cloud service may result in unique behavior for devices that a manufacturer assumes will be
402 connected to a cloud service). Initial connection to cloud/internet-based services may be required
403 for some devices. This phase may also include connection of an IoT device to a smartphone or
404 another device that is expected to manage the device (such as a controller/smart hub).

405 Setup failure situations can also produce connectivity behaviors different from those anticipated
406 by the manufacturer. For example, a device that is configured to connect with a controller, smart
407 hub, or cloud service may be unable to do so for any number of reasons, including lack of
408 internet connection and blocked ports.

409 **2.1.1.2 Normal Operation**

410 The “normal operation” phase captures an IoT device’s behavior for the majority of its service
411 life after it has been set up and is performing its intended functions. This phase covers a wide
412 range of behaviors, such as human-to-device interactions, controller or smart hub-to-device
413 interactions, and cloud service-to-device interactions. It also covers device-initiated behaviors
414 that can occur without human interaction. Software and firmware updates may occur with or
415 without human initiation or interaction and can cause an intended change in device behavior.
416 Capture of both human-initiated updates and automatic updates is important, though capture of
417 automatic updates may be the more challenging. Other types of interactions during normal
418 operation may include remote control through smartphones and cloud-based services. Normal
419 operation failure situations, such as being unable to access required resources, can also produce
420 anomalous behaviors. “Unexpected” scenarios, including removing essential devices, removing
421 the controller/smart hub, or performing a hard reset on the IoT device, are still considered normal
422 operation and should also be examined.

423 **2.1.1.3 Decommissioning/Removal**

424 The final phase in an IoT device’s life cycle (before the device is reused elsewhere or reaches
425 end-of-life) includes the process of de-registering the IoT device from other devices, such as
426 controllers/smart hubs, and/or cloud services (decommissioning) and removing it from the
427 network (removal). If manufacturer instructions for this process exist, they should be included as
428 part of the capture-planning process if possible. If no instructions exist, a factory reset is
429 generally the recommended procedure for decommissioning and removal. In either case, a
430 factory reset should be included as part of the capture-planning process. Factory reset brings the
431 device back to its initial configuration. (Note: Firmware updates may not be rolled back during
432 the factory reset process.)

433 This report treats the factory reset process as an element of the decommissioning/removal phase
434 because a factory reset can sometimes de-register the device from a cloud service and/or
435 disconnect the IoT device from the network. Inclusion of other types of removal situations is also
436 recommended because IoT devices can sometimes be removed from a network without taking
437 prior decommissioning actions. If the device is used in a different role or by a new owner,
438 subsequent actions are treated here as falling within a new setup phase. Capture plans should
439 cover both device-initiated behaviors and behaviors triggered by human interaction during
440 decommissioning and removal.

441 **2.1.2 Environmental Variables**

442 The IoT device should be examined under a wide variety of environmental conditions to capture
443 the largest possible range of intended device behaviors. For example, if an IoT device is not
444 permitted access to the internet, it may not be able to complete some of the communications on
445 which it relies to function as intended (e.g., cloud-based manufacturer support services or
446 network time services). This can cause the IoT device to exhibit different behaviors on the
447 network than those originally anticipated or documented by the manufacturer. As discussed in
448 Section 1.3, there is currently no guarantee that the manufacturer-provided MUD file will cover
449 every communication pattern that the device may exhibit. For example, it is possible that the

450 device's apparent behavior may have changed due to updates of third-party libraries. Behaviors
451 like this need to be captured to provide a more accurate characterization of the IoT device.

452 This subsection provides an example set of environmental variables that can be applied during
453 each of the three life-cycle phases described in Section 2.1.1. This is not a complete list, but
454 depending on the device type and design, each of the variables has the potential to change the
455 behavior of an IoT device. For consistency and to limit confusion, these variables should persist
456 throughout the duration of a network traffic capture process and should not be added or removed
457 after the capture has begun. There are exceptions to this rule, such as capturing behaviors when
458 emulating an internet outage. Any deviations from persistent variables should be clearly
459 documented.

- 460 • **No internet** removes internet access from the local network to which the IoT device is
461 connected. This can limit an IoT device's access to resources.
- 462 • **Preferred DNS servers blocked** tests a device's behavior when its preferred Domain
463 Name System (DNS) servers have been blocked. For example, an IoT device may be
464 configured to rely on DNS servers managed by the manufacturer. If access to these DNS
465 servers is restricted, the IoT device's functionality will be reduced unless compensating
466 measures are taken.
- 467 • **Device isolation** indicates that the device is alone on the local network; that is, no other
468 devices are connected except essential network or other communication components
469 needed for the IoT device to function properly. For example, if the IoT device needs to be
470 controlled by a controller/smart hub or smartphone, this device may also be connected
471 during the capture.
- 472 • **No human interaction** means that no human interaction or configuration of the device
473 has taken place for the duration of the capture activity. The device will not be
474 preprogrammed by the analysts to take any actions prior to the start of the capture
475 process.
- 476 • **Controller/smart hub control** indicates that the device has been or will be connected to
477 a controller/smart hub during the capture. An IoT device connected to a controller/smart
478 hub will typically display different behavior than a device that is not connected.
- 479 • **Same manufacturer** means that at least one device from the same manufacturer has been
480 connected to the network before the capture has begun. It is likely that a network may
481 have two IoT devices from the same manufacturer. Additionally, many manufacturers
482 have been working to create their own IoT "ecosystems." Because some IoT devices are
483 designed to communicate with other IoT devices from the same manufacturer, connecting
484 multiple devices from the same manufacturer may reveal additional behavior not seen
485 when only one device from that manufacturer is connected to the network.
- 486 • **Full network** indicates that enough active devices to simulate an IoT application are
487 connected to the local network before the beginning of the capture. As the purpose and
488 scope of networks that support IoT devices can vary widely and are often application-
489 dependent, it is up to the analyst to determine how many and/or what variety of devices is
490 considered a full network. The presence of other devices on the same network may affect
491 the behavior of IoT devices being characterized.

- 492 • **Notable physical environment** indicates that the physical environment has changed
493 significantly before or during the packet capture. Many IoT devices contain sensors that
494 track aspects of the physical environment, including light, temperature, and sound.

495 **2.1.3 Activity-Based and Time-Based Capture Approaches**

496 Activity-based captures are focused on IoT device behavior solely during a specified set of
497 actions. For example, capturing IoT device setup behaviors does not require a specific amount of
498 time; its beginning and completion are determined only by the duration of the setup process.

499 Time-based captures are focused on capturing IoT device behavior during a specific time period.
500 For example, capturing IoT device behaviors throughout an entire day of normal operation can
501 allow observation and documentation of a wide range of behaviors (e.g., device-initiated
502 behaviors). Some behaviors may be observed only over a longer term. One example of this
503 property involves devices that “learn” the user’s behavior and modify functionality accordingly.
504 These devices may behave in a different way over the weekend than during the week or when the
505 learned pattern is broken, such as on a holiday or when the user is traveling for an extended
506 period.

507 **2.1.4 Network Architecture and Capture Approach**

508 The ideal capture activity will capture the network traffic among all hosts on the local network
509 and all communications entering and leaving the local network. In cases of smaller and/or
510 simpler networks, capture of network traffic directly from a single gateway may be sufficient
511 because the gateway will receive all communication both to and from the local network and
512 among all network devices. An example of a capture setup using a single gateway can be found
513 in Appendix A. In larger or more complex networks where network traffic does not flow through
514 a single gateway, capture of network traffic from multiple locations throughout the network is
515 recommended where possible. These capture locations should be carefully chosen to ensure that
516 all relevant traffic can be properly captured.

517 The capture approach adopted may depend on the hardware available. The capture device will
518 need sufficient resources to store all captured traffic. The absence of sufficient processing power,
519 memory, or storage is likely to cause network packets to be dropped and may compromise the
520 accuracy and integrity of the capture.

521 In some cases, a device may have additional network interfaces that enable communication that
522 cannot be observed by the local network gateways. For example, a ZigBee hub may interface
523 with a ZigBee network as well as a Wi-Fi network. Ideally the traffic on both networks should be
524 captured for analysis. In some instances, a device’s secondary interface enables communication
525 to an entirely external network, as in the case of 3G, 4G, and 5G devices. It is ideal to capture
526 this communication as well, but it may be difficult or impossible to do so. In any case, however,
527 it is important to document any additional network interfaces a device may have, as they may be
528 alternative vectors for information to travel. Documentation procedures are discussed in depth in
529 Section 2.3.

530 Once network capture locations have been determined, the method of capture should be chosen.
531 Capture of traffic directly on the chosen gateway/router/switch is ideal if the network device’s

532 resources are sufficient for the task. This allows capturing network traffic from any or all of the
533 Ethernet ports and wireless radios managed by the network device and saving the captured
534 information directly. It is not always possible to capture traffic directly on the network device,
535 but alternatives are available for situations that do not permit capture in this manner. For
536 example, placing a network tap in-line on a wired IoT device can provide access to the desired
537 communication. Another alternative is using a mirrored or switched port analyzer (SPAN) port to
538 send all traffic from a port or virtual local area network to a capture device that is listening on a
539 selected port. For IoT devices that communicate over a wireless network, using a wireless
540 network adapter in promiscuous mode will allow capture of wireless traffic. However, wireless
541 capture is not always an ideal option, as there may be instances where interference with
542 capturing wireless traffic is unavoidable (e.g., due to wireless isolation being used).

543 **2.1.5 Capture Tools**

544 Various tools are available for capturing network traffic. Two of the most widely used are
545 tcpdump and Wireshark.

546 **2.1.5.1 tcpdump**

547 tcpdump is a lightweight command-line-based tool that can be used on Cisco IOS, Junos OS, and
548 many Linux-based router and switch operating systems. Packet captures (pcaps) can be saved to
549 a standard pcap file format, which is commonly used to store network traffic data. The following
550 command demonstrates tcpdump usage:

```
551 bash$ tcpdump -i eth0 -s0 -n -B 2000000 -w capture.pcap
```

- 552 • “tcpdump” starts the capture program.
- 553 • “-i eth0” instructs tcpdump to start capturing packets from the interface eth0.
- 554 • “-s0” sets the snapshot length to an unlimited size, allowing capture of larger packets.
555 tcpdump normally truncates IPv4 packets that are larger than 68 bytes.
- 556 • “-n” turns off host name resolution, which reduces the processing and buffer resources
557 needed to capture properly.
- 558 • “-B 2000000” sets the operating system capture buffer size to 2,000,000 kibibytes,
559 allowing capture of a greater amount of network traffic. Packet drops can still occur in
560 the driver and in the kernel, so it is important to ensure the capture hardware is adequate
561 to the task.
- 562 • “-w capture.pcap” saves network traffic to a file named capture.pcap.

563 **2.1.5.2 Wireshark**

564 Wireshark is one of the most readily available packet capture and analysis tools, and it is open
565 source. Wireshark provides a graphical user interface (GUI) during both capture and analysis. It
566 also has a command-line-based capture utility called tshark, which can perform both capture and
567 analysis functions.

568 Wireshark is supported by Windows, macOS, and a wide range of Unix and Unix-like platforms,
569 including Linux and Berkeley Software Distribution (BSD). Use of Wireshark as a capture tool
570 often involves setting up a mirrored/SPAN port or a network tap to ensure that Wireshark can
571 capture as much relevant network traffic as possible. Wireshark also supports putting network
572 interfaces into promiscuous mode, which is often necessary to properly capture wireless network
573 traffic. Wireshark supports the PCAP Next Generation Dump (PcapNg) file format, which allows
574 addition of metadata to network traffic captures. See Section 2.3 for further details.

575 **2.2 Capture Procedure**

576 This section lists example procedures for capturing network traffic. These examples focus on
577 capturing directly from a router. They are purposely generalized to be applicable to many
578 situations and may be modified/customized as required. See Appendix A for an example of a
579 network in which these procedures could be used.

580 **2.2.1 Device Setup Capture**

581 Device setup captures are mainly activity-based. An example process for this capture type is as
582 follows:

- 583 1. Select, implement, and document environmental variables to be used for this capture.
- 584 2. Start packet capture on router.
- 585 3. Begin device setup according to manufacturer instructions.
- 586 4. Complete device setup.
- 587 5. End packet capture.
- 588 6. Transfer packet capture file from router to external storage for analysis.

589 **2.2.2 Normal Operation Capture**

590 Capture of normal operation can be either activity-based or time-based. An example process for
591 this capture type is as follows:

- 592 1. Select, implement, and document environmental variables to be used for this capture.
- 593 2. Start packet capture on router.
- 594 3. Begin normal operation for device (following manufacturer directions, if available).
- 595 4. Document actions/activity taken.
- 596 5. End device operations.
- 597 6. End packet capture.
- 598 7. Transfer packet capture file from router to external storage for analysis.

599

600 2.2.3 Decommissioning/Removal Capture

601 Decommissioning/removal captures are mainly activity-based. An example process for this
602 capture type is as follows:

- 603 1. Select, implement, and document environmental variables to be used for this capture.
- 604 2. Start packet capture on router.
- 605 3. Begin decommissioning process for device (remove from smartphone application/smart
606 hub/cloud service).
- 607 4. End decommissioning process.
- 608 5. Remove the device from the network.
- 609 6. End packet capture.
- 610 7. Transfer packet capture file from router to external storage for analysis.

611 2.3 Documentation Strategy

612 After each network traffic capture has been completed, it is important to ensure that the
613 conditions and other applicable details are thoroughly documented and linked to each packet
614 capture. Documenting the life-cycle phase, environmental variables involved, and other
615 important factors can greatly help with subsequent analysis of the network traffic. Options for
616 recording this information include editing the file name, using a text document, storing
617 information in a database, or recording metadata to the capture file itself.

618 Note that the MUD specification does not include mechanisms for allowing or blocking traffic
619 under specific conditions. However, it may be useful to a network administrator to be able to
620 trace network activity to a particular event. For a situation like this, and to gain a better
621 understanding of a device's behavior, it is important to keep a log of the activities, actions
622 performed, and environmental variables during each capture.

623 There are a number of ways to document this information. The simplest is to manually write
624 descriptions for each capture and store the text documents along with the captures. This approach
625 is not scalable and may lead to mistakes where capture-document pairs are separated. An
626 alternative is to use the comment field in the PcapNg. PcapNg extends the capabilities of the
627 libpcap format. Wireshark can convert pcap files to PcapNg, and comments can be added by
628 using the GUI. The terminal-based interface to Wireshark, tshark, allows inclusion of comments
629 while taking a network capture. The following command allows insertion of a text description of
630 the capture environment and variables. This way, the information is contained within the capture
631 itself.

```
632 bash$ tshark -w capture.pcapng --capture-comment "Example comment."
```

- 633 • The same `-i`, `-s`, `-n`, and `-B` options used in Section 2.1.5.1 (tcpdump) can be used
634 here.
- 635 • The default file type for tshark captures is PcapNg.
- 636 • The `--capture-comment` option allows text comments to be added during a capture.

637 Use of the comment field in PcapNg may still not be an optimal solution. PcapNg is limited in
638 that it requires further manual interaction for the information to be consumed and used by
639 interested parties. As the comment field allows arbitrary text input, it is possible to embed
640 information in JavaScript Object Notation (JSON) format. JSON is computer parsable/readable.
641 Consequently, the NCCoE developed a Python-based tool to format the desired information as
642 JSON and insert it into the comment field of a PcapNg file. This tool is included with MUD-PD,
643 which is described in Section 3.2. This can be initiated at the start of a capture or inserted
644 afterwards; however, the tool inserts the information into an existing file. As JSON is somewhat
645 human-readable and the data being added is fairly simple, a user can still understand the
646 necessary information from the output. An example format is as follows:

```
647     {  
648         "details": "Example of capture details",  
649         "lifecyclePhase": "normal operation",  
650         "internet": "True",  
651         "humanInteraction": "True",  
652         "preferredDNS": "True",  
653         "isolated": "False",  
654         "controllerHub": "False",  
655         "mfrSame": "True",  
656         "fullNetwork": "False",  
657         "physicalChanges": "False",  
658         "durationBased": "True",  
659         "duration": "60 seconds",  
660         "actionBased": "False",  
661         "action": ""  
662     }
```

663 This format aligns with the Python dictionary (dict) datatype which enables easy reading and
664 writing. As such, if a dictionary object, `envi_vars`, is defined as the example above, it can be
665 inserted into a packet capture file as follows:

```
666     python3> import src.pcapng_comment  
667     python3> insert_comment(filename_in="./capture.pcap",  
668                             comment=envi_vars,  
669                             filename_out="./capture_commented.pcapng")
```

- 670
- 671 • `filename_in` is the existing pcap or PcapNg capture file.
 - 672 • `comment` is the Python dict object containing the metadata.
 - 673 • `filename_out` is an optional input that defines the name of the outputted file. If omitted,
674 the output filename will be the input filename without the file extension and
“_commented.pcapng” appended to the end.

675 The tool can insert the metadata into either a pcap file after converting it to PcapNg, or a direct
676 copy of a PcapNg file. This tool is integrated graphically in MUD-PD as described in Section
677 3.2.2.1.

678 **3 Analysis Use Cases and Tools**

679 This section describes several use cases for the characterization methodology along with useful
680 analysis tools.

681 **3.1 Manual MUD File Generation**

682 Currently, MUD files are often generated manually. Although there are tools such as MUD
683 Maker [8] that allow a user to input the necessary values without concern for the computer
684 syntax, most MUD files are still written by hand and require significant effort to complete. After
685 capturing the necessary data through network traffic captures (as described in Section 2), manual
686 analysis is needed to extract the information needed. Relevant information often includes
687 network destinations with which the IoT device has communicated, ports and protocols utilized,
688 and other data regarding the device's behavior. This may be achieved using network traffic-
689 analysis tools like Wireshark and NetworkMiner, which enable extraction of the information
690 necessary for a MUD file.

691 **3.1.1 Wireshark**

692 Wireshark is a well-known open-source tool for network traffic analysis (as well as for packet
693 capture, as discussed in Section 2.1.5.2). It can be run on Windows, OSX/macOS, and Linux. It
694 supports deep packet inspection for hundreds of protocols, which allows the user to sift through
695 packet bytes and extract the relevant information. Analysis can be performed using a wide array
696 of display filters, and results can be exported in a variety of formats. In addition, Wireshark
697 includes decryption support for Secure Sockets Layer (SSL)/Transport Layer Security (TLS) and
698 Wi-Fi Protected Access (WPA)/WPA2. The combination of capabilities allows analysis needed
699 to generate a MUD file from the packet capture file generated as described in Section 2.

700 **3.1.2 NetworkMiner**

701 NetworkMiner is another popular open-source network traffic-analysis tool, and it is built and
702 maintained by Netresec. It is officially supported only on Windows but can be run in macOS
703 through Mono. While it can also be used for packet capture, NetworkMiner's strengths lie in
704 processing network traffic captures and displaying relevant information quickly and easily. It
705 automatically displays network hosts involved and extracts files, images, messages, and
706 credentials. NetworkMiner also compiles a list of individual sessions between hosts and DNS
707 requests throughout the network traffic capture. NetworkMiner does not have the deep packet
708 inspection capabilities that Wireshark has, but it is a quick and helpful tool that complements
709 Wireshark's depth.

710 **3.1.3 Overview of Manual MUD File Generation Process**

711 The process for generating/developing a MUD file begins with a set of network communication
712 capture files. The assumption is that this set includes diverse behaviors such as those described in
713 Section 2. For each network communication capture file, the following steps may be performed:

- 714 1. Inspect packets to locate and record:
 - 715 a. IoT device (source) addresses (media access control [MAC], IPv4/6)

- 716 b. destinations
- 717 i. addresses (MAC, IPv4/6)
- 718 ii. domain names
- 719 c. protocols and ports (Transmission Control Protocol [TCP]/User Datagram
- 720 Protocol [UDP], IPv4/6)
- 721 i. source-initiated (the IoT device being characterized)
- 722 ii. destination-initiated (a device outside the IoT device being characterized)
- 723 2. Identify the destination devices and servers:
- 724 a. type of device
- 725 b. manufacturer

726 Once all of this information has been collected for every packet capture, the final steps are to
727 consolidate it and write the MUD file. The information should be consolidated into a unique list,
728 as some devices and protocols may appear in multiple network communication capture files and
729 each device may have been assigned different IP addresses over time. While IP addresses are not
730 used in MUD files, capturing them can be useful for tracking source and destination pairs. As
731 mentioned above, writing the MUD file may be done manually in a simple text editor or through
732 text entry into MUD Maker [8]. Before any MUD file is deployed, it should be manually
733 verified, and the contents of the MUD file should be confirmed to accurately depict the intended
734 and accepted communication requirements of the IoT device.

735 3.2 MUD-PD

736 The NCCoE developed an open-source tool, MUD-PD, as a proof-of-concept for how to reduce
737 the barrier to entry for vendors to create accurate MUD files for their devices. MUD-PD
738 supplements currently available methodologies for writing MUD files that use packet inspection
739 tools like Wireshark and NetworkMiner. Several approaches to automated MUD file generation
740 currently exist. These include one devised by a researcher at the University of Twente [9], an
741 open-source tool created by the University of New South Wales (UNSW) called MUDgee [10],
742 and an open-source tool called muddy [11], which was created by Lucas Estienne and Daniel
743 Innes at the IETF 105 Hackathon.

744 The MUDgee tool takes a single network traffic capture file and generates a MUD file based on
745 the observed network behavior. MUDgee assumes that all the activity seen is intended and is
746 nonmalicious. While the core of the MUD file generation function in MUD-PD was originally
747 built upon MUDgee, it is now built upon a fork of muddy. The project, lstn/muddy, was
748 developed as a Python and command-line tool to mirror the functionality of MUD Maker. The
749 fork of muddy, usnistgov/muddy [12], leverages the rich code base of lstn/muddy to create a
750 Python object that is more portable and easier to integrate. This fork allows the user to input the
751 desired rules in any order and formats the output according to the format outlined in IETF RFC
752 8520 [1]. Some optional features are not included, but the core data and format are supported.
753 The NCCoE uses this fork of muddy to generate a complete MUD file based on a collection of
754 network traffic captures.

755 The initial version of MUD-PD required that the user manually enter all the metadata as the files
756 are imported. While this functionality is still present, it has been enhanced and the user interface
757 has been simplified. Because MUD-PD now supports the PcapNg file format, JSON-formatted
758 data about the capture environment can be embedded and extracted. This enhancement simplifies
759 the import process and embeds information on the nature of the capture within the packet capture
760 itself to enable metadata to automatically be extracted and imported. The combination of
761 network capture data and documentation allows both for greater portability of the data and for
762 the MUD file-generation process to be more comprehensive and to be automated, requiring little
763 user input.

764 MUD-PD parses and extracts data from packet captures and organizes it in a relational database.
765 The GUI allows the user to examine individual packets or any combination of packets when
766 inspecting the network communications of specific devices. As the metadata about the physical
767 actions and activities that occurred during the network captures are also stored, the user can gain
768 greater insight as to how the network activity and physical world may be associated. In addition
769 to being an exploratory tool intended to aid MUD file development, the database at its core can
770 be queried through any MySQL interface. This allows more potential uses.

771 Additional functions built into MUD-PD include generation of a human-readable device report
772 that summarizes what is discovered on the network and general metadata for each individual
773 network traffic capture. Another significant added function is the automated generation of a
774 MUD file. The MUD file can then be used as is or adjusted and tweaked by the developer or
775 network administrator as they see fit to protect the device and MUD-enabled network. MUD
776 files are generated through a custom user interface to a fork of muddy. This interface leverages
777 an enhanced version of muddy's data pipeline while using the rich preprocessed data stored in
778 the database.

779 3.2.1 Current Feature Set

780 This subsection provides a high-level overview of MUD-PD as it stands. In Section 3.2.2, a tour
781 of the tool illustrates its finer details. MUD-PD has three main functions:

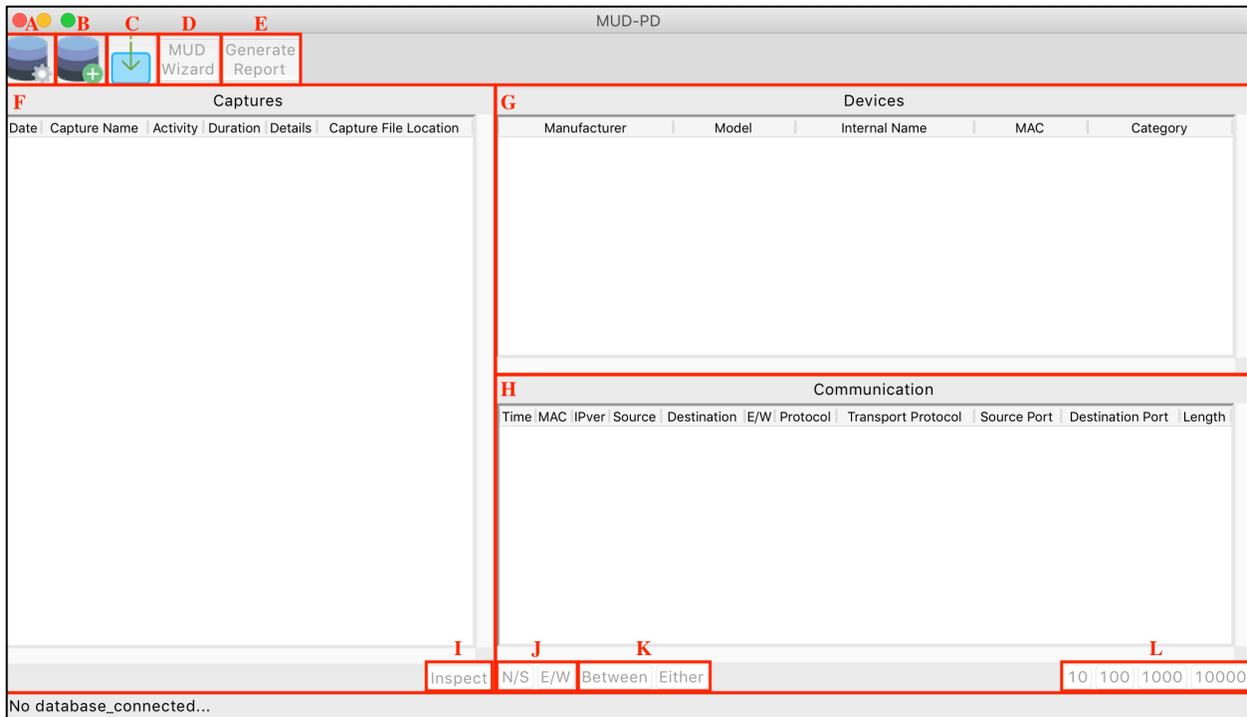
- 782 • **Information import:** The first function is to import network traffic captures. During this
783 step, the user is provided the opportunity to input metadata about the capture. The goal of
784 importing the network traffic capture is to parse the packets—extracting features of
785 interest such as the source, destination, ports, and protocols. This information is at the
786 heart of MUD files. Parsing and importing the network traffic captures permits MUD-PD
787 to extract local network devices and allows them to be labeled as devices of interest.
- 788 • **Database viewing:** The second function is to present a user with a view of information of
789 interest that has been imported into the database. The user can view a list of all the
790 imported packet captures and the devices seen in any and all of the selected network
791 traffic capture files. The user can then select a device or combination of devices to view
792 some information about the packets coming from or to them. For deeper inspection, the
793 user can open the file separately in a packet capture analyzer such as Wireshark or
794 NetworkMiner.
- 795 • **File generation:** The third and most useful function is to generate device reports and
796 MUD files. The device reports summarize the captures in which the device is found,

797 including metadata of the capture environment and a summary of what other devices
798 were communicating on the local network. A wizard walks the user through generating a
799 MUD file from the data in the database and user input. It is up to the user to determine
800 whether the MUD files created are accurate enough to be put in service.

801 **3.2.2 GUI Overview**

802 Upon starting MUD-PD for the first time (installation instructions can be found at
803 <https://github.com/usnistgov/MUD-PD>), the user is greeted with the MUD-PD main window
804 (Figure 1). The labels contained in Figure 1 highlight the components of this window:

- 805 • (A) button to connect to an existing database
- 806 • (B) button to create and (re)initialize a database
- 807 • (C) button to import a capture file
- 808 • (D) button to generate a MUD file
- 809 • (E) button to generate a device report
- 810 • (F) box to show a list of imported capture files
- 811 • (G) box to show a list of active local network devices
- 812 • (H) box to show a list of communications
- 813 • (I) button to inspect a previously imported capture file
- 814 • (J) toggle to limit view of communications to north/south (i.e., external) traffic or
815 east/west (i.e., internal) traffic
- 816 • (K) toggle for a future feature described below
- 817 • (L) buttons to select how many packets to view in the communication box



818

819

Figure 1: MUD-PD main window with buttons and list boxes labeled

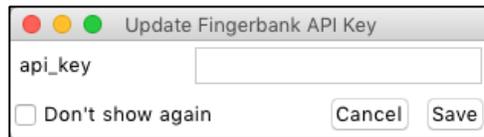
820

When running MUD-PD for the first time, and until dismissed or completed, the user is prompted to provide a locally stored Fingerbank application programming interface (API) key (Figure 2). This enables some useful automation features described in Section 3.2.2.2, but is optional.

821

822

823



824

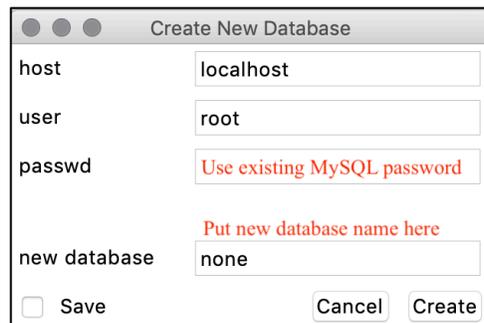
825

Figure 2: Prompt for providing Fingerbank API key

826

The next step is to select the button labeled B to initiate the prompt to create a new database (Figure 3).

827

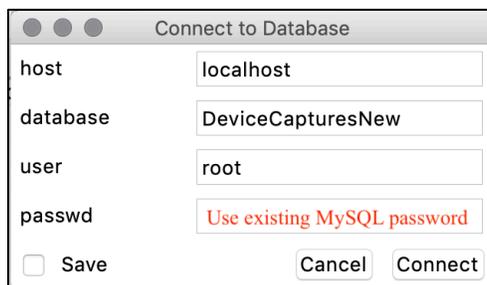


828

829

Figure 3: Prompt for creating a new database

830 Every time MUD-PD is run from this point forward, the user can select the button labeled A to
 831 connect to an existing database (see Figure 1 and Figure 4). When connected to an existing
 832 database, the button for creating a new database may also be used to reinitialize the database,
 833 wiping all existing data. The process is irreversible, so this should be done with caution.



834

835

Figure 4: Prompt for connecting to an existing database

836 After connecting to a database, the user can examine any data contained within it. Referring to
 837 Figure 1, the user can view a list of packet capture files (pcap or PcapNg) that have been
 838 imported thus far in the Captures box (F) on the left side. On the upper right is the section called
 839 Devices (G), which contains a list of local devices communicating in the selected capture files.
 840 The lower right section called Communication (H) contains a list of the packets sent by the
 841 selected devices in the capture files. Above these boxes is a short toolbar with some options.
 842 From left to right, these are: connect to a database (A), create a new database (B), import a
 843 capture file (C), generate a MUD file (D), and generate a device report (E).

844 The Captures list (F) contains metadata for the imported capture files, including the time of
 845 capture, the event captured, the duration of the capture (in seconds), the file location, and any
 846 additional details input during the import process. Below the list is an option to inspect (I) the
 847 currently selected packet capture. If more than one capture is selected, only the capture closest to
 848 the top will be opened. Inspecting a packet capture presents the same window that is opened
 849 when importing a capture file but allows the user to update/modify the details in the database.
 850 The details are identical to the import process, which is covered in detail in Section 3.2.2.1. The
 851 user can select any number of capture files, which will modify the list of devices to show any/all
 852 local devices that have sent or received packets during the captures.

853 The Devices list (G) includes information that either can be inferred from capture information or
 854 that has been input by the user during the import process. This includes the manufacturer, the
 855 model, a unique name for internal/lab use, the MAC address, and the general category of the
 856 device. The selection of an entry in the Devices list will determine what is listed in the
 857 Communication box. The user can either select “All...” to view all the packets communicated
 858 across the network, or a single device to view only the communication to/from that device. The
 859 user may select multiple devices to view the communication to/from any of the selected devices.

860 The Communication list (H) displays parsed packet information such as the time, MAC address
 861 of the sender, IP version, source and destination addresses, scope of traffic, innermost protocol
 862 layer, transport protocol, source and destination ports, and packet length. The IP version is given
 863 as either 4 or 6. If it is blank, the packet is below the IP layer (i.e., layer 3). By scope of traffic,
 864 we mean whether it would be considered east/west (i.e., internal/local network) traffic indicated

865 by a value of 1, or north/south (i.e., to/from an external address/network) indicated by a value of
 866 0. The source and destination ports are those of TCP or UDP. The user can choose to filter by
 867 north/south (N/S) or east/west (E/W) traffic and can select the number of packets displayed (J).
 868 When two devices are selected, the two additional buttons (K) allow the user to view traffic
 869 either *between* the two devices or involving *either* device. Last, the user may select to view the
 870 first 10, 100, 1000, or 10,000 packets that satisfy the above filters (L).

871 3.2.2.1 Importing a New Packet Capture

872 The potential of this tool begins to be realized when importing a packet capture file. Here, the
 873 user is prompted to select the file to import (Figure 5). If the file is a PcapNg file, then MUD-PD
 874 will automatically search for embedded metadata, otherwise the user can input metadata
 875 regarding the capture. This includes the phase of the device life cycle being captured. In most
 876 cases, this will be normal operation. The other two options are setup and removal, as described in
 877 Section 2.1.1. The user can also select all the environmental variables that apply, including
 878 whether internet connectivity was enabled, the device's preferred DNS was blocked, the device
 879 was isolated on the network, there were notable physical environmental changes, the capture was
 880 of a full network of devices, a controller or hub was involved, a device of the same manufacturer
 881 as the primary device of interest was connected, and/or there was human interaction with the
 882 device. Whether the capture was duration-based or action-based should also be selected. The
 883 specific duration (in seconds) or action can be input, which is highly recommended for
 884 auditability and ease of use.

The screenshot shows a dialog box titled "Import Packet Capture". It features a "File" input field with a browse button, a "Notes (optional)" text area, and three sections of options:

- Lifecycle Phase:** Three radio buttons: "Setup", "Normal Operation" (selected), and "Removal".
- Environmental Variables:** A grid of checkboxes: "Internet" (checked), "Preferred DNS Enabled" (checked), "Isolated", "Notable Physical Changes", "Full Network", "Controller/Hub", "Same Manufacturer", and "Human Interaction".
- Capture Type:** Two radio buttons: "Duration-based" (selected) and "Action-based". Below "Duration-based" is a "Duration" text field. Below "Action-based" is an "Action" text field.

At the bottom right, there are "Cancel" and "Import" buttons.

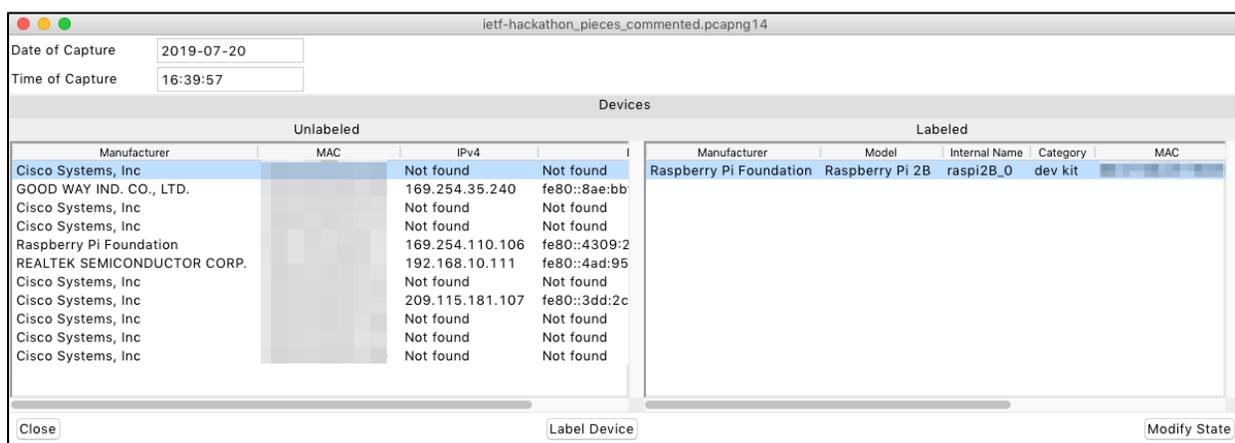
885

886

Figure 5: Prompt for importing packet captures into database

887 **3.2.2.2 Viewing and Importing Devices**

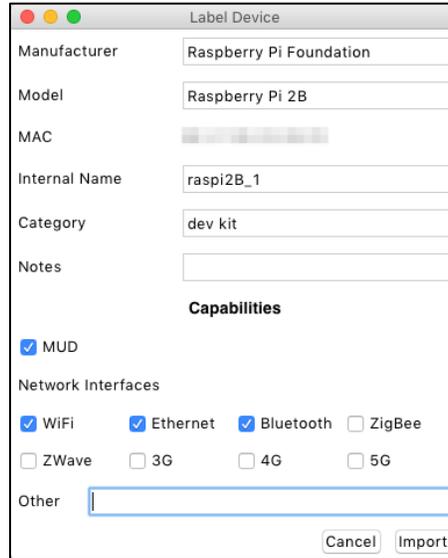
888 During the packet capture import process, the user is presented with lists of the labeled and
 889 unlabeled devices that were seen in the capture file (Figure 6). A *labeled device* is one that has
 890 been seen in a previously imported capture and has had related data imported to the database. An
 891 *unlabeled device* may have been seen in a previous capture but has not yet had any additional
 892 data imported. This packet capture import window also includes the time and date of the capture,
 893 which is extracted from the capture file, but can be edited if the user believes either or both are
 894 incorrect for some reason. The left list is the unlabeled devices. MUD-PD attempts to look up the
 895 manufacturer based on the Organizationally Unique Identifier (OUI), which is the first 24 bits of
 896 the MAC address, and also lists the IP addresses (both v4 and v6 when available). The user can
 897 select any device in this list and import additional details into the database, moving it to the list
 898 of labeled devices on the right. In addition to the information found in the unlabeled list, this one
 899 includes all the information available in the device list of the main window (Figure 1). The state
 900 of the device (i.e., the firmware version) can also be updated here. This field is not used in any
 901 automated processes of MUD-PD but can be queried through MySQL.



902

903 **Figure 6: Window listing devices imported and to import during the packet capture import process**

904 Selecting the Import Device button presents the user with a window with fields for adding or
 905 modifying the manufacturer, model, internal name, category, notes, and list of capabilities
 906 (Figure 7). The manufacturer and model are required fields. In addition to being required, the
 907 internal name must be unique. The device category and notes are optional fields but may be
 908 useful for documentation and future analyses. The capabilities consist of MUD, Wi-Fi, Ethernet,
 909 Bluetooth, ZigBee, ZWave, 3G, 4G, 5G, and other. Other than MUD, all the capabilities are
 910 network interfaces, of which at least one must be selected. The MAC address of the device is
 911 also listed but may not be modified, as this is determined from the capture itself and is used as an
 912 identifier. In addition, integration with the Fingerbank API is included, assisting the user by
 913 identifying the device model based on the Dynamic Host Configuration Protocol (DHCP)
 914 fingerprint and MAC address. To enable this feature, the user must obtain and enter a valid
 915 Fingerbank API key.



916

917

Figure 7: Window prompt for importing a device

918 After the metadata has been input and the Import button has been selected, the user is prompted
 919 to input the firmware version of the device (Figure 8).



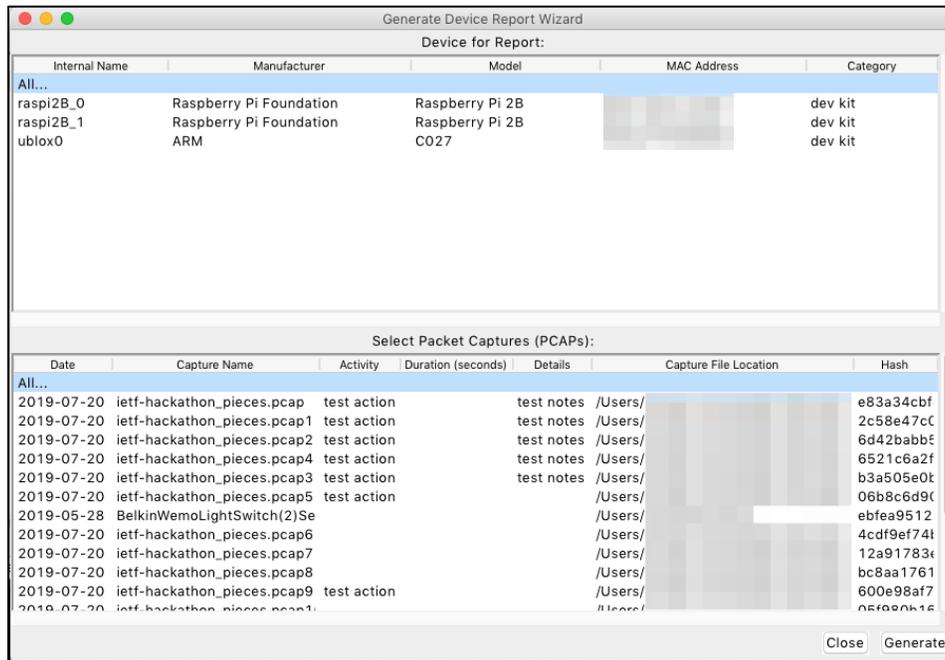
920

921

Figure 8: Window prompting to update the firmware version logged in the database

922 **3.2.2.3 Generating Device Reports**

923 The process for generating a device report is straightforward (Figure 9). The user may generate
 924 reports for any combination of devices or a single device. After selecting the device(s) for which
 925 to generate the report, the list of packet captures is updated to only those in which the device has
 926 sent or received packets. The user may select all or any combination of packets to report on.



927

928

Figure 9: Prompt for generating a human-readable device report

929 The generated report lists the packet captures in which the device is seen, including the hash of
 930 the file. The example report, shown in Figure 10, contains only one file, whereas a typical report
 931 may contain many. The capture metadata is also listed for each file. In addition, listed under each
 932 capture file are the other local devices seen on the network during the capture. The internal name
 933 (if the device is labeled) is also given. Eventually, this report may include more specific
 934 information about the communication to/from the device, similar to what would be listed in the
 935 device’s MUD file (if it had one).

```

This document serves to indicate the devices and operations captured in addition
to any specific procedures or environmental details of the captures.

Device: raspi0
MAC:    b8:27:eb:01:23:45

Capture File:  example_file.pcap
SHA256 Hash:  e83a34cbf4eab7bd8726bb9f4fce1db89b3928625c27a300d3c557ea7056466f
Device Phase:  Normal Operation
Environmental Variables:
  Internet enabled      True
  Human Interaction     False
  Preferred DNS Enabled True
  Device Isolated      False
Action-based Capture:  False
Duration-based Capture: True
  Intended Duration:   600
  Actual Duration:     754
Start Time:    2020-02-02 12:34:56
End Time:     2020-02-02 12:47:30
Other Devices:
  Name:  router0
  MAC:  01:23:45:67:89:ab
  Name:  controller0
  MAC:  fe:dc:ba:98:76:54
  Name:  rasp11
  MAC:  d8:27:eb:67:89:ab
Notes:
  Example capture with made-up devices
    
```

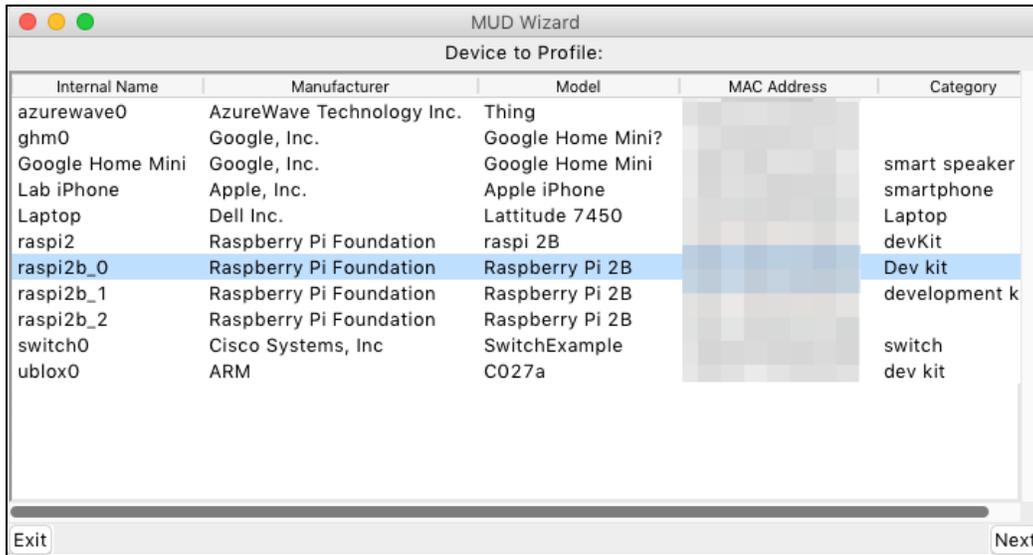
936

937

Figure 10: Example device report showing the details of a single packet capture

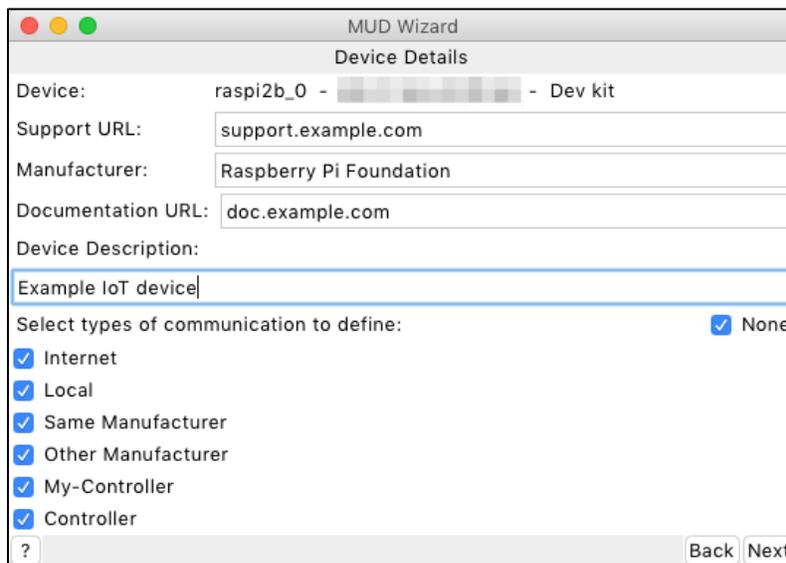
938 **3.2.2.4 Generating a MUD File**

939 Generating the MUD file requires more user input and involves more steps than generating the
 940 device report. The user is first prompted to select a device for which to generate a MUD file
 941 (Figure 11).



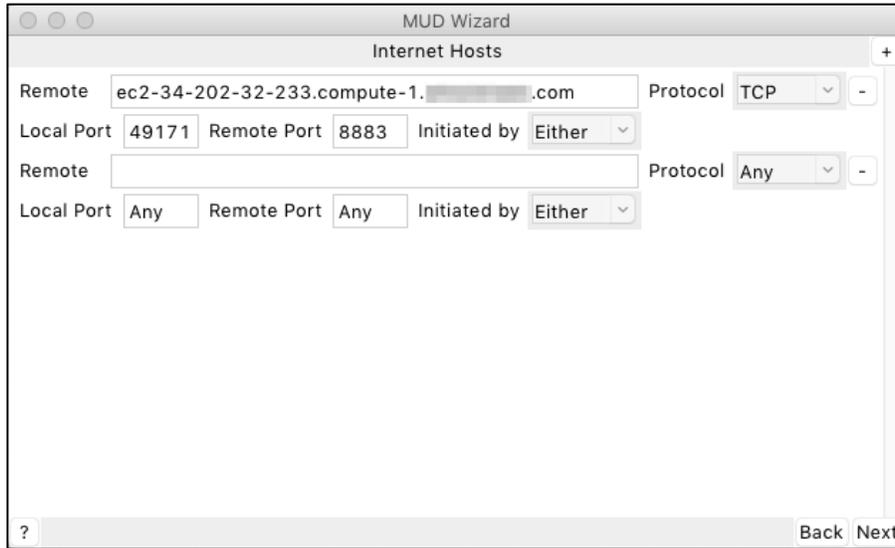
942
 943 **Figure 11: Prompt for selecting a device for which the MUD file will be generated**

944 The next step is to fill in the device details (Figure 12). Here the support URL, documentation
 945 URL, and device description can be provided. Additionally, the user may select which types of
 946 communication to define in the MUD file: internet, local, same manufacturer, other named
 947 manufacturers, network-defined controller (my-controller), and controller. Internet and local
 948 communication may automatically be selected if traffic involving the device has been identified
 949 going N/S or E/W, respectively.



950
 951 **Figure 12: Prompt for providing device details including the support and document URLs**

952 Proceeding to the next page provides the user with the ability to define a list of rules for the
 953 given communication type (Figure 13). The layout of the window is the same for each of the six
 954 communication types. For internet and local communication, the window is populated with a list
 955 of hosts that were observed to have communicated with the device in any of the packet capture
 956 files stored in the database. DNS resolution is attempted for all internet hosts.

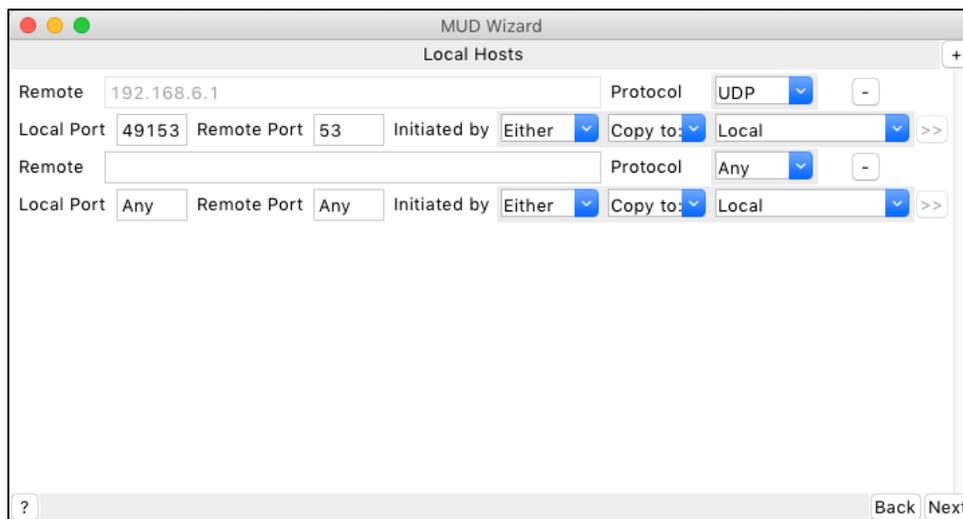


957

958

Figure 13: Prompt for providing internet communication rules

959 The window for local communication rules (Figure 14) differs slightly from the internet
 960 communication rules window. As the local traffic observed may have been from a device from
 961 one of the other communication types selectable in the “Device Details” window, these local
 962 rules can be copied or moved to any of the other non-internet communication types that were
 963 selected to be defined. Each of the windows for the remaining communication types also allow
 964 rules to be copied or moved.



965

966

Figure 14: Prompt for providing local communication rules

967 For each host, the communication protocol (i.e., TCP or UDP) is automatically selected based on
 968 the observed communication. The local and remote ports are also automatically filled based on
 969 the observed communication. If the ports are left blank or listed as “any”, any port will be
 970 allowed. The initiation direction can be manually specified by selecting “thing” or “remote” to
 971 indicate whether the IoT device or the host, respectively, must be the party to start the
 972 communication. By default, “either” is selected, indicating that either side may initiate the
 973 communication.

974 The MUD specification defines several conditions that apply to the protocol, ports, and initiation
 975 direction. If “any” protocol is allowed (i.e., both TCP and UDP), the ports and initiation direction
 976 are ignored. If TCP is selected, the ports and initiation direction can be specified. If UDP is
 977 selected, the initiation direction is ignored, and communication can be initiated by either side.

978 There is a difference in how a few of the communication types process the host fields. These
 979 include local, same manufacturer, and network-defined controllers. For these types, the host
 980 cannot be specified or is ignored. Local communication rules allow any local device to follow
 981 the indicated rules, same manufacturer uses the manufacturer hostname specified on the second
 982 page, and network-defined controllers are defined by the local network administrator.

983 Once all the desired communication type rules have been defined, the user is provided with a
 984 preview of the resulting MUD file (Figure 15). There is an option for advanced users to manually
 985 modify the outputted MUD file at their own risk. There is no guarantee that a modified output
 986 file will be formatted or defined correctly.

```

{
  "ietf-mud:mud": {
    "mud-version": 1,
    "mud-url": "support.example.com",
    "is-supported": true,
    "mfg-name": "Raspberry Pi Foundation",
    "documentation": "doc.example.com",
    "systeminfo": "raspi2b_0",
    "cache-validity": 48,
    "last-update": "2021-05-03T16:14:56",
    "from-device-policy": {
      "access-lists": {
        "access-list": [
          {
            "name": "mud-70313-v4fr"
          }
        ]
      }
    }
  }
}

```

987 Advanced mode Back Save

988 **Figure 15: Preview of the MUD file to be generated**

989 3.2.3 MUD-PD Uses

990 MUD-PD is primarily intended to be a tool in support of MUD. It may be one component of a
 991 greater pipeline that enables MUD research and deployment. There are several possible pipelines
 992 that depend on specific use-cases, each of which are described in greater detail in Section 4.2.

993 While there are MUD-specific features to MUD-PD, it is relatively purpose-agnostic. While its
994 primary intention is to assist in generating MUD files for IoT devices, the data it stores can be
995 analyzed in other ways for other purposes. Because the data set will inevitably get large and is
996 labeled, machine learning techniques could be applied in an effective manner. The applications
997 of machine learning and this data set are plentiful, including those not foreseen.

998 As the next section discusses, the same data collected for generating MUD files can be used to
999 examine the privacy implications of these devices. Investigation into what the devices are
1000 communicating (the content of the communication) rather than simply how they are
1001 communicating can lead to a deeper understanding and greater awareness of the implications of
1002 putting smart devices in our homes.

1003 **3.3 MUD-PD Support for Privacy Analysis**

1004 As mentioned above, MUD-PD is a tool that can be applied for more purposes than generating
1005 MUD files for IoT devices. While MUD files define the suggested behavior of a device, and one
1006 could argue that the content communicated is a component of a device's behavior, they do not
1007 necessarily capture the privacy implications associated with the device or its associated
1008 networks. In the case where the intended use of MUD-PD is to investigate privacy, the NCCoE
1009 recommends this tool be used only in a research and development setting, as there are no security
1010 guarantees for the stored data. Use in an uncontrolled setting may result in a violation of
1011 confidentiality. To understand the privacy implications in such a setting requires understanding
1012 the data content being transmitted from the device to outside services. This may be challenging,
1013 depending on the device and the protocols implemented, as the content in the network packets
1014 may or may not be encrypted. Even where they are encrypted, the protocol under analysis may
1015 be susceptible to a man-in-the-middle attack that reveals some or all of the contents of the
1016 packets. Utilizing such an attack may be useful for an investigation into privacy, but again,
1017 should be implemented with caution and only in a controlled research and development setting.

1018 There may be some moral, ethical, and privacy implications in implementing such an evaluation
1019 technique. These should be mitigated by limiting use of the tool to a controlled environment (i.e.,
1020 a laboratory) and by adhering to the NIST Privacy Framework [\[13\]](#) and the Common Rule for
1021 the Protection of Human Subjects [\[14\]](#). The same techniques for collection and logging can be
1022 beneficial to privacy investigations—tracking what potentially private information is transmitted
1023 and tracing the risks to all the devices and parties involved.

4 Future Work

1025 The NCCoE has decided to conclude feature development for MUD-PD unless there becomes a
1026 substantial demand for additional features. This does not mean there is no more room for work or
1027 development in this area. The NCCoE encourages continued community-driven research of
1028 device characterization and development of MUD-PD.

1029 A few open problems include ensuring that any methodology prescribed for characterizing
1030 devices is robust from security and reliability points of view. Additional analysis use cases and
1031 tools, including alternative device characterization approaches (e.g., fingerprinting), could also
1032 be demonstrated and documented to help expand and confirm the efficacy of the methodology.
1033 Additional situations and environmental variables that could modify the behavior of an IoT
1034 device need to be identified and documented. Support for storing capture environment variables
1035 within a PcapNp file as an official option would also benefit the community of packet capture
1036 analysts.

4.1 Extending MUD-PD Features

1038 Existing plans for development of MUD-PD have been concluded. That said, the NCCoE
1039 encourages any interested members of the community to consider continuing the development of
1040 enhancements and additional features. As MUD-PD was to serve primarily as a proof-of-
1041 concept, there is room for improvement of existing features including streamlining, simplifying,
1042 and speeding the collection, logging, and file generation processes. The usefulness of generated
1043 device reports could also be improved from community and user feedback. These reports could
1044 be expanded to list the ports used, as well as the specific hosts and servers with which the device
1045 has communicated.

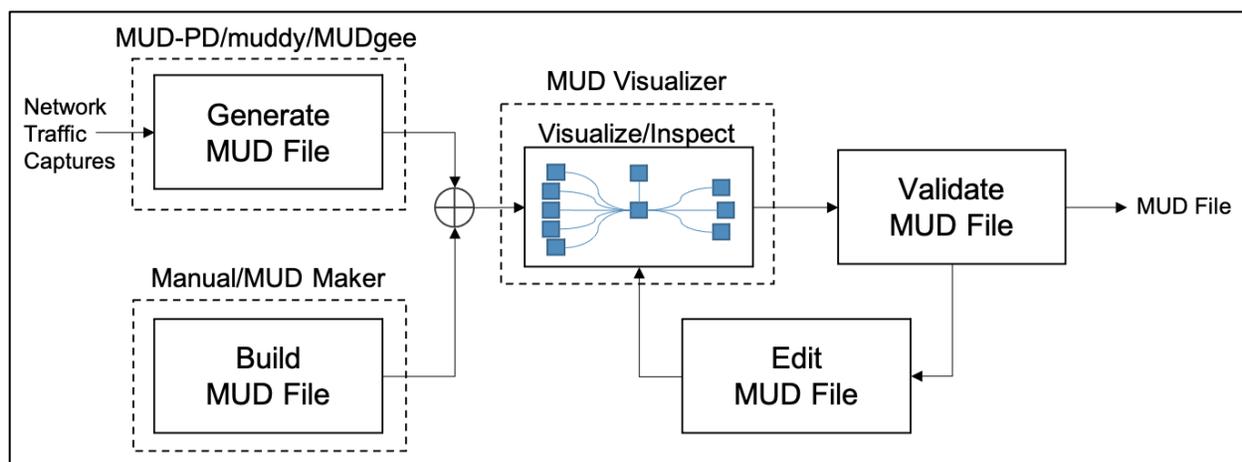
1046 A number of enhancements to the usability and user experience of the MUD-file generation
1047 process itself should also be considered. This includes presenting the user with coarse estimates
1048 or warnings of the potential quality of the produced MUD file that can be expected based upon
1049 the network traffic captures selected, the goal being to highlight where gaps and deficiencies
1050 may exist in the resulting MUD file.

1051 MUD-PD could be extended to extract and catalog additional data from capture files to
1052 investigate the privacy implications of IoT devices. To do so will require that packet payload
1053 information be extracted and stored. This includes strings, images, credentials seen, and
1054 certificates. It may also be worth logging whether packets are encrypted as well as the type and
1055 strength of the algorithm.

4.2 Developing a MUD Pipeline

1057 The NCCoE is proposing the creation of a set of pipelines focused on MUD file development,
1058 which address different use cases for MUD. Three use cases have been considered: (1) a device
1059 manufacturer or developer that needs to provide a MUD file for its users; (2) a network
1060 administrator who may wish to inspect an official MUD file or a device's adherence to said file
1061 and who may wish to augment or modify its allowed behavior; and (3) a researcher who may be
1062 interested in all of the above in addition to investigating the intricacies of existing MUD rules
1063 and proposed extensions.

1064 In the first use case, a device manufacturer or developer might find it useful to have access to a
 1065 suite of interoperable tools that make the generation, inspection, and validation of MUD files
 1066 easy and straightforward (Figure 16). To begin the process, the two options are to build a MUD
 1067 file by hand by using a tool like MUD Maker [8] or to generate one from a capture of network
 1068 traffic by using MUD-PD or MUDgee. The next steps are to inspect the MUD file, which can be
 1069 done visually using the MUD Visualizer [15], and validate that no rules are missing that should
 1070 be present and no rules are present that should not be—and to edit rules where necessary. After a
 1071 number of iterations through these steps, manufacturers may reach a point where they are
 1072 confident in the MUD files and publish them for user consumption. The process depicted in
 1073 Figure 16 can also be used to generate MUD files for devices without a manufacturer-provided
 1074 MUD file.

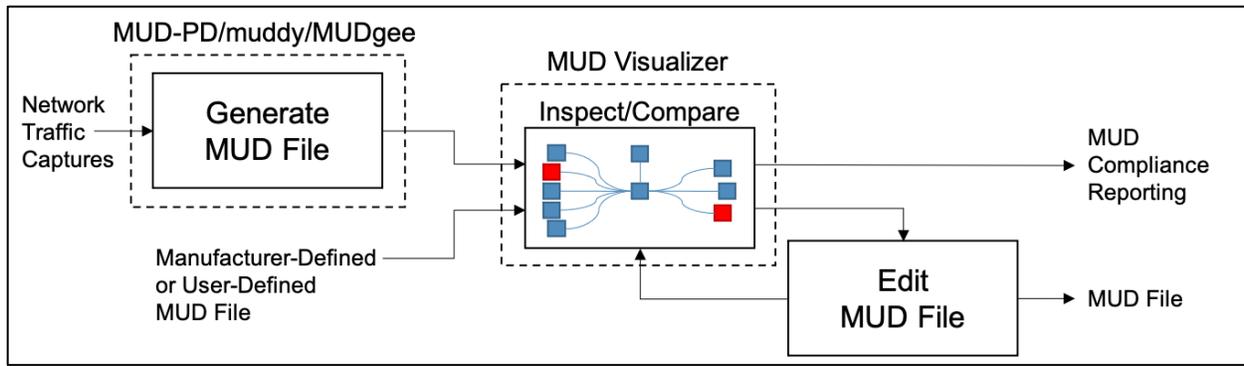


1075

1076

Figure 16: MUD pipeline for the device manufacturer or developer use case

1077 In the second use case, it may be useful for network administrators to have a view of the network
 1078 with an overlay of the MUD rules that have been defined by a manufacturer (Figure 17). To
 1079 drive this capability, they must be able to ingest a MUD file and compare it against the behavior
 1080 observed on the network. The MUD file may be manufacturer-defined or user-defined. When the
 1081 MUD file and observed behavior are inspected and compared, the network administrator could
 1082 be presented with a diagram highlighting where the observed behavior does not comply with the
 1083 MUD file. The UNSW researchers have developed a tool for comparing a provided MUD file
 1084 with observed activity [16]. One also could imagine the MUD Visualizer tool being extended to
 1085 include this capability. Because the network administrator may also be interested in reducing or
 1086 expanding a device's capabilities, tailoring it to their specific network, the ability to build and/or
 1087 edit MUD files would be desirable. MUD files can currently be built/written using MUD Maker,
 1088 but there is not a dedicated tool for editing MUD files. To assist in live network administration
 1089 and monitoring, it may be useful for the comparisons to be done on the fly on a live network,
 1090 issuing live reports or warnings when noncompliance is detected.

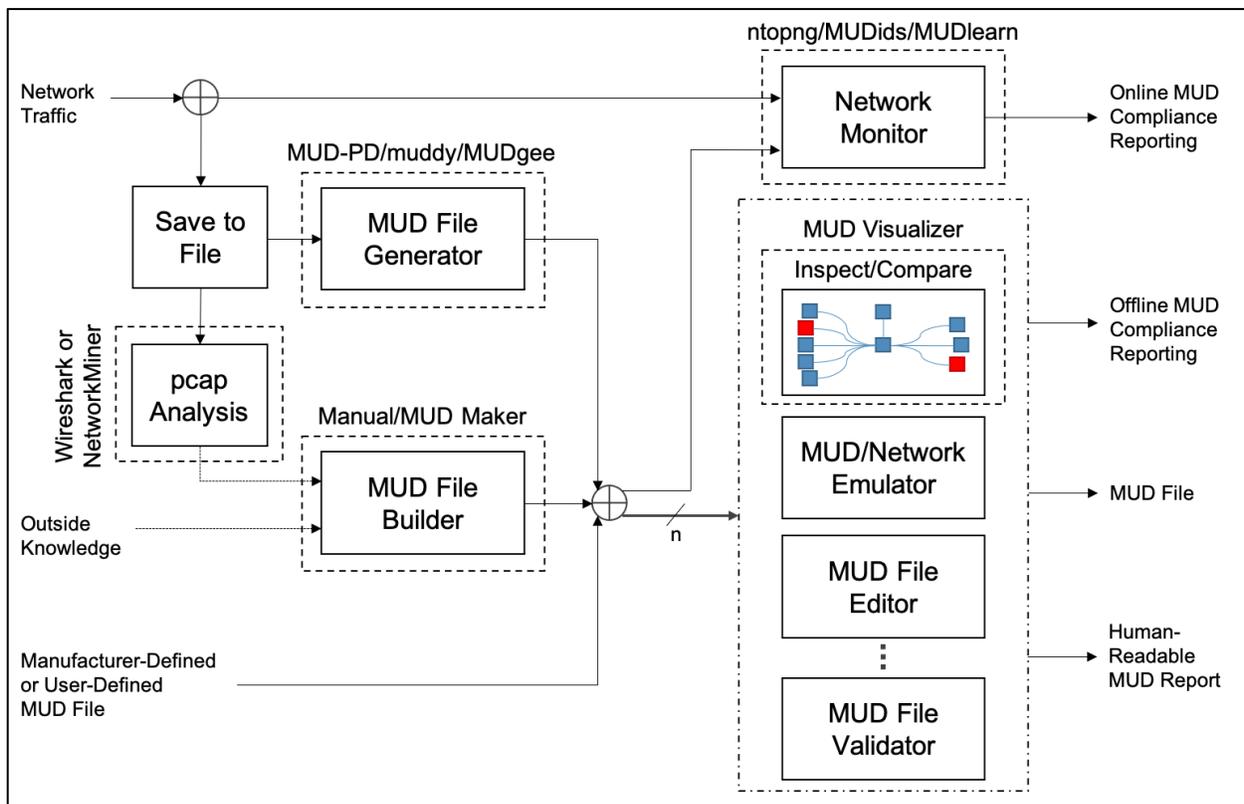


1091

1092

Figure 17: MUD pipeline for the network administrator use case

1093 The third use case is more open-ended. Researchers may also want access to all the same tools
 1094 useful to manufacturers and network administrators, and even more. There could be interest in
 1095 studying existing MUD files or investigating the implications of various MUD rules or offering
 1096 extensions (see Figure 18). For researchers, it may be useful to emulate a network of devices
 1097 based on the MUD files to understand how networks scale and devices interact.



1098

Figure 18: The overarching MUD pipeline, particularly as it may be used for research and development

1100 Figure 18 demonstrates how many existing and proposed future tools relevant to MUD can be
 1101 leveraged to achieve the research and development goals of the use cases described above.
 1102 Several boxes in Figure 18 are labeled with existing tools that could potentially fill the associated

1103 roles in their current state or with future development. The boxes that lack a dashed outline have
1104 not been associated with any existing tools that could potentially fill the role.

1105 There are a number of ways in which a MUD file may be generated or selected. MUD files may
1106 come from the manufacturer, be generated by the user using network captures through MUD-PD
1107 or MUDgee, or be written by hand with assistance from MUD Maker and Wireshark and/or
1108 NetworkMiner. These MUD files can then be used for several purposes or processed in a number
1109 of ways. Some may require using one version while others may require two or more, as indicated
1110 by the n in Figure 18.

1111 A MUD plug-in is in development for the ntopng network monitoring tool [17]. When using a
1112 MUD file with live analysis of network activity, there is the potential for real-time MUD
1113 compliance reporting. Additionally, extensions to MUD's functionality are being proposed for
1114 use within the tool. Interest has been expressed in developing other MUD reporting tools. For
1115 example, the UNSW researchers have been using MUD in combination with software-defined
1116 networking to develop an intrusion detection system as well as a tool for detecting volumetric
1117 attacks, both of which have the potential for live reporting. These are called MUDids and
1118 MUDlearn, respectively [18], [19]. MUD files can also be visualized using the MUD Visualizer
1119 tool that is paired with MUD Maker. This tool could potentially be extended to compare two
1120 MUD files for offline compliance and manual validation. Additionally, tools are being proposed
1121 for automated validation of MUD files and network emulation based on these files. Development
1122 of APIs for these tools would greatly enhance interoperability and future development. The
1123 NCCoE hopes that the community of IoT manufacturers, developers, network administrators,
1124 and researchers will continue to contribute to improvements in this area.

1125 **4.3 Open Problems for the Community**

1126 The NCCoE encourages members of this community of interest to consider addressing the
1127 following open problems and questions:

- 1128 • Because it may be impossible to capture all potential aspects of an IoT device's behavior,
1129 how can the accuracy of a MUD file be measured?
 - 1130 ○ What other situations and environmental variables could modify the behavior of a
1131 device?
 - 1132 ○ How can the correctness of a MUD file be verified (and ensure that unnecessary
1133 behavior is not included)?
 - 1134 ○ What combination of captures is needed to create a comprehensive MUD file (and
1135 ensure behavior that should be permissible is not omitted)?
- 1136 • What are other applications of a MUD-PD tool or its data sets?
- 1137 • What other tools should be considered for connecting in the MUD pipeline (or other
1138 pipelines)?
- 1139 • What features are desirable for a tool like this?
- 1140 • What other extractable features of packet captures might be of use to developers, network
1141 administrators, and researchers?

- 1142 • How can the quality and efficiency of the tool be improved?
- 1143 • How can a prescribed methodology for characterizing devices be ensured to be robust in
1144 its security and reliability?
 - 1145 ○ How can its efficacy be objectively demonstrated? How do alternative device
1146 characterization approaches (e.g., fingerprinting) compare?
- 1147 • Are there widespread use-cases for including capture environment variables within a
1148 PcapNg file such that it should be included as an official option in the specification?
 - 1149 ○ What environmental variables should be included in such an option?

1150 **References**

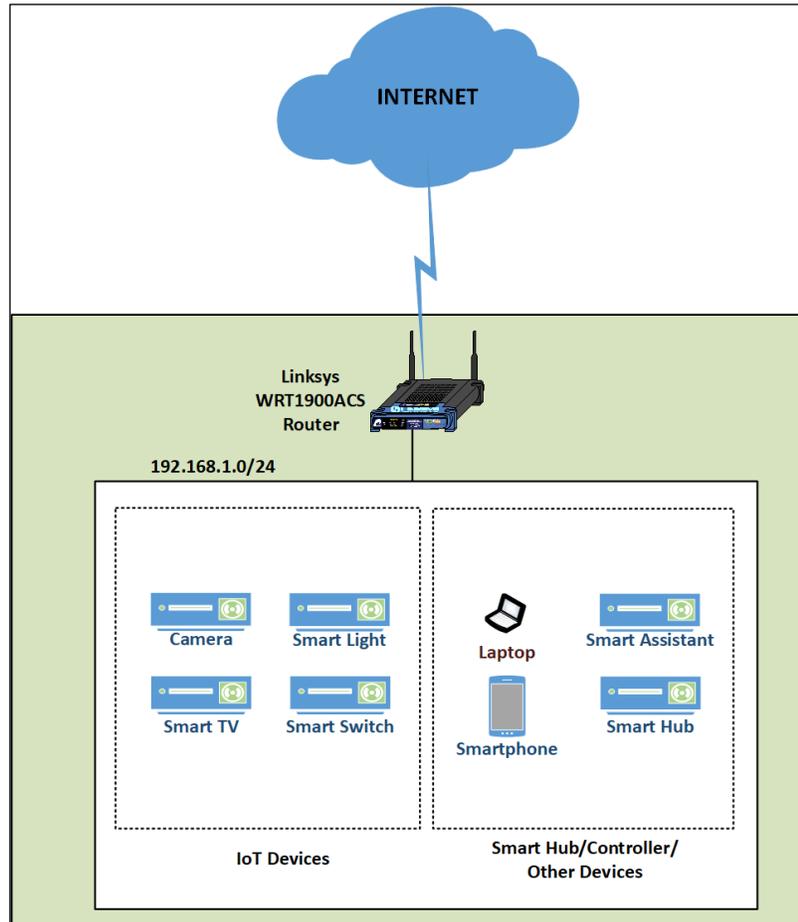
- [1] Lear E, Droms R, Romascanu D (2019) Manufacturer Usage Description Specification. (Internet Engineering Task Force (IETF)), IETF Request for Comments (RFC) 8520. <https://doi.org/10.17487/RFC8520>
- [2] Thangavelu V, Divakaran DM, Sairam R, Bhunia SS, Gurusamy M (2019) DEFT: A Distributed IoT Fingerprinting Technique. *IEEE Internet of Things Journal* 6(1):940-952. <https://doi.org/10.1109/JIOT.2018.2865604>
- [3] Huang DY, Apthorpe N, Acar G, Li F, Feamster N (2019) IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *arXiv preprint*. <https://arxiv.org/abs/1909.09848>
- [4] Meidan Y, Bohadana M, Shabtai A, Guarnizo JD, Ochoa M, Tippenhauer NO, Elovici Y (2017) ProfilloT: A Machine Learning Approach for IoT Device Identification Based on Network Traffic Analysis. *Proceedings of the Symposium on Applied Computing (SAC '17)* (ACM, Marrakech, Morocco), pp 506-509. <https://doi.org/10.1145/3019612.3019878>
- [5] Bezawada B, Bachani M, Peterson J, Shirazi H, Ray I, Ray I (2018) Behavioral Fingerprinting of IoT Devices. *Proceedings of the 2018 Workshop on Attacks and Solutions in Hardware Security (ASHES '18)* (ACM, Toronto, Canada), pp 41-50. <https://doi.org/10.1145/3266444.3266452>
- [6] Aneja S, Aneja N, Islam MS (2018) IoT Device Fingerprint using Deep Learning. *2018 IEEE International Conference on Internet of Things and Intelligence System (IOTAIS)* (IEEE, Bali, Indonesia), pp 174-179. <https://doi.org/10.1109/IOTAIS.2018.8600824>
- [7] Dodson D, Montgomery D, Polk T, Ranganathan M, Souppaya M, Johnson S, Kadam A, Pratt C, Thakore D, Walker M, Lear E, Weis B, Barker W, Coclin D, Hojjati A, Wilson C, Jones T, Baykal A, Cohen D, Yelch K, Fashina Y, Grayeli P, Harrington J, Klosterman J, Mulugeta B, Symington S, Singh J (2021) Special Publication 1800-15: Securing Small-Business and Home Internet of Things (IoT) Devices: Mitigating Network-Based Attacks Using Manufacturer Usage Description (MUD) (National Institute of Standards and Technology, Gaithersburg, MD). <https://doi.org/10.6028/NIST.SP.1800-15>
- [8] Lear E (2020) *MUD Maker Tool*. Available at <https://mudmaker.org/mudmaker.html>
- [9] Schutijser CJTM (2018) Towards Automated DDoS Abuse Protection Using MUD Device Profiles. (University of Twente, Enschede, The Netherlands). Available at <http://purl.utwente.nl/essays/76207>
- [10] Hamza A, Ranathunga D, Gharakheili HH, Roughan M, Sivaraman V (2018) Clear as MUD: Generating, Validating and Applying IoT Behavioral Profiles. *Proceedings*

- of the 2018 Workshop on IoT Security and Privacy (IoT S&P '18)* (ACM, Budapest, Hungary), pp 8-14. <https://doi.org/10.1145/3229565.3229566>
- [11] Estienne L, Innes D (2019) *muddy*. Available at <https://github.com/lstn/muddy>
- [12] Watrobski P, Klosterman J (2021) *muddy*. Forked from *muddy* (<https://github.com/lstn/muddy>), Estienne L and Innes D. (National Institute of Standards and Technology, Gaithersburg, MD). Available at <https://github.com/usnistgov/muddy>
- [13] National Institute of Standards and Technology (2020) NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, Version 1.0. (National Institute of Standards and Technology, Gaithersburg, MD). Available at <https://www.nist.gov/privacy-framework>
- [14] U.S. Department of Health and Human Services, Office for Human Research Protections (2020) *Federal Policy for the Protection of Human Subjects ('Common Rule')*. Available at <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html>
- [15] Andalibi V, Lear E (2020) *MUD Visualizer Tool*. Available at <https://www.mudmaker.org/mudvisualizer.php>
- [16] Hamza A, Ranathunga D, Gharakheili HH, Benson TA, Roughan M, Sivaraman V (2019) Verifying and Monitoring IoTs Network Behavior using MUD Profiles. *arXiv preprint*. <https://arxiv.org/abs/1902.02484>
- [17] ntop (2020) *ntopng: High-Speed Web-based Traffic Analysis and Flow Collection*. Available at <https://www.ntop.org/products/traffic-analysis/ntop/>
- [18] Hamza A, Gharakheili HH, Sivaraman V (2018) Combining MUD Policies with SDN for IoT Intrusion Detection. *Proceedings of the 2018 Workshop on IoT Security and Privacy (IoT S&P '18)* (ACM, Budapest, Hungary), pp 1-7. <https://doi.org/10.1145/3229565.3229571>
- [19] Hamza A, Gharakheili HH, Benson TA, Sivaraman V (2019) Detecting Volumetric Attacks on IoT Devices via SDN-Based Monitoring of MUD Activity. *Proceedings of the 2019 ACM Symposium on SDN Research (SOSR '19)* (ACM, San Jose, California), pp 36-48. <https://doi.org/10.1145/3314148.3314352>

1151 **Appendix A—Example Capture Environment**

1152 This appendix presents an example capture environment that supports analysis of both wired and
1153 wireless IoT devices. Example procedures for capture are identified in Section 2.2. The
1154 following components compose the example environment and are depicted in Figure 19:

- 1155 • home router with tcpdump capability for capturing all network traffic, both wired and
1156 wireless (Linksys WRT1900ACS running OpenWRT)
- 1157 • external storage (such as a flash drive) to increase capture storage capacity of the home
1158 router
- 1159 • computer running Linux or macOS (can be used for both capture and analysis as needed)
- 1160 • IoT devices to characterize (camera, smart light, smart TV, smart switch)
- 1161 • other devices that interact/communicate with the IoT devices (such as smart
1162 hubs/controllers/smartphones)



1163

1164

Figure 19: Example capture architecture

1165 Appendix B—Acronyms

1166 Selected acronyms and abbreviations used in this paper are defined below.

1167	API	Application Programming Interface
1168	BSD	Berkeley Software Distribution
1169	DHCP	Dynamic Host Configuration Protocol
1170	DNS	Domain Name System
1171	E/W	East/West
1172	FOIA	Freedom of Information Act
1173	GUI	Graphical User Interface
1174	IETF	International Engineering Task Force
1175	IoT	Internet of Things
1176	IP	Internet Protocol
1177	IT	Information Technology
1178	ITL	Information Technology Laboratory
1179	JSON	JavaScript Object Notation
1180	MAC	Media Access Control
1181	MUD	Manufacturer Usage Description
1182	NCCoE	National Cybersecurity Center of Excellence
1183	NIST	National Institute of Standards and Technology
1184	N/S	North/South
1185	OUI	Organizationally Unique Identifier
1186	pcap	Packet Capture
1187	PcapNg	Packet Capture Next Generation Dump
1188	RFC	Request for Comments
1189	SDN	Software-Defined Networking
1190	SPAN	Switched Port Analyzer
1191	SSL	Secure Sockets Layer
1192	TCP	Transmission Control Protocol
1193	TLS	Transport Layer Security
1194	UDP	User Datagram Protocol
1195	UNSW	University of New South Wales
1196	URL	Uniform Resource Locator
1197	WPA	Wi-Fi Protected Access