

# Withdrawn Draft

## Warning Notice

The attached draft document has been withdrawn, and is provided solely for historical purposes. It has been superseded by the document identified below.

**Withdrawal Date** June 5, 2019

**Original Release Date** September 21, 2018

## Superseding Document

**Status** Final

**Series/Number** NISTIR 8221

**Title** A Methodology for Enabling Forensic Analysis Using Hypervisor Vulnerabilities Data

**Publication Date** June 2019

**DOI** <https://doi.org/10.6028/NIST.IR.8221>

**CSRC URL** <https://csrc.nist.gov/publications/detail/nistir/8221/final>

## Additional Information

**A Methodology for Determining  
Forensic Data Requirements for  
Detecting Hypervisor Attacks**

Ramaswamy Chandramouli  
Anoop Singhal  
Duminda Wijesekera  
Changwei Liu

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13

Draft NISTIR 8221

# A Methodology for Determining Forensic Data Requirements for Detecting Hypervisor Attacks

Ramaswamy Chandramouli  
Anoop Singhal  
Duminda Wijesekera  
Changwei Liu

*Computer Security Division  
Information Technology Laboratory*

September 2018



U.S. Department of Commerce  
*Wilbur L. Ross, Jr., Secretary*

National Institute of Standards and Technology  
*Walter Copan, NIST Director and Under Secretary of Commerce for Standards and Technology*

46  
47  
48  
49  
50

National Institute of Standards and Technology Internal Report 8221  
27 pages (September 2018)

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

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

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

64

**Public comment period: *September 21, 2018 through October 12, 2018***

65  
66  
67  
68  
69  
70

National Institute of Standards and Technology  
Attn: Computer Security Division, Information Technology Laboratory  
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930  
Email: [NISTIR8221@nist.gov](mailto:NISTIR8221@nist.gov)

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

72

## Reports on Computer Systems Technology

73  
74  
75  
76  
77  
78  
79  
80  
81  
82

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

83

### Abstract

84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96

Hardware/Server Virtualization is a key feature of data centers used for cloud computing services and enterprise computing that enables ubiquitous access to shared system resources. Server virtualization is typically performed by a hypervisor, which provides mechanisms to abstract hardware and system resources from an operating system. Hypervisors are large pieces of software with several thousand lines of code and are therefore known to have vulnerabilities. This document analyzes the recent vulnerabilities associated with two open-source hypervisors—Xen and KVM—as reported by the National Institute of Standards and Technology’s (NIST) National Vulnerability Database (NVD), and develops a profile of those vulnerabilities in terms of hypervisor functionality, attack type, and attack source. Based on the predominant number of vulnerabilities in a hypervisor functionality (attack vector), two sample attacks using those attack vectors were launched to exploit those vulnerabilities, and the associated system calls were logged. The objective was to determine the evidence coverage for detecting and reconstructing those attacks and identify techniques required to gather missing evidence.

97

### Keywords

98  
99

cloud computing; forensic analysis; hypervisors; KVM; vulnerabilities; Xen

100  
101

## **Acknowledgments**

102 The authors thank Ms. Isabel Van Wyk for her valuable editorial review.

103

## **Audience**

104 The target audience for this document includes security staff and Chief Information Security  
105 Officers (CISO) in virtualized infrastructures used for enterprise computing needs or for offering  
106 cloud services.

107

## **Trademark Information**

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

**109 Executive Summary**

110 Hypervisors provide the mechanism that both creates and runs multiple operating systems (also  
111 called guest virtual machines) on a single physical platform (a host) in cloud environments. The  
112 increasing popularity of cloud services and the complex nature of hypervisors, which are  
113 essentially large software modules, have led to malicious attackers exploiting hypervisor  
114 vulnerabilities in order to attack cloud services. To discover recent trends in hypervisor attacks  
115 and prevent future hypervisor exploitation, recent vulnerability reports associated with two popular  
116 open-source hypervisors in the NIST National Vulnerability Database (NIST-NVD), Xen and  
117 KVM, were analyzed and classified based on the hypervisor functionalities (attack vector), attack  
118 type and attack source.

119 Ten functionalities traditionally provided by hypervisors re considered for the classification of  
120 hypervisor vulnerabilities. These functionalities include: (1) virtual CPUs, (2) symmetric  
121 multiprocessing, (3) soft memory management unit, (4) interrupt and timer mechanisms, (5) I/O  
122 and networking, (6) paravirtualized I/O, (7) VM exits, (8) hypercalls, (9) VM management and  
123 remote management software, and (10) hypervisor Add-ons. Based on functionalities, the  
124 vulnerability profile reveals that most attacks were caused by vulnerabilities in the soft memory  
125 management unit and I/O and networking functionalities. It also reveals that two most common  
126 hypervisor attacks are denial-of-service (DoS) and privilege escalation attacks launched primarily  
127 by guest OS users. Using vulnerabilities related to the hypervisor functionality of the soft memory  
128 management unit, two sample attacks were launched to obtain the evidence needed to perform  
129 forensic analysis on hypervisor attacks in which corresponding system calls were captured. The  
130 objective was to determine the evidence coverage for detecting and reconstructing those attacks  
131 and identify techniques required to gather missing evidence. A close analysis of these system calls  
132 reveals that more evidence regarding the execution path for the attacks is found in the run-time  
133 memory.

134 The methodology outlined in this document can assist cloud providers in enhancing the security  
135 of their virtualized infrastructure, help cloud service customers discover recent hypervisor attack  
136 trends, identify information that reveals the presence of such attacks, and provide guidance on  
137 taking proactive steps to prevent those attacks in the operating environment.

138

139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159

## Table of Contents

<b>Executive Summary</b> .....		<b>iv</b>
<b>1 Introduction</b> .....		<b>6</b>
<b>2 Background and Related Work</b> .....		<b>7</b>
2.1 Hypervisors .....		7
2.1.1 Xen .....		7
2.1.2 KVM .....		8
2.2 Related Work .....		8
<b>3 Deriving a Profile of Hypervisor Vulnerabilities</b> .....		<b>10</b>
3.1 The Vulnerabilities in the NIST-NVD .....		10
3.2 Associating Hypervisor Functionalities with Vulnerabilities .....		10
3.3 Deriving the Hypervisor Vulnerability Profile .....		12
<b>4 Sample Attacks and Forensic Analysis</b> .....		<b>15</b>
4.1 The Two Sample Attacks .....		15
4.2 Identifying Evidence Coverage for Forensic Analysis .....		16
4.3 Use of Virtual Machine Introspection (VMI) for Forensics .....		17
<b>5 Conclusions</b> .....		<b>18</b>
<b>Appendix A— Description of Hypervisor Functionality</b> .....		<b>19</b>
<b>Appendix B— The Syscalls Intercepted from the Attacking Program</b> .....		<b>22</b>
<b>Appendix C— References</b> .....		<b>24</b>

**1 Introduction**

161 Most cloud services are provided in a virtualized environment. Since virtualization of all system  
162 resources—including processors, memory, and I/O devices—makes it possible to run multiple  
163 operating systems on a single physical platform (host), virtualization is a key feature of cloud  
164 computing that enables ubiquitous access to shared pools of system resources and high-level  
165 services provisioned with minimal management effort [1, 2]. An Operating System (OS) directly  
166 controls hardware resources in a non-virtualized system, but virtualization, typically performed by  
167 a hypervisor (also called a virtual machine monitor or VMM) [3] within a cloud environment,  
168 provides a mechanism that abstracts the hardware and system resources from an OS. As a software  
169 layer that lies between the physical hardware and the Virtual Machines (VMs or guest machines),  
170 a hypervisor supports the guest machines by presenting the guest OSs with a virtual operating  
171 platform and managing their execution.

172 However, hypervisors are large pieces of software with many lines of code and known  
173 vulnerabilities [4]. While there is published research dedicated to characterizing and assessing  
174 hypervisor vulnerabilities as well as detecting and forensically analyzing the corresponding attacks  
175 [4-8], there is no formal framework for conducting forensic analysis on popular hypervisors, such  
176 as KVM and XEN. Motivated by the work presented in [4], which characterized hypervisor  
177 vulnerabilities as of July 2012 with the objective of preventing their exploitation, this document  
178 considers the recent vulnerability reports associated with Xen and KVM in the NIST National  
179 Vulnerability Database (NIST-NVD). The objective is to discover recent trends in hypervisor  
180 attacks, provide suggestions for mitigating hypervisor attack risks, and identify evidence of those  
181 attacks. The main contributions are as follows: (1) all recent hypervisor vulnerabilities of Xen and  
182 KVM (from years of 2016 and 2017) in NIST-NVD were analyzed and classified based on the  
183 hypervisor functionalities, the attack types, and the sources of attacks; (2) classifications of the  
184 recent Xen and KVM hypervisor vulnerabilities can provide suggestions for mitigating potential  
185 hypervisor attacks and enhancing the hypervisor resilience against known hypervisor  
186 vulnerabilities; (3) some sample attacks were simulated to show the methodology of determining  
187 the forensic data for detecting hypervisor attacks.

188 The rest of the publication is organized as follows. Section 2 presents the background of  
189 hypervisors and discusses related work. Section 3 lists typical hypervisor functionalities and shows  
190 analysis of the recent two-year hypervisor vulnerabilities listed in NIST-NVD. Section 4 describes  
191 the sample attacks and the forensic evidence used for reconstructing the sample attacks. Section 5  
192 provides conclusions.

## 193 **2 Background and Related Work**

194 This section provides an outline of the architectures of the two open-source hypervisors and  
195 discusses related work in the area of cloud forensic analysis.

### 196 **2.1 Hypervisors**

197 Hypervisors are software and/or firmware modules that virtualize system resources such as CPU,  
198 memory, and devices. In [9], Popek and Goldberg classify hypervisors as Type 1 hypervisor and  
199 Type 2 hypervisor. Type 1 hypervisors run directly on the host's hardware to control the hardware  
200 and manage guest operating systems (Guest OS). For this reason, Type 1 hypervisors are  
201 sometimes called bare metal hypervisors and include Xen, Microsoft Hyper-V, and VMware  
202 ESX/ESXi. Type 2 hypervisors are similar to other computer programs that run on an OS as a  
203 process. VMware Player, VirtualBox, Parallels Desktop for Mac, and QEMU are Type 2  
204 hypervisors. Some systems have features of both. For example, Linux's Kernel-based Virtual  
205 Machine (KVM) is a kernel module that effectively converts the host OS to a Type 1 hypervisor  
206 but is also categorized as a Type 2 hypervisor because Linux distributions are still general-purpose  
207 OSs with other applications competing for VM resources [10].

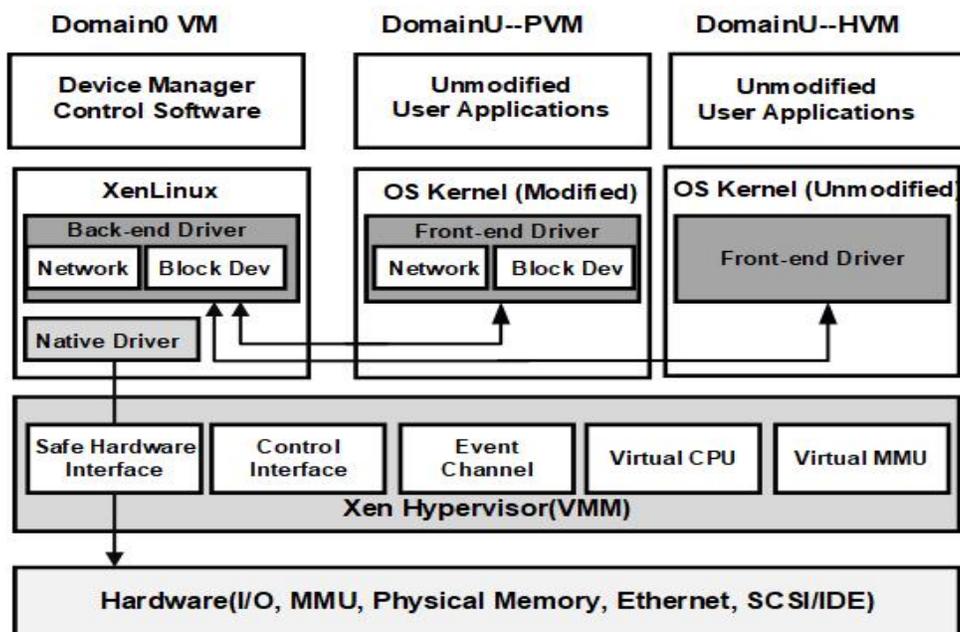
208 According to the 2015 State of Hyperconverged Infrastructure Market Report by ActualTech  
209 media [23], there are four popular hypervisors: Microsoft Hyper-V, VMware VSphere/ESX, Citrix  
210 XenServer/Xen, and KVM. Since Microsoft Hyper-V and VMware VSphere/ESX are commercial  
211 products, this document and research focus on the vulnerabilities on two widely used open-source  
212 hypervisors, Xen and KVM. Their architectures are briefly discussed below.

#### 213 **2.1.1 Xen**

214 Figure 1 shows the architecture of Xen. In this design, the Xen hypervisor manages three kinds of  
215 VMs including the control domain (also called Dom0) and guest domains (also called DomU) that  
216 support two different virtualization modes: Paravirtualization (PV) and Hardware-assisted  
217 Virtualization (HVM) [11]. Dom0 is the initial domain started by the Xen hypervisor on booting  
218 up a privileged domain that plays the administrator role and supplies services for DomU VMs. For  
219 the two kinds of DomU guests, PV is a highly efficient and lightweight virtualization technology  
220 introduced by XEN in which Xen PV does not require virtualization extensions from the host  
221 hardware. Thus, PV enables virtualization on hardware architectures that do not support HVM,  
222 but it requires PV-enabled kernels and PV drivers to power a high performance virtual server.  
223 HVM requires hardware extensions, and Xen typically uses QEMU (Quick Emulator), a generic  
224 hardware emulator [15], for simulating PC hardware (e.g., CPU, BIOS, IDE, VGA, network cards,  
225 and USBs). Because of the use of simulation technologies, HVM VMs' performance is inferior to  
226 PV VMs. Xen 4.4 provides a new virtualization mode named PVH. PVH guests are lightweight  
227 HVM-like guests that use virtualization extensions in the host hardware. Unlike HVM guests,  
228 instead of using QEMU to emulate devices, PVH guests use PV drivers for I/O and native OS  
229 interfaces for virtualized timers, virtualized interrupts, and a boot. PVH guests require PVH-  
230 enabled guest OS [11].

231 **2.1.2 KVM**

232 In the open-source hypervisor projects, the Kernel-based Virtual Machine (KVM) is a relatively  
 233 new product which was first introduced in 2006 and soon merged into the Linux kernel (2.6.20).  
 234 KVM is a full virtualization solution for Linux on x86 hardware containing virtualization  
 235 extensions (Intel VT or AMD-V) where VMs run as normal Linux processes [12]. Figure 2 shows  
 236 the KVM architecture, in which the KVM module uses QEMU to create guest VMs running as  
 237 separate user processes. Because KVM is installed on top of the host OS, it is considered a Type  
 238 2 hypervisor. However, KVM kernel module turns Linux kernel into a Type 1 bare-metal  
 239 hypervisor, providing the power and functionality of even the most complex and powerful Type 1  
 240 hypervisors.



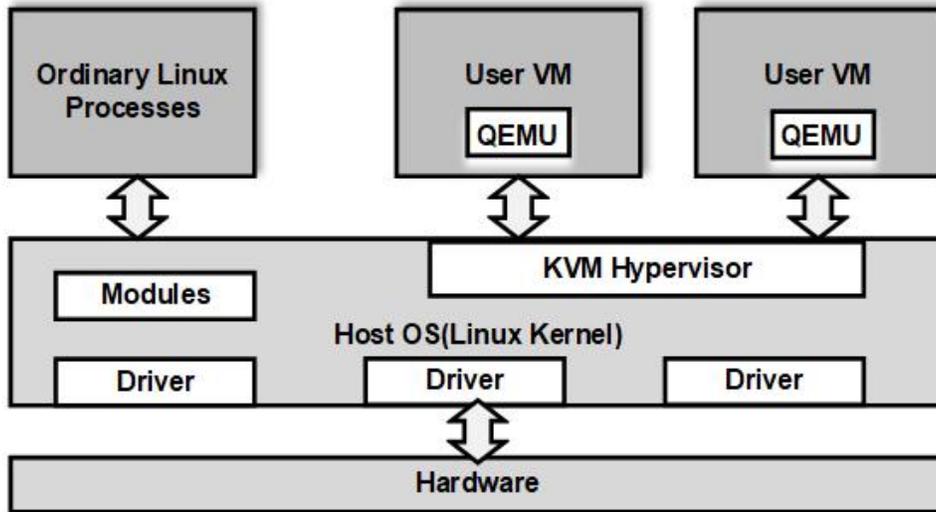
241

242

Figure 1: The Xen architecture

243 **2.2 Related Work**

244 Hypervisor attacks are categorized as external attacks and defined as exploits of the hypervisor's  
 245 vulnerabilities which allow attackers to gain accessibility and authorization over the hypervisors  
 246 [13]. In support of hypervisor defense, Perez-Botero et al. characterized Xen and KVM  
 247 vulnerabilities based on hypervisor functionalities in 2012 [4]. However, these cannot be used to  
 248 predict recent attack trends. To assess the weakness, severity scores, and attack impacts,  
 249 Thongthua et al. assessed the vulnerabilities of widely used hypervisors, including VMware ESXi,  
 250 Citrix XenServer, and KVM, using the NIST 800-115 security testing framework and performed  
 251 some sample experiments [5]. In an effort to develop hypervisor forensic methods, researchers  
 252 discussed the attacks on hypervisors, their forensic mechanisms and challenges [8], and leveraged  
 253 existing memory forensic techniques to perform forensic analysis on hypervisor attacks [7].  
 254



255

256

Figure 2: The KVM architecture

### 257 **3 Deriving a Profile of Hypervisor Vulnerabilities**

258 As a prelude to developing a methodology for determining forensic data requirements for detecting  
259 hypervisor attacks, it is necessary to derive a profile of recent hypervisor vulnerabilities in terms  
260 of the following classification criteria:

- 261 • Hypervisor Functionality where the vulnerability exists (attack vector)
- 262 • Attack Type (impact of the attack by exploiting the vulnerability)
- 263 • Attack Source (the component in the hypervisor platform from which the attack is  
264 launched)

265 The approach adopted for deriving the vulnerability profile involved obtaining all vulnerabilities  
266 (tagged with CVE numbers) in two open-source hypervisors (Xen and KVM) from the NIST-NVD  
267 for years 2016 and 2017. The hypervisor functionality (attack vector) was then associated with the  
268 attack type (impact) that resulted from exploiting each vulnerability and the attack source based  
269 on the description of vulnerabilities in that database. The total number of vulnerabilities for the  
270 two chosen open-source hypervisors in each of the three categories (attack vector, attack type, and  
271 attack source) thus provided a recent vulnerability profile for those hypervisor offerings.

272 A brief description of the information sources that were used and the steps adopted as part of the  
273 approach for deriving the vulnerability profile is given in sections 3.1, 3.2, and 3.3.

#### 274 **3.1 The Vulnerabilities in the NIST-NVD**

275 The NIST-NVD is the U.S. government repository of standards-based vulnerability management  
276 data and includes databases of security checklist references, security-related software flaws,  
277 misconfigurations, product names, and impact metrics [14]. A search of the NIST-NVD for the  
278 vulnerabilities posted during the years 2016 and 2017 revealed 83 Xen hypervisor vulnerabilities  
279 and 20 KVM hypervisor vulnerabilities. These vulnerabilities were then associated with the  
280 following:

- 281 • Hypervisor functionality where the vulnerability arises
- 282 • Potential attack type
- 283 • Attack source (i.e., the component/associated user from which the potential attack can  
284 be launched)

#### 285 **3.2 Associating Hypervisor Functionalities with Vulnerabilities**

286 To better understand different hypervisor vulnerabilities, Perez-Botero et al. considered 11  
287 functionalities that a traditional hypervisor provides and mapped vulnerabilities to them [4]. These  
288 functionalities include:

- 289 1) Virtual CPUs (vCPU)
- 290 2) Symmetric Multiprocessing (VSMP)
- 291 3) Soft Memory Management Unit (MMU)
- 292 4) I/O and Networking
- 293 5) Paravirtualized I/O

- 294           6)     Interrupt and Timer mechanisms  
 295           7)     Hypercalls  
 296           8)     VMEExit  
 297           9)     VM Management  
 298           10)    Remote Management Software  
 299           11)    Hypervisor Add-ons

300   Based on the common function provided by numbers four and five above, these were merged into  
 301   a single functionality. (A detailed description of all these functionalities can be found in Appendix  
 302   A). All reported Xen and KVM vulnerabilities during the years 2016 and 2017 were mapped to  
 303   these hypervisor functionalities based on the approach in [4]. A brief description of a sample  
 304   vulnerability associated with each functionality is given in Table 1 below:

305                                   **Table 1: A sample vulnerability for each hypervisor functionality**

Hypervisor Functionality	Sample Vulnerability
vCPU	CVE-2017-10923 is an example of vCPU vulnerability in which Xen through 4.8.x does not validate a vCPU array index upon sending a software generated interrupt(SGI), which allows a guest OS user to cause a denial-of-service(DoS) attack, finally resulting in crashing the hypervisor.
VSMP	NONE
Soft MMU	An example of soft MMU vulnerability is CVE-2017-17565, which existed up to Xen version 4.9.x. Due to an incorrect assertion related to M2P, this vulnerability allows a paravirtualized guest OS user to cause a DoS attack when both the shadow mode and log-dirty mode are set up and working.
I/O and Networking	CVE-2017-15589 is an example of an I/O and networking vulnerability discovered in Xen versions through 4.9.x which allows x86 HVM guest OS users to obtain sensitive information from the host OS (or an arbitrary guest OS). In these versions of Xen, at least one write path was found wherein the data that had been stored in an internal structure could contain bits from an uninitialized hypervisor stack slot. A subsequent emulated read would retrieve these bits.
Interrupt/Timer	CVE-2018-7542 is an example of an interrupt/timer vulnerability caused by leveraging the mishandling of configurations that lack a local APIC. It was discovered in Xen 4.8.x through 4.10.x. This vulnerability allows an x86 PVH guest OS user to cause a DoS attack (a NULL pointer dereference and hypervisor crash).
Hypercalls	An example of hypercall vulnerability is CVE-2017-8903, which is reported through Xen 4.8.x on 64-bit platforms that might allow a PV guest OS user to execute arbitrary code on the host OS by mishandling page tables after an IRET hypercall.

Hypervisor Functionality	Sample Vulnerability
VMExit	The exploit on VM Exit-handling code usually leads to a DoS attack. An example of VMExit vulnerability is CVE-2017-2596, in which the “nested_vmx_check_vmptr” function in arch/x86/kvm/vmx.c in the Linux kernel through 4.9.8 improperly emulates the VMXON instruction that puts the processor in VMX root mode. This then allows a KVM L1 guest OS user to cause a DoS attack (the host OS memory consumption) by leveraging the mishandling of page references.
VM Management	The exploit of the management functionality may allow a host compromise. An example of VM management functionality vulnerability is CVE-2016-5302. When a deployment has been upgraded from an earlier release, XenServer 7.0 before the vendor's Hot x XS70E003 may allow a remote attacker on the management network to compromise a host by leveraging credentials for an active directory account.
Remote Management Software	NONE
Hypervisor Add-ons	CVE-2016-0749 is an example vulnerability of hypervisor add-ons. By leveraging the smartcard interaction in SPICE as KVM add-ons, a remote attacker can cause a DoS attack (QEMU-KVM process crash) or possibly execute arbitrary code via vectors related to connecting to a guest VM, which triggers a heap-based buffer overflow.

306 **3.3 Deriving the Hypervisor Vulnerability Profile**

307 With the goal of deriving the hypervisor security vulnerability profile, 83 Xen and 20 KVM  
 308 vulnerabilities listed in the NIST-NVD for the years 2016 and 2017 were analyzed and classified  
 309 according to functionalities, attack types (impacts), and attack sources.

310 **Table 2: The vulnerabilities of Xen and KVM classified by functionality**

Number	Hypervisor Functionality	Xen	KVM
1	vCPU	6 (7%)	4 (20%)
2	VSMP	0 (0%)	0 (0%)
3	Soft MMU	34 (40%)	5 (25%)
4	I/O and Networking	24 (29%) Five are fully-virtualized; 19 are paravirtualized; none are direct access or self-virtualized.	4 (20%) All are fully-virtualized.
5	Interrupt/Timer	7 (8%)	3 (15%)

Number	Hypervisor Functionality	Xen	KVM
6	Hypercalls	3 (4%)	1 (5%)
7	VMExit	1 (1%)	2 (10%)
8	VM Management	8 (10%)	0 (0%)
9	Remote Management Software	0 (0%)	0 (0%)
10	Hypervisor Add-ons	0 (0%)	1 (5%)

311 Classifications based on the hypervisor functionalities are shown in Table 2. With the exception  
312 of the two functionalities of virtual symmetric multiprocessing and remote management software,  
313 all functionalities were reported as having vulnerabilities. The number of vulnerabilities and the  
314 percentages within each hypervisor offering are listed. The table reveals that there are more  
315 reported Xen vulnerabilities than KVM, which can be attributed to a broader user base for Xen.  
316 Furthermore, approximately 69% of the vulnerabilities in Xen and 45% of the vulnerabilities in  
317 KVM are concentrated in two functionalities—Soft MMU and I/O and Networking. A detailed  
318 reading of CVE reports reveals that these vulnerabilities primarily originated in page tables and  
319 I/O grant table emulation. Additionally, the vulnerabilities based on the I/O and Networking  
320 functionality were also associated with each of the four types of I/O virtualization: (1) fully  
321 virtualized devices, (2) paravirtualized devices, (3) direct access devices, and (4) self-virtualized  
322 devices. Table 2 shows that most of the I/O and networking vulnerabilities in Xen came from  
323 paravirtualized devices, while all I/O and networking vulnerabilities in KVM came from fully-  
324 virtualized devices. This is due to the fact that in most Xen deployments, I/O and networking  
325 functionality is configured using a paravirtualized device, while in KVM, that functionality is  
326 configured using a fully virtualized device.

327

**Table 3: The types of attacks caused by Xen and KVM vulnerabilities**

Type of Attack	Xen	KVM
Denial-of-service (DoS)	48 (four have other impacts) (44%)	17 (three have other impacts) (63%)
Privilege escalation	33 (16 have other impacts) (30%)	3 (two have other impacts) (11%)
Information leakage	15 (five have other impacts) (14%)	5 (19%)
Arbitrary code execution	8 (two have other impacts) (7%)	2 (all have other impacts) (7%)
Reading/modifying/deleting a file	3 (3%)	0 (0%)
Others including compromising a host, canceling other administrators' operations and corrupting data	3 (3%)	0 (0%)

328 Classifications based on the attack types and the sources of attacks are listed in Table 3 and  
 329 Table 4. Table 3 reveals that the most common attack was DoS (44% for Xen and 63% for  
 330 KVM), indicating that attacking cloud services' availability has been the most serious cloud  
 331 security problem. The other top attacks were privilege escalation (30% for Xen and 11% for  
 332 KVM), information leakage (14% for Xen and 19% for KVM), and arbitrary code execution (7%  
 333 for Xen and 7% for KVM). Although each of these three attacks occurs with less frequency than  
 334 a DoS attack, they all result in more serious damage by allowing attackers to obtain sensitive  
 335 user information or compromise the hosts or guest VMs. Table 4 shows that the greatest source  
 336 of all attacks was guest OS users (76% for Xen and 85% for KVM), though other sources  
 337 included cloud administrators, guest OS administrators, and remote users. This suggests that  
 338 cloud providers must closely monitor guest users' activities in order to reduce attack risks.

339

**Table 4: Attack Sources and Number of Exploits**

Source of Attack	Xen	KVM
Administrator	2 (Management) (2%)	0 (0%)
Guest OS administrator	17 (including HVM and PV administrators) (20%)	1 (5%)
Guest OS user	63 (including ARM, X86, HVM and PV users) (76%)	17 (including KVM L1, L2, and privileged users) (85%)
Remote attacker	1 (1%)	1 (including an authenticated remote guest user) (5%)
Host OS user	0 (0%)	1 (5%)

340

## 341 **4 Sample Attacks and Forensic Analysis**

342 Since numerous vulnerabilities are related to Xen soft MMU functionality, this section will show  
343 two sample attacks, including those that exploit vulnerabilities CVE-2017-7228 and CVE-2016-  
344 6258, to demonstrate how the evidence for detecting and reconstructing hypervisor attacks is  
345 determined.

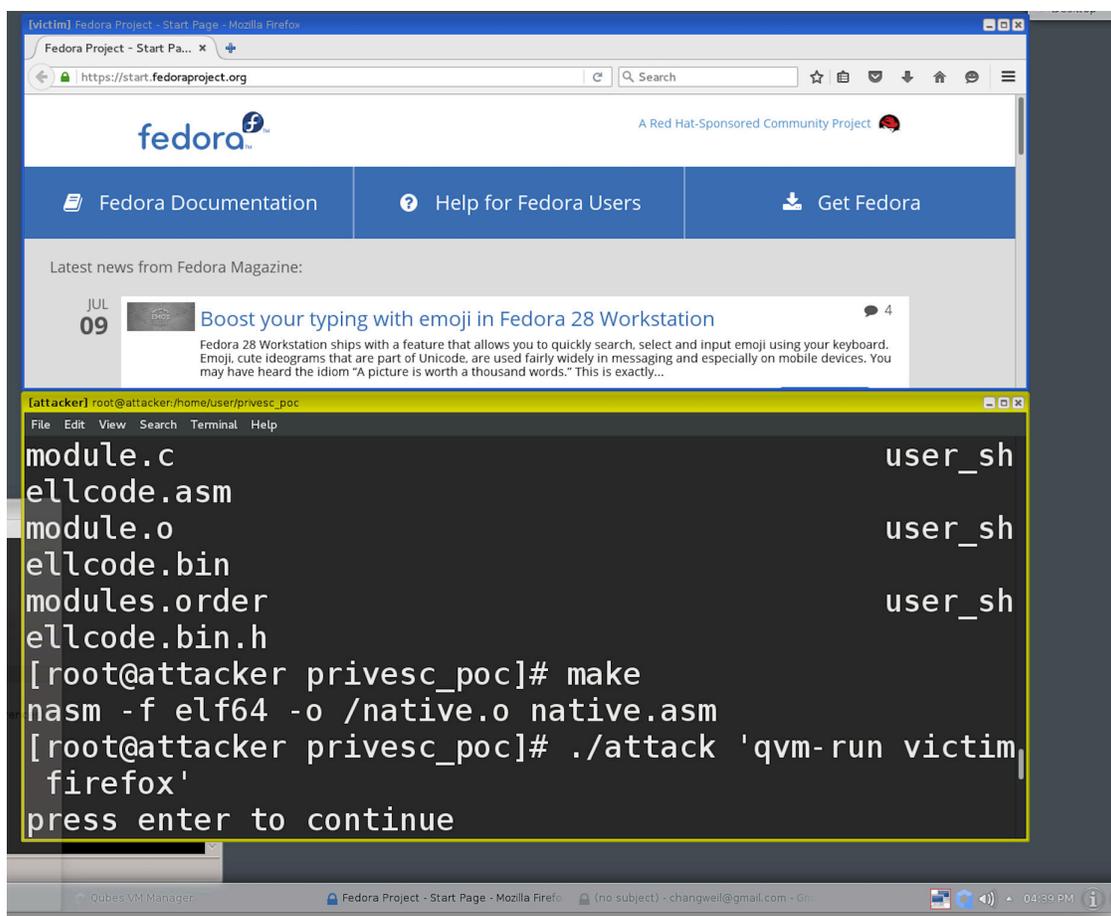
### 346 **4.1 The Two Sample Attacks**

347 As presented in Section 2.1.1., the Xen hypervisor manages three kinds of VMs, including the  
348 control domain (also called Dom0) and guest domains (also called DomU). These then support  
349 two different virtualization modes: Paravirtualization (PV) and Hardware-assisted Virtualization  
350 (HVM). The PV module has been widely utilized for its higher performance [25]. However,  
351 because the Xen PV model uses complex code to emulate the MMU, it introduces many  
352 vulnerabilities, such as CVE-2017-7228 and CVE-2016-6258.

353 Known by Xen as XSA-212, CVE-2017-7228 was first reported by Jann Horn of Google's Project  
354 Zero in 2017 [20]. Horn discovered that this vulnerability in X86 64 bit Xen (including 4.8.x, 4.7.x,  
355 4.6.x, 4.5.x, and 4.4.x versions) was caused by an insufficient check on the function  
356 "XENMEM\_exchange", which allows the PV guest user as the function caller to access hypervisor  
357 memory outside of the PV guest VM's provisioned memory. Therefore, a malicious 64-bit PV  
358 guest who can make a hypercall "HYPERVISOR\_memory\_op" function to invoke the  
359 "XENMEM\_exchange" function may be able to access all of a system's memory, allowing for  
360 VM escape (the process of breaking out of a guest VM and interacting with the hypervisor's host  
361 operating system) from DomU to Dom0, hypervisor host crash, and information leakage. With  
362 these attacks, the PV guest from "attacker" (the green terminal) could execute commands like  
363 "qvm-run victim firefox" to open a Firefox web-browser in "victim" guest VM, which can only be  
364 executed by Dom0 as shown in Figure 3.

365 CVE-2016-6258 is also known as XSA-182, which was reported by Jeremie Boutoille from  
366 Quarklab in 2016 [21]. In the PV module, page tables are used to map pseudo-physical/physical  
367 addresses seen by the guest VM to the underlying memory of the machine. Since there is a  
368 vulnerability in XEN PV page tables that allows updates to be made to pre-existing page table  
369 entries, the malicious PV guests can access the page directory with an updated write privilege to  
370 execute the VM escape, breaking out of DomU to control Dom 0.

371 Both types of attacks were launched on the PV module configured in Qubes 3.1 with Xen 4.6 [22].  
372 As illustrated in Figure 3, the attacker impersonating the PV guest root user could execute a  
373 command, "qvm-run victim firefox," that can only be executed by Dom0 to open the victim PV  
374 guest's Firefox web browser. Both attacks allowed the PV guest users to gain the control of Dom0.



375

376

Figure 3: CVE-2017-7228 and CVE-2016-6258 Attacks

377

## 4.2 Identifying Evidence Coverage for Forensic Analysis

378 Both attacks used vulnerabilities related to hypercalls and soft MMU in Xen in addition to using  
 379 Xen's device activity logs. The affected processes' runtime syscalls were therefore logged to  
 380 perform a forensic analysis. As an example, Appendix B illustrates the syscalls obtained by using  
 381 the "strace" Linux command on the running "attack" program of CVE-2017-7228. Analysis of the  
 382 device activity logs and runtime syscalls showed the relevant evidence originated from the syscalls  
 383 captured from the attackers' VMs. Despite the noise among syscalls that can be found in most  
 384 programs, other syscalls revealed that the attack program injected a loadable kernel module into  
 385 the kernel space which exploited the vulnerability to control the Dom0. This then opened the  
 386 Firefox browser in the victim's guest VM.

387 Evidence acquisition plays an important role in forensic analysis by determining and  
 388 reconstructing attacks. As presented in a previous work which illustrated the use of a layered  
 389 graphical framework to reconstruct attack scenarios [24], relevant evidence was identified and  
 390 collected to reconstruct the corresponding attack path(s) representing the attack scenarios. During  
 391 this process, an attack path with missing attack steps led to the collection of additional supporting  
 392 evidence. An analysis of the syscalls captured for two sample attacks revealed that while the  
 393 syscalls obtained using "strace" Linux command were useful for forensic analysis, they lacked

394 attack details and had the following deficiencies: (1) the syscalls did not provide details of how  
395 features of the loadable kernel module used Xen's memory management to launch the attack; and  
396 (2) the syscalls were collected from the attacker's guest VM, which could easily be tampered with  
397 or removed by the attacker. The VM introspection technique and corresponding memory analysis  
398 tools are therefore recommended to obtain more supporting and admissible evidence from the run-  
399 time memory.

#### 400 **4.3 Use of Virtual Machine Introspection (VMI) for Forensics**

401 The VMI is a process that allows for the external viewing of the state of a VM, either from a  
402 privilege VM or VMM itself. The state information includes CPU state (e.g., registers), all  
403 memory, and all I/O device states such as the contents of storage devices or register states of I/O  
404 controllers. Leveraging this capability, VMI-based applications can be built to perform forensic  
405 analysis in the following ways:  
406

- 407 1. The VMI-based application can capture the entire memory and I/O state of a VM that is  
408 suspected of being compromised or attacked by taking a checkpoint (taking a snapshot).  
409 The captured state of the running VM under observation can be compared to either: (a) a  
410 suspended VM in a known good state or (b) the original VM image from which the running  
411 VM was instantiated. [26].
- 412 2. A VMI-based application can be built to perform execution path analysis on the monitored  
413 VM. This is achieved by tracing—analyzing the sequence of VM activities and the  
414 corresponding complete VM state (e.g., memory map, IO access). This aids in the  
415 construction a detailed attack graph with the VM state as nodes and the VM activities as  
416 edges, thereby tracing the path through which the current compromised state was reached  
417 [27]. This approach addresses deficiencies in performing forensic analysis that simply uses  
418 the system calls from the compromised VMs as follows:
  - 419 • There is the possibility that syscalls/hypercalls from the compromised VM could  
420 be tampered with or entirely removed by the attacker. In this approach, the sequence  
421 of VM states and VM activities are captured from outside the compromised VM,  
422 thus eliminating this possibility.
  - 423 • All variables that characterize a VM state and a VM activity are captured, helping  
424 to reconstruct the attack details based on memory access information with the  
425 ability to detect even malicious attacks, such as code and data modification.

426

**427 5 Conclusions**

428 An analysis of all reported vulnerabilities on Xen and KVM in the last two years was conducted,  
429 and two sample attacks were launched to identify evidence for a forensic analysis. Data  
430 subsequently showed that most attacks on the two hypervisors were caused by vulnerabilities that  
431 existed in soft MMU and I/O and Networking functionalities. The two most common hypervisor  
432 attacks were DoS and privilege escalation attacks. Most attackers are guest OS users. The collected  
433 evidence on the sample attacks showed that most valuable evidence remains in the run-time system  
434 memory. Therefore, to obtain valuable evidence with guaranteed integrity, VM introspection  
435 technique and secure logging systems showing memory access should be implemented and used.

## 436 **Appendix A—Description of Hypervisor Functionality**

437 **Virtual CPUs (vCPU):** A vCPU, also known as a virtual processor, abstracts a portion or share  
438 of a physical CPU that is assigned to a virtual machine (VM). The hypervisor uses a portion of the  
439 physical CPU cycle and allocates it to a vCPU assigned to a VM. The hypervisor schedules vCPU  
440 tasks to the physical CPUs.

441 **Virtual Symmetric Multiprocessing (VSMP):** VSMP is a method of symmetric multiprocessing  
442 (SMP), which enables multiple vCPU belonging to the same VM to be scheduled to a physical  
443 CPU that has at least two logical processors.

444 **Soft Memory Management Unit (Soft MMU):** The Memory Management Unit (MMU) is the  
445 hardware responsible for managing memory by translating the virtual addresses manipulated by  
446 the software into physical addresses. In an OS running on bare metal, the MMU translates the  
447 virtual addresses manipulated by the software into physical addresses. The mappings from virtual  
448 to physical addresses are kept in page tables (PT) and managed by the OS. In a virtualized  
449 environment, the hypervisor emulates the MMU (therefore called the soft MMU) for the guest  
450 OSs. This is done by mapping what the guest OS sees as physical memory (often called pseudo-  
451 physical/physical address in Xen) to the underlying memory of the machine (called machine  
452 addresses in Xen). The mapping table from the physical address to machine address (P2M) is  
453 typically maintained in the hypervisor and hidden from the guest OS by using a shadow page table  
454 for each guest VM. Each shadow page table mapping translates virtual addresses of programs in a  
455 guest VM to guest (pseudo) physical addresses and is placed in the guest OS [16, 17]. The Xen  
456 paravirtualized MMU model requires that the guest OS be directly aware of mapping between  
457 (pseudo) physical and machine addresses (the P2M table). Additionally, in order to read page table  
458 entries that contain machine addresses and convert them back into (pseudo) physical addresses, a  
459 translation from machine to (pseudo) physical addresses provided by the M2P table is required in  
460 Xen paravirtualized MMU model [17].

461 **I/O and Networking:** There are three common approaches that provide I/O services to guest  
462 VMs. Using the Xen I/O structures illustrated in Figure 4 as an example, these common approaches  
463 include:

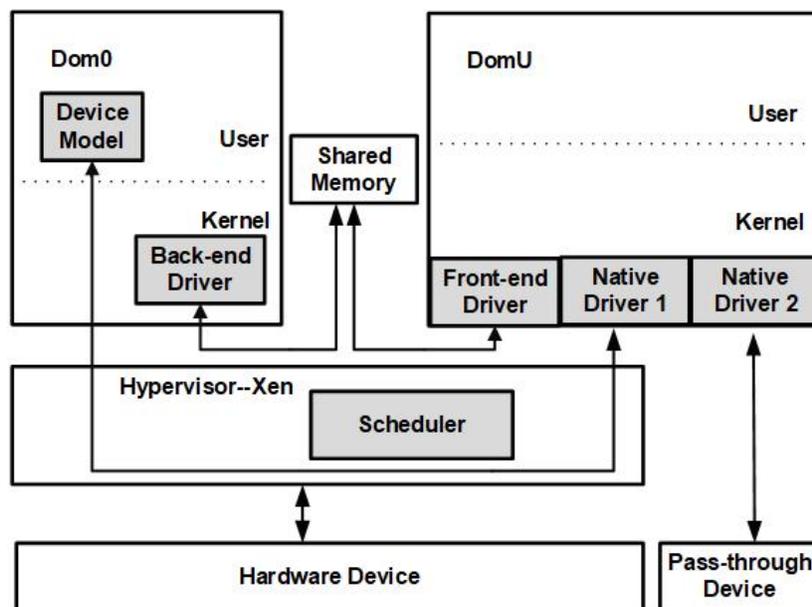
464 (1) the hypervisor emulates a known I/O device in a fully virtualized system, and the guests  
465 use an unmodified driver (called a native driver) to interact with it (illustrated as “Native  
466 Driver 1” in DomU to “Device Model” in Dom0 in Figure 4);

467 (2) a paravirtual driver (known as a front-end driver) in a paravirtualized system is installed  
468 in the modified guest OS in DomU, which uses shared-memory—asynchronous buffer-  
469 descriptor rings—to communicate with the back-end I/O driver in the hypervisor  
470 (illustrated as “Front-end Driver” in DomU to “Back-end Driver” to Dom0 in Figure 4);

471 (3) the host assigns a device (known as a pass-through device) directly to the guest VM  
472 (illustrated as “Native Driver 2” in DomU to “Pass-through Device” in Figure 4).

473 To reduce I/O virtualization overhead, improve virtual machine performance, and provide I/O  
474 services to guest VMs, scalable self-virtualizing I/O devices that allow direct access interface to  
475 multiple VMs are also used. However, the two approaches do not virtualize the I/O since they

476 include direct access, and self-virtualized I/O devices allow the device driver within a guest OS to  
 477 interact with the hardware directly. Furthermore, they scale poorly due to challenges, performance,  
 478 and cost [22].



479

480

Figure 4: Xen I/O structures

481 In paravirtualized Xen systems, the front-end and back-end drivers communicate with each other  
 482 using two producer-consumer ring buffers (standard lockless shared memory data structures built  
 483 on grant tables and event channels), where one is used for packet reception and the other is used  
 484 for packet transmission. Though hypervisors enforce isolation across VMs residing within a single  
 485 physical machine, the grant mechanism provides inter-domain communications in Xen, allowing  
 486 shared-memory communications between unprivileged domains by using grant tables [16]. Grant  
 487 tables are used to protect the I/O buffer in a guest domain's memory and share the I/O buffer with  
 488 Dom0 properly, which underpin the split device drivers for block and network I/O. Each domain  
 489 has its own grant table that allows the domain to inform Xen with the kind of permissions other  
 490 domains have on their pages. KVM typically uses Virtio, a virtualization standard for network and  
 491 disk drivers, which is architecturally similar to Xen paravirtualized device drivers which are  
 492 composed of front-end drivers and back-end drivers.

493 **Interrupt/Timer:** Hypervisors should be able to virtualize and manage interrupts/timers [18], the  
 494 interrupt/timer controller of the guest OS, and the guest OS's access to the controller. The  
 495 interrupt/timer mechanism in a hypervisor includes a programmable interval timer (PIT), the  
 496 advanced programmable interrupt controller (APIC), and the interrupt request (IRQ) mechanisms  
 497 [4].

498 **Hypercall:** Hypercalls are similar to system calls (syscalls) that provide user-space applications  
 499 with kernel-level operations. They are performed using the syscall instruction with up to six  
 500 arguments passed in registers. A hypercall layer is commonly available and allows guest OSs to  
 501 make requests of the host OS. Domains will use hypercalls to request privileged operations such

502 as updating page tables from the hypervisors. Thus, an attacker can use hypercalls to attack the  
503 hypervisor from a guest VM.

504 **VMExit**: According to Belay et al. [19], the mode change from Virtual Machine Extension (VMX)  
505 root mode to VMX non-root mode is called VMEntry, and the mode change from VMX non-root  
506 mode to VMX root mode is called VMExit. VM exits are a response to some instructions and  
507 events (e.g., page fault) from guest VMs and are the main cause of performance degradation in a  
508 virtualized system. These events could include external interrupts, triple faults, task switches, I/O  
509 operation instructions (e.g., INB, OUTB), and accesses to control registers.

510 **VM management functionality**: Hypervisors support basic VM management functionalities,  
511 including starting, pausing, or stopping VMs. These tasks are implemented in Xen Dom0 and  
512 KVM's libvirt driver.

513 **Remote Management Software**: Remote management software is employed as a user-friendly  
514 interface that connects directly to the hypervisor in order to provide additional management and  
515 monitoring tools. With an intuitive user interfaces that visualizes the status of a system, the remote  
516 management software allows administrators to tweak or manage the virtualized environment.

517 **Add-ons**: The add-ons of hypervisors use modular designs to add extended functions. By  
518 leveraging the interaction between the add-ons and hypervisors, an attacker can cause a host to  
519 crash (a DoS attack) or even compromise the host.

520

**521 Appendix B—The Syscalls Intercepted from the Attacking Program**

522 The syscalls in this appendix were obtained by employing Linux command “strace” on the running  
 523 attack program using the vulnerability CVE-2017-7228 (the attack program is named “attack”).  
 524 These syscalls show: (1) the attacker executed the attack program with arguments aimed at the  
 525 victim guest VM (Line 1); (2) the attack program and required Linux libraries have been loaded  
 526 to the memory for the program execution (Line 2 to Line 16); (3) the memory pages of the attack  
 527 program have been protected from accessed by other processes (Line 17 to Line 23); and (4) the  
 528 attack program injected a loadable Linux module named “test.ko” to the kernel space to exploit  
 529 the vulnerability (Line 24 to Line 31).

```

530     1. execve("./attack", ["/./attack", "qvm-run victim firework"], [/* 30 vars */]) = 0
531     2. brk(NULL) = 0x8cd000
532     3. mmap(NULL, 4096, PROT_READ|PROT_WRITE,
533        MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fa3a3022000
534     4. access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
535     5. open("/etc/ld.so.cache", O_RDONLY|O_CLOEXEC) = 3
536     6. fstat(3, {st_mode=S_IFREG|0644, st_size=74105, ...}) = 0
537     7. mmap(NULL, 74105, PROT_READ, MAP_PRIVATE, 3, 0) = 0x7fa3a300f000
538     8. close(3) = 0
539     9. open("/lib64/libc.so.6", O_RDONLY|O_CLOEXEC) = 3
540    10. read(3, "\177ELF\2\1\1\3\0\0\0\0\0\0\0\3\0>\0\1\0\0\0\240\6\2\0\0\0\0"..., 832) = 832
541    11. fstat(3, {st_mode=S_IFREG|0755, st_size=2104216, ...}) = 0
542    12. mmap(NULL, 3934688, PROT_READ|PROT_EXEC,
543        MAP_PRIVATE|MAP_DENYWRITE, 3, 0) = 0x7fa3a2a42000
544    13. mprotect(0x7fa3a2bf9000, 2097152, PROT_NONE) = 0
545    14. mmap(0x7fa3a2df9000, 24576, PROT_READ|PROT_WRITE,
546        MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x1b7000) = 0x7fa3a2df9000
547    15. mmap(0x7fa3a2dff000, 14816, PROT_READ|PROT_WRITE,
548        MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x7fa3a2dff000
549    16. close(3) = 0
550    17. mmap(NULL, 4096, PROT_READ|PROT_WRITE,
551        MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fa3a300e000
552    18. mmap(NULL, 4096, PROT_READ|PROT_WRITE,
553        MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fa3a300d000
554    19. mmap(NULL, 4096, PROT_READ|PROT_WRITE,
555        MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fa3a300c000
556    20. arch_prctl(ARCH_SET_FS, 0x7fa3a300d700) = 0
557    21. mprotect(0x7fa3a2df9000, 16384, PROT_READ) = 0
558    22. mprotect(0x600000, 4096, PROT_READ) = 0
559    23. mprotect(0x7fa3a3023000, 4096, PROT_READ) = 0
560    24. munmap(0x7fa3a300f000, 74105) = 0
561    25. open("test.ko", O_RDONLY) = 3
562    26. finit_module(3, "user_shellcmd_addr=1407334317317"..., 0) = 0
563    27. fstat(1, {st_mode=S_IFCHR|0620, st_rdev=makedev(136, 0), ...}) = 0
564    28. mmap(NULL, 4096, PROT_READ|PROT_WRITE,
565        MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fa3a3021000
  
```

```
566 29. mmap(0x600000000000, 4096, PROT_READ|PROT_WRITE,  
567     MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS|MAP_LOCKED, -1, 0) =  
568     0x600000000000  
569 30. delete_module("test", O_NONBLOCK)    = 0  
570 31. exit_group(0)                        = ?  
571
```

572 **Appendix C—References**

- 573 [1] R. Uhlig, G. Neiger, D. Rodgers, A. L. Santoni, F. C. Martins, A. V. Anderson, S. M.  
574 Bennett, A. Kagi, F. H. Leung and L. Smith, Intel virtualization technology, *Computer*,  
575 38(5), pp.48-56, 2005.
- 576 [2] P. Mell and T. Grance, The NIST definition of cloud computing, *Communications of the*  
577 *ACM*, 53(6), p.50, 2010.
- 578 [3] R. P. Goldberg, Survey of virtual machine research, *Computer*, 7(6), pp.34-45, 1974.
- 579 [4] D. Perez-Botero, J. Szefer and R. B. Lee, Characterizing hypervisor vulnerabilities in cloud  
580 computing servers, In *Proceedings of the 2013 International Workshop on Security in*  
581 *Cloud Computing* (pp. 3-10). ACM, May 2013.
- 582 [5] A. Thongthua and S. Ngamsuriyaroj, Assessment of hypervisor vulnerabilities, In  
583 *International Conference Cloud Computing Research and Innovations (ICCCRI)*, pp. 71-  
584 77, May 2016.
- 585 [6] J. Szefer, E. Keller, R. B. Lee and J. Rexford, Eliminating the hypervisor attack surface for  
586 a more secure cloud, In *Proceedings of the 18th ACM Conference on Computer and*  
587 *Communications Security* (pp. 401-412). ACM, Oct 2011.
- 588 [7] M. Graziano, A. Lanzi and D. Balzarotti, Hypervisor memory forensics, In *International*  
589 *Workshop on Recent Advances in Intrusion Detection* (pp. 21-40). Springer, Berlin,  
590 Heidelberg. 2013.
- 591 [8] L. M. Joshi, M. Kumar and R. Bharti, Understanding threats in hypervisor, its forensics  
592 mechanism and its research challenges, *International Journal of Computer Applications*  
593 119.1 (2015).
- 594 [9] G. J. Popek and R. P. Goldberg, Formal requirements for virtualizable third generation  
595 architectures, *Communications of the ACM* 17.7 (1974): 412-421.
- 596 [10] B. Pariseau, KVM reignites Type 1 vs Type 2 hypervisor debate, retrieved from  
597 [https://searchservervirtualization.techtarget.com/news/2240034817/KVM-reignites-Type-](https://searchservervirtualization.techtarget.com/news/2240034817/KVM-reignites-Type-1-vs-Type-2-hypervisor-debate)  
598 [1-vs-Type-2-hypervisor-debate](https://searchservervirtualization.techtarget.com/news/2240034817/KVM-reignites-Type-1-vs-Type-2-hypervisor-debate) on Apr-11-2018.
- 599 [11] Xen project software overview, retrieved from  
600 [https://wiki.xen.org/wiki/Xen\\_Project\\_Software\\_Overview](https://wiki.xen.org/wiki/Xen_Project_Software_Overview) on Apr-11- 2018.
- 601 [12] KVM, retrieved from [https://www.linux-kvm.org/page/Main\\_Page](https://www.linux-kvm.org/page/Main_Page) on Apr-12-2018.
- 602 [13] J. Shi, Y. Yang and C. Tang, Hardware assisted hypervisor introspection, SpringerPlus,  
603 vol. 5, no. 1, 2016.
- 604 [14] NIST National Vulnerability Database, retrieved from <https://nvd.nist.gov> on Apr-12-  
605 2018.
- 606 [15] QEMU--the FAST! processor emulator, retrieved from <https://www.qemu.org> on Apr-12-  
607 2008.
- 608 [16] J.F. Kloster, J. Kristensen and A. Mejlholm, Efficient memory sharing in the Xen virtual  
609 machine monitor, Department of Computer Science, Aalborg University (Jan. 2006).

- 610 [17] X86 paravirtualised memory management, retrieved from  
611 [https://wiki.xen.org/wiki/X86\\_Paravirtualised\\_Memory\\_Management](https://wiki.xen.org/wiki/X86_Paravirtualised_Memory_Management).
- 612 [18] Y. Song, H. Wang and T. Soyata, Hardware and software aspects of VM-based mobile-  
613 cloud offloading, Enabling Real-Time Mobile Cloud Computing through Emerging  
614 Technologies, pp.247-271, 2015.
- 615 [19] A. Belay, A. Bittau, A. J. Mashtizadeh, D. Terei, D. Mazieres and C. Kozyrakis, Dune:  
616 safe user-level access to privileged CPU features, October 2012, In Osdi (Vol. 12, pp. 335-  
617 348).
- 618 [20] Pandavirtualization: exploiting the Xen hypervisor, Retrieved from  
619 <https://googleprojectzero.blogspot.com/2017/04/pandavirtualization-exploiting-xen.html>  
620 on May-30-2018.
- 621 [21] Xen exploitation part 3: XSA-182, Qubes escape, Retrieved from  
622 <https://blog.quarkslab.com/xen-exploitation-part-3-xsa-182-qubes-escape.html>.
- 623 [22] J. Satran, L. Shalev, M. Ben-Yehuda and Z. Machulsky, Scalable I/O-a well-architected  
624 way to do scalable, secure and virtualized I/O, In Workshop on I/O Virtualization, Dec  
625 2008.
- 626 [23] 2015 state of hyperconverged infrastructure market report, Retrieved from  
627 [https://www.actualtechmedia.com/wp-content/uploads/2015/05/2015-State-of-  
628 Hyperconverged-Infrastructure-Market-Report.pdf](https://www.actualtechmedia.com/wp-content/uploads/2015/05/2015-State-of-Hyperconverged-Infrastructure-Market-Report.pdf) on Aug-3-2018.
- 629 [24] C. Liu, A. Singhal and D. Wijesekera, A layered graphical model for mission attack impact  
630 analysis, In Communications and Network Security (CNS), Oct, 2017 IEEE Conference  
631 on (pp. 602-609).
- 632 [25] H. Fayyad-Kazan, L. Perneel and M. Timmerman, Full and para-virtualization with Xen:  
633 a performance comparison. Journal of Emerging Trends in Computing and Information  
634 Sciences, 4(9), 2013.
- 635
- 636