

Withdrawn Draft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The attached draft document has been withdrawn, and is provided solely for historical purposes. It has been superseded by the document identified below.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 May 27, 2021

Original Release Date April 28, 2020

Superseding Document

Status Draft

Series/Number NIST Interagency or Internal Report 8320

Title Hardware-Enabled Security: Enabling a Layerd Approach to Platform Security for Cloud and Edge Computing Use Cases

Publication Date May 2021

DOI <https://doi.org/10.6028/NIST.IR.8320-draft>

CSRC URL <https://csrc.nist.gov/publications/detail/nistir/8320/draft>

Additional Information

Hardware-Enabled Security for Server Platforms:

Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases

Michael Bartock
Murugiah Souppaya
*Computer Security Division
Information Technology Laboratory*

Ryan Savino
Tim Knoll
Uttam Shetty
Mourad Cherfaoui
Raghu Yeluri
*Intel Data Platforms Group
Santa Clara, CA*

Karen Scarfone
*Scarfone Cybersecurity
Clifton, VA*

April 28, 2020

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.CSWP.04282020-draft>

29

Abstract

30 In today's cloud data centers and edge computing, attack surfaces have significantly increased,
31 hacking has become industrialized, and most security control implementations are not coherent
32 or consistent. The foundation of any data center or edge computing security strategy should be
33 securing the platform on which data and workloads will be executed and accessed. The physical
34 platform represents the first layer for any layered security approach and provides the initial
35 protections to help ensure that higher-layer security controls can be trusted. This white paper
36 explains hardware-based security techniques and technologies that can improve platform security
37 and data protection for cloud data centers and edge computing.

38

Keywords

39 confidential computing; container; hardware-enabled security; hardware security module (HSM);
40 secure enclave; trusted execution environment (TEE); trusted platform module (TPM);
41 virtualization.

42

Disclaimer

43 Any mention of commercial products or reference to commercial organizations is for information
44 only; it does not imply recommendation or endorsement by NIST, nor does it imply that the
45 products mentioned are necessarily the best available for the purpose.

46

Additional Information

47 For additional information on NIST's Cybersecurity programs, projects, and publications, visit
48 the [Computer Security Resource Center](#). Information on other efforts at [NIST](#) and in the
49 [Information Technology Laboratory](#) (ITL) is also available at www.nist.gov and
50 www.nist.gov/itl.

51

Public Comment Period: April 28, 2020 through June 2, 2020

52

National Institute of Standards and Technology

53

Attn: Computer Security Division, Information Technology Laboratory

54

100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

55

Email: hwsec@nist.gov

56

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

57

Acknowledgments

58 The authors thank their colleagues from Intel Corporation who contributed their time and
59 expertise to the development of this white paper, including Alex Eydelberg, Sugumar
60 Govindarajan, Kapil Sood, Jeanne Guillory, David Song, Scott Raynor, Scott Huang, Matthew
61 Arenó, Charlie Stark, Subomi Laditan, Kamal Natesan, Haidong Xia, Jerry Wheeler, Dhinesh
62 Manoharan, and John Pennington.

63

Audience

64 The primary audiences for this white paper are security professionals, such as security engineers
65 and architects; system administrators and other information technology (IT) professionals for
66 cloud service providers; and hardware, firmware, and software developers who may be able to
67 leverage hardware-based security techniques and technologies to improve platform security for
68 cloud data centers and edge computing.

69

Trademark Information

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

71

72

73 **Table of Contents**

74 **1 Introduction 1**

75 **2 Hardware Platform Security Overview 2**

76 **3 Platform Integrity Verification 4**

77 3.1 Hardware Security Module (HSM) 4

78 3.2 The Chain of Trust (CoT) 5

79 3.2.1 Technology Example: Intel Trusted Execution Technology (TXT) 6

80 3.2.2 Technology Example: Intel Boot Guard 6

81 3.2.3 Technology Example: UEFI Secure Boot (SB) 7

82 3.2.4 Technology Example: Intel Platform Firmware Resilience (PFR) 8

83 3.3 Supply Chain Protection 10

84 3.3.1 Technology Example: Intel Transparent Supply Chain (TSC) 10

85 3.3.2 Technology Example: PFR with Protection in Transit (PIT) 10

86 3.4 Technology Example Summary 11

87 **4 Data Protection and Confidential Computing 13**

88 4.1 Memory Isolation 13

89 4.2 Application Isolation 14

90 4.2.1 Technology Example: Intel Software Guard Extensions (SGX) 14

91 4.3 VM Isolation 15

92 4.4 Cryptographic Acceleration 15

93 4.4.1 Technology Example: Intel QuickAssist Technology (QAT) with Intel

94 Key Protection Technology (KPT) 15

95 4.5 Technology Example Summary 16

96 **5 Remote Attestation Services 17**

97 5.1 Platform Attestation 17

98 5.2 TEE Attestation 19

99 5.3 Technology Example: Intel Security Libraries for the Data Center (ISecL-DC)

100 20

101 5.4 Technology Summary 20

102 **6 Cloud Use Case Scenarios Leveraging Hardware-Based Security 21**

103 6.1 Visibility to Security Infrastructure 21

104 6.2 Workload Placement on Trusted Platforms 21

105	6.3 Asset Tagging and Trusted Location	23
106	6.4 Workload Confidentiality	24
107	6.5 Protecting Keys and Secrets.....	26
108	7 Next Steps	28
109	References.....	29
110	Appendix A – Acronyms.....	32
111		

112 **1 Introduction**

113 In today's cloud data centers and edge computing, there are three main forces that impact
114 security: (1) the introduction of billions of connected devices and increased adoption of the cloud
115 have significantly increased attack surfaces; (2) hacking has become industrialized with
116 sophisticated and evolving techniques to compromise data; and (3) solutions composed of
117 multiple technologies from different vendors result in a lack of coherent and consistent
118 implementations of security controls. Given these forces, the foundation for a data center or edge
119 computing security strategy should have a consolidated approach to comprehensively secure the
120 entire hardware platform on which workloads and data are executed and accessed.

121 In the scope of this document, the *hardware platform* is a server (e.g., application server, storage
122 server, virtualization server) in a data center or edge compute facility. The hardware platform
123 represents the first part of the layered security approach. Hardware security can provide a
124 stronger foundation than one offered by software or firmware, which can be modified with
125 relative ease. Existing security implementations can be enhanced by providing a base-layer,
126 immutable hardware module that chains software and firmware verifications from the hardware
127 all the way to the application space or specified security control. In that manner, existing security
128 mechanisms can be trusted even more to accomplish their security goals without compromise,
129 even when there is a lack of physical security or attacks originate from the software layer.

130 This white paper explains hardware-based security techniques and technologies that can improve
131 server platform security and data protection for cloud data centers and edge computing. The rest
132 of this white paper covers the following topics:

- 133 • Section 2 provides an overview of hardware platform security.
- 134 • Section 3 discusses the measurement and verification of platform integrity.
- 135 • Section 4 considers protecting data in use, also known as confidential computing.
- 136 • Section 5 examines remote attestation services, which can collate platform integrity
137 measurements to aid in integrity verification.
- 138 • Section 6 describes a number of cloud use case scenarios that take advantage of
139 hardware-based security.
- 140 • Section 7 states the next steps for this white paper and how others can contribute.

141 Although this document does not address other platforms like laptops, desktops, mobile devices,
142 or Internet of Things (IoT) devices, the practices in this white paper can be adapted to support
143 those platforms and their associated use cases.

144

145 **2 Hardware Platform Security Overview**

146 The data center threat landscape has evolved in recent years to encompass more advanced attack
147 surfaces with more persistent attack mechanisms. With increased attention being applied to high-
148 level software security, attackers are pushing lower in the platform stack, forcing security
149 administrators to address a variety of attacks that threaten the platform firmware and hardware.
150 These threats can result in:

- 151 • Unauthorized access to and potential extraction of sensitive platform or user data,
152 including direct physical access to dual in-line memory modules (DIMMs)
- 153 • Modification of platform firmware, such as that belonging to the Unified Extensible
154 Firmware Interface (UEFI)/Basic Input Output System (BIOS), Board Management
155 Controller (BMC), Manageability Engine (ME), Peripheral Component Interconnect
156 Express (PCIE) device, and various accelerator cards
- 157 • Supply chain interception through the physical replacement of firmware or hardware with
158 malicious versions
- 159 • Access to data or execution of code outside of regulated geopolitical or other boundaries
- 160 • Circumvention of software and/or firmware-based security mechanisms

161 For example, LoJax, discovered in August 2018, leveraged a UEFI loophole to continuously
162 reinstall a malicious piece of code at the firmware layer, thus remaining invisible to standard
163 kernel-based virus scans [1]. These attacks can be devastating to cloud environments because
164 they often require server-by-server rebuilds or replacements, which can take weeks. Although
165 still rare, these attacks are increasing as attackers become more sophisticated.

166 Regulated or sensitive workloads and data present additional security challenges for multi-tenant
167 clouds. While virtualization and containers significantly benefit efficiency, adaptability, and
168 scalability, these technologies consolidate workloads onto fewer physical platforms and
169 introduce the dynamic migration of workloads and data across platforms. Consequently, cloud
170 adoption results in a loss of visibility and control over the platforms that host virtualized
171 workloads and data, and introduces the usage of third-party infrastructure administrators. Cloud
172 providers often have data centers that span multiple geopolitical boundaries, subjecting workload
173 owners to complicated legal and regulatory compliance requirements from multiple countries.
174 Hybrid cloud architectures, in particular, utilize multiple infrastructure providers, each with their
175 own infrastructure configurations and management.

176 Without physical control over or visibility into platform configurations, traditional security best
177 practices and regulatory requirements become difficult or impossible to implement. With new
178 regulatory structures like the European General Data Protection Regulation (GDPR) introducing
179 high-stakes fines for noncompliance, having visibility and control over where data may be
180 accessed is more important than ever before. Top concerns among cloud security professionals
181 include the protection of workloads from general security risks, the loss or exposure of data in
182 the event of a data breach, and regulatory compliance.

183 Existing mitigations of threats against cloud servers are often rooted in firmware or software,
184 making them vulnerable to the same attack strategies. For example, if the firmware can be
185 successfully exploited, then the firmware-based security controls can most likely be
186 circumvented in the same fashion. Hardware-based security techniques can help mitigate these
187 threats by establishing and maintaining *platform trust*—an assurance in the integrity of the
188 underlying platform configuration, including hardware, firmware, and software. By providing
189 this assurance, security administrators can gain a level of visibility and control over where access
190 to sensitive workloads and data is permitted. Platform security technologies that establish
191 *platform trust* can provide notification or even self-correction of detected integrity failures.
192 Platform configurations can automatically be reverted back to a trusted state and give the
193 platform resilience against attack.

194 All security controls must have a *root of trust (RoT)*—a starting point that is implicitly trusted.
195 Hardware-based controls can provide an immutable foundation for establishing platform
196 integrity. Combining these functions with a means of producing verifiable evidence that these
197 integrity controls are in place and have been executed successfully is the basis of creating a
198 trusted platform. Minimizing the footprint of this RoT translates to reducing the number of
199 modules or technologies that must be implicitly trusted. This substantially reduces the attack
200 surface.

201 Platforms that secure their underlying firmware and configuration provide the opportunity for
202 trust to be extended higher in the software stack. Verified platform firmware can, in turn, verify
203 the operating system (OS) boot loader, which can then verify other software components all the
204 way up to the OS itself and the hypervisor or container runtime layers. The transitive trust
205 described here is consistent with the concept of the *chain of trust (CoT)*—a method where each
206 software module in a system boot process is required to measure the next module before
207 transitioning control.

208 Rooting platform integrity and trust in hardware security controls can strengthen and
209 complement the extension of the CoT into the dynamic software category. There, the CoT can be
210 extended even further to include data and workload protection. Hardware-based protections
211 through CoT technology mechanisms can form a layered security strategy to protect data and
212 workloads as they move to multi-tenant environments in a cloud data center or edge computing
213 facility.

214 In addition, there are other hardware platform security technologies that can protect data at rest,
215 in transit, and in use by providing hardware-accelerated disk encryption or encryption-based
216 memory isolation. By using hardware to perform these tasks, the attack surface is mitigated,
217 preventing direct access or modification of the required firmware. Isolating these encryption
218 mechanisms to specific hardware can allow performance to be addressed and enhanced
219 separately from other system processes as well.

220 **3 Platform Integrity Verification**

221 A key concept of trusted computing is verification of the underlying platform’s integrity.
222 Platform integrity is typically comprised of two parts:

- 223 • **Cryptographic measurement of software and firmware.** In this white paper, the term
224 *measurement* refers to calculating a cryptographic hash of a software or firmware
225 executable, configuration file, or other entity. If there is any change in an entity, a new
226 measurement will result in a different hash value than the original [2]. By measuring
227 software and firmware prior to execution, the integrity of the measured modules and
228 configurations can be validated before the platform launches or before data or workloads
229 are accessed. These measurements can also act as cryptographic proof for compliance
230 audits.
- 231 • **Firmware and configuration verification.** When firmware and configuration
232 measurements are made, local or remote attestations can be performed to verify if the
233 desired firmware is actually running and if the configurations are authorized. Attestation
234 can also serve as the foundation for further policy decisions that fulfill various cloud
235 security use case implementations. For instance, encryption keys can be released to client
236 workloads if a proof is performed that the platform server is trusted and in compliance
237 with policies.

238 In some cases, a third part is added to platform integrity:

- 239 • **Firmware and configuration recovery.** If the verification step fails (i.e., the attestations
240 do not match the expected measurements), the firmware and configuration can
241 automatically be recovered to a known good state, such as rolling back firmware to a
242 trusted version. The process by which these techniques are implemented affects the
243 overall strength of the assertion that the measured and verified components have not been
244 accidentally altered or maliciously tampered. Recovery technologies allow platforms to
245 maintain resiliency against firmware attacks and accidental provisioning mistakes.

246 There are many ways to measure platform integrity. Most technologies center around the
247 aforementioned concept of the CoT. In many cases, a hardware security module is used to store
248 measurement data to be attested at a later point in time. The rest of this section discusses
249 hardware security modules and various chain of trust technology implementations.

250 **3.1 Hardware Security Module (HSM)**

251 A *hardware security module (HSM)* is “a physical computing device that safeguards and
252 manages cryptographic keys and provides cryptographic processing” [3]. Cryptographic
253 operations such as encryption, decryption, and signature generation/verification are typically
254 hosted on the HSM device, and many implementations provide hardware-accelerated
255 mechanisms for cryptographic operations.

256 A *trusted platform module (TPM)* is a special type of HSM that can generate cryptographic keys
257 and protect small amounts of sensitive information, such as passwords, cryptographic keys, and
258 cryptographic hash measurements. [4] The TPM is a standalone device that can be integrated
259 with server platforms, client devices, and other products. One of the main use cases of a TPM is

260 to store digest measurements of platform firmware and configuration during the boot process.
261 Each firmware module is measured by generating a digest, which is then extended to a TPM
262 platform configuration register (PCR). Multiple firmware modules can be extended to the same
263 PCR, and the TPM specification provides guidelines for which firmware measurements are
264 encompassed by each PCR [5].

265 TPMs also host functionality to generate binding and signing keys that are unique per TPM and
266 stored within the TPM non-volatile random-access memory (NVRAM). The private portion of
267 this key pair is decrypted inside the TPM, making it only accessible by the TPM hardware or
268 firmware. This can create a unique relationship between the keys generated within a TPM and a
269 platform system, restricting private key operations to the platform firmware that has ownership
270 and access to the specified TPM. Binding keys are used for encryption/decryption of data, while
271 signing keys are used to generate/verify cryptographic signatures.

272 There are two versions of TPMs: 1.2 and 2.0. The 2.0 version supports additional security
273 features and algorithms [5]. TPMs also meet the National Institute of Standards and Technology
274 (NIST) Federal Information Processing Standard (FIPS) 140 validation criteria and support
275 NIST-approved cryptographic algorithms [6].

276 **3.2 The Chain of Trust (CoT)**

277 The *chain of trust (CoT)* is a method for maintaining valid trust boundaries by applying a
278 principle of transitive trust. Each firmware module in the system boot process is required to
279 measure the next module before transitioning control. Once a firmware module measurement is
280 made, it is recommended to immediately extend the measurement value to an HSM register for
281 attestation at a later point in time [5]. The CoT can be extended further into the application
282 domain, allowing for files, directories, devices, peripherals, etc. to be measured and attested.

283 Every CoT starts with an RoT module. It can be composed of different hardware and firmware
284 components. For several platform integrity technologies, the RoT core firmware module is
285 rooted in the central processing unit (CPU) microcode. However, not all technologies define
286 their RoTs in this manner [5]. The RoT is typically separated into components that verify and
287 measure. The core root of trust for verification (CRTV) is responsible for verifying the first
288 component before control is passed to it. The core root of trust for measurement (CRTM) is the
289 first component that is executed in the CoT and extends the first measurement to the TPM. The
290 CRTM can be divided into a static portion (SCRTM) and dynamic portion (DCRTM). The
291 SCRTM is composed of elements that measure firmware at system boot time, creating an
292 unchanging set of measurements that will remain consistent across reboots. The DRTM allows a
293 CoT to be established without rebooting the system, permitting the root of trust for measurement
294 to be reestablished dynamically.

295 An RoT that is built with hardware protections will be more difficult to change, while an RoT
296 that is built solely in firmware can easily be flashed and modified.

297 Various platform integrity technologies build their own CoTs. Some of these are discussed below
298 to illustrate the concept.

299 3.2.1 Technology Example: Intel Trusted Execution Technology (TXT)

300 Intel Trusted Execution Technology (TXT) in conjunction with a TPM provides a hardware RoT
301 available on Intel server and client platforms that enables “security capabilities such as measured
302 launch and protected execution” [7]. TXT utilizes *authenticated code modules (ACMs)* that
303 measure various pieces of the CoT during boot time and extend them to the platform TPM [2][7].
304 TXT’s ACMs are chipset-specific signed binaries that are called to perform functions required to
305 enable the TXT environment. An ACM is loaded into and executed from within the CPU cache
306 in an area referred to as the *authenticated code RAM (AC RAM)*. CPU microcode, which acts as
307 the *core root of trust for measurement (CRTM)*, authenticates the ACM by verifying its included
308 digital signature against a manufacturer public key with its digest hard-coded within the chipset.
309 The ACM code, loaded into protected memory inside the processor, performs various tests and
310 verifications of chipset and processor configurations.

311 The ACMs needed to initialize the TXT environment are the BIOS and the Secure Initialization
312 (SINIT) ACMs. Both are typically provided within the platform BIOS image. The SINIT ACM
313 can be provisioned on disk as well [2][8]. The BIOS ACM is responsible for measuring the
314 BIOS firmware to the TPM and performs additional BIOS-based security operations. The latest
315 version of TXT converged with Intel Boot Guard Technology labels this ACM as the Startup
316 ACM to differentiate it from the legacy BIOS ACM. The SINIT ACM is used to measure the
317 system software or operating system to the TPM, and it “initializes the platform so the OS can
318 enter the secure mode of operation” [8].

319 When the BIOS startup procedures have completed, control is transitioned to the OS loader. In a
320 TXT-enabled system, the OS loader is instructed to load a special module called Trusted Boot
321 before loading the first kernel module [8]. Trusted Boot (tboot) is an open-source, pre-
322 kernel/virtual machine manager (VMM) module that integrates with TXT to perform a measured
323 launch of an OS kernel/VMM. The tboot design typically has two parts: a preamble and the
324 trusted core. The tboot preamble is most commonly executed by the OS loader but can be loaded
325 at OS runtime. The tboot preamble is responsible for preparing SINIT input parameters and is
326 untrusted by default. It executes the processor instruction that passes control to the CPU
327 microcode. The microcode loads the SINIT into AC RAM, authenticates it, measures SINIT to
328 the TPM, and passes control to it. SINIT verifies the platform configuration and enforces any
329 present Launch Control Policies, measuring them and tboot trusted core to the TPM. The tboot
330 trusted core takes control and continues the CoT, measuring the OS kernel and additional
331 modules (like initrd) before passing control to the OS [9].

332 Intel TXT includes a policy engine feature that provides the capability to specify known good
333 platform configurations. These *Launch Control Policies (LCPs)* dictate which system software is
334 permitted to perform a secure launch. LCPs can enforce specific platform configurations and
335 tboot trusted core versions required to launch a system environment [8].

336 3.2.2 Technology Example: Intel Boot Guard

337 Intel Boot Guard provides a hardware RoT for authenticating the BIOS. An original equipment
338 manufacturer (OEM) enables Boot Guard authentication on the server manufacturing line by
339 permanently fusing a policy and OEM-owned public key into the silicon. When an Intel

340 processor identifies that Boot Guard has been enabled on the platform, it authenticates and
341 launches an ACM. The ACM loads the initial BIOS or Initial Boot Block (IBB) into the
342 processor cache, authenticates it using the fused OEM public key, and measures it into the TPM.

343 If the IBB authenticates properly, it verifies the remaining BIOS firmware, loads it into memory,
344 and transfers execution control. The IBB is restricted to this limited functionality, which allows it
345 to have a small enough size to fit in the on-die cache memory of Intel silicon. If the Boot Guard
346 authentication fails, the system is forced to shut down. When the Boot Guard execution
347 completes, the CoT can continue for other components by means of UEFI Secure Boot. TXT can
348 be used in conjunction with these technologies to provide a dynamic trusted launch of the OS
349 kernel and software.

350 Because Boot Guard is rooted in permanent silicon fuses and authenticates the initial BIOS from
351 the processor cache, it provides resistance from certain classes of physical attacks. Boot Guard
352 also uses fuses to provide permanent revocation of compromised ACMs, BIOS images, and input
353 polices.

354 **3.2.3 Technology Example: UEFI Secure Boot (SB)**

355 “UEFI Secure Boot (SB) is a verification mechanism for ensuring that code launched by a
356 computer’s UEFI firmware is trusted” [10]. SB prevents malware from taking “advantage of
357 several pre-boot attack points, including the system-embedded firmware itself, as well as the
358 interval between the firmware initiation and the loading of the operating system” [11].

359 The basic idea behind SB is to sign executables using a public-key cryptography scheme. The
360 public part of a *platform key* (PK) can be stored in the firmware for use as a root key. Additional
361 key exchange keys (KEKs) can also have their public portion stored in the firmware in what is
362 called the *signature database*. This database contains public keys that can be used to verify
363 different components that might be used by UEFI (e.g., drivers), as well as bootloaders and OSs
364 that are loaded from external sources (e.g., disks, USB devices, network). The signature database
365 can also contain *forbidden signatures*, which correspond to a revocation list of previously valid
366 keys. The signature database is meant to contain the current list of authorized and forbidden keys
367 as determined by the UEFI organization. The signature on an executable is verified against the
368 signature database before the executable can be launched, and any attempt to execute an
369 untrusted program will be prevented [10][11].

370 Before a PK is loaded into the firmware, UEFI is considered to be in *setup mode*, which allows
371 anyone to write a PK or KEK to the firmware. Writing the PK switches the firmware into *user*
372 *mode*. Once in user mode, PKs and KEKs can only be written if they are signed using the private
373 portion of the PK. Essentially, the PK is meant to authenticate the platform owner, while the
374 KEKs are used to authenticate other components of the distribution (distro), like OSs [11].

375 Shim is a simple software package that is designed to work as a first-stage bootloader on UEFI
376 systems. It is a common piece of code that is considered safe, well-understood, and audited so
377 that it can be trusted and signed using PKs. This means that firmware certificate authority (CA)
378 providers only have to worry about signing shim and not all of the other programs that vendors
379 might want to support [10]. Shim then becomes the RoT for all the other distro-provided UEFI

380 programs. It embeds a distro-specific CA key that is itself used to sign additional programs (e.g.,
381 Linux, GRUB, fwupdate). This allows for a clean delegation of trust; the distros are then
382 responsible for signing the rest of their packages. Ideally, shim will not need to be updated often,
383 which should reduce the workload on the central auditing and CA teams [10].

384 A key part of the shim design is to allow users to control their own systems. The distro CA key is
385 built into the shim binary itself, but there is also an extra database of keys that can be managed
386 by the user—the so-called *Machine Owner Key (MOK)*. Keys can be added and removed in the
387 MOK list by the user, entirely separate from the distro CA key. The mokutil utility can be used
388 to help manage the keys from Linux OS, but changes to the MOK keys may only be confirmed
389 directly from the console at boot time. This helps remove the risk of OS malware potentially
390 enrolling new keys and therefore bypassing SB [10].

391 On systems with a TPM chip enabled and supported by the system firmware, shim will extend
392 various PCRs with the digests of the targets it is loading [12]. Certificate hashes are also
393 extended to the TPM, including system, vendor, MOK, and shim blacklisted and whitelisted
394 certificate digests.

395 **3.2.4 Technology Example: Intel Platform Firmware Resilience (PFR)**

396 Intel Platform Firmware Resilience (PFR) technology is a platform-level solution that creates an
397 open platform RoT based on a programmable logic device. It is designed to provide firmware
398 resiliency (in accordance with NIST Special Publication [SP] 800-193 [13]) and comprehensive
399 protection for various platform firmware components, including BIOS, Server Platform Services
400 Firmware (SPS FW), and Board Management Controllers (BMCs). PFR provides the platform
401 owner with a minimal trusted compute base (TCB) under full platform-owner control. This TCB
402 provides cryptographic authentication and automatic recovery of platform firmware to help
403 guarantee correct platform operation and to return to a known good state in case of a malicious
404 attack or an operator error such as a failed update.

405 NIST SP 800-193 [13] outlines three guiding principles to support the resiliency of platforms
406 against potentially destructive attacks:

- 407 • **Protection:** Mechanisms for ensuring that platform firmware code and critical data
408 remain in a state of integrity and are protected from corruption, such as the process for
409 ensuring the authenticity and integrity of firmware updates
- 410 • **Detection:** Mechanisms for detecting when platform firmware code and critical data have
411 been corrupted
- 412 • **Recovery:** Mechanisms for restoring platform firmware code and critical data to a state
413 of integrity in the event that any such firmware code or critical data are detected to have
414 been corrupted or when forced to recover through an authorized mechanism. Recovery is
415 limited to the ability to recover firmware code and critical data.

416 In addition, NIST SP 800-193 [13] provides guidance on meeting those requirements via three
417 main functions of a Platform Root of Trust:

- 418 • **Root of Trust for Update (RTU)**, which is responsible for authenticating firmware
419 updates and critical data changes to support platform protection capabilities; this includes
420 signature verification of firmware updates as well as rollback protections during update.
- 421 • **Root of Trust for Detection (RTD)**, which is responsible for firmware and critical data
422 corruption detection capabilities.
- 423 • **Root of Trust for Recovery (RTRec)**, which is responsible for recovery of firmware and
424 critical data when corruption is detected or when instructed by an administrator.

425 PFR is designed to support NIST guidelines and create a resilient platform that is able to self-
426 recover upon detection of attack or firmware corruption. This includes verification of all
427 platform firmware and configuration at platform power-on time, active protection of platform
428 non-volatile memory at runtime, and active protection of the Serial Peripheral Interface (SPI
429 flash) and System Management Bus (SMBus). PFR functionality also incorporates monitoring
430 the platform component's boot progress and providing automatic firmware recovery to a known
431 good state upon detection of firmware or configuration corruption. PFR achieves this goal by
432 utilizing a Field-Programmable Gate Array (FPGA) to establish an RoT.

433 PFR technology defines a special pre-boot mode (T-1) where only the PFR FPGA is active. Intel
434 Xeon processors and other devices that could potentially interfere with the boot process, such as
435 the Platform Controller Hub (PCH)/Manageability Engine (ME) and BMC, are not powered.
436 Boot critical firmware, like the BIOS, ME, and BMC, are cryptographically verified during T-1
437 mode. In case of corruption, a recovery event is triggered, and the corrupted firmware in the
438 active regions of the SPI flash is erased and restored with a known-good recovery copy. Once
439 successful, the system proceeds to boot in a normal mode, leveraging Boot Guard for static RoT
440 coverage.

441 The PFR FPGA RoT leverages a key hierarchy to authenticate data structures residing in SPI
442 flash. The key hierarchy is based on a provisioned Root Key (RK) stored in the NVRAM of the
443 FPGA RoT and a Code Signing Key (CSK) structure, which is endorsed by the RK, stored in the
444 SPI flash, and used for the signing of lower-level data structures. The PFR FPGA uses this CSK
445 to verify the digital signature of the Platform Firmware Manifest (PFM), which describes the
446 expected measurements of the platform firmware. The PFR FPGA RoT verifies those
447 measurements before allowing the system to boot. When a recovery is needed, either because
448 measurements do not match the expected value or because a hang is detected during system
449 bootup, the PFR FPGA RoT uses a recovery image to recover the firmware. The recovery image
450 and any update images are stored in a compressed capsule format and verified using a digital
451 signature.

452 Each platform firmware storage is divided into three major sections: Active, Recovery, and
453 Staging. The Recovery regions, as well as the static parts of the Active regions, are write-
454 protected from other platform components by the PFR FPGA RoT. The Staging region is open to
455 the other platform components for writing in order to provide an area to place digitally signed
456 and compressed update capsules, which are then verified by the PFR FPGA RoT before being
457 committed to the Active or Recovery regions. The Recovery copy can be updated in T-1 mode
458 once the PFR FPGA has verified the digital signature of the update capsule and confirmed that
459 the recovery image candidate is bootable.

460 **3.3 Supply Chain Protection**

461 Organizations are increasingly at risk of supply chain compromise, whether intentional or
462 unintentional. Managing cyber supply chain risks requires, in part, ensuring the integrity, quality,
463 and resilience of the supply chain, its products, and its services. Cyber supply chain risks may
464 include counterfeiting, unauthorized production, tampering, theft, and insertion of malicious or
465 otherwise unexpected software and hardware, as well as poor manufacturing and development
466 practices in the cyber supply chain [14][15].

467 Special technologies have been developed to help ascertain the authenticity and integrity of
468 platform hardware, including its firmware and configuration. These technologies help ensure that
469 platforms are not tampered with or altered from the time that they are assembled at the
470 manufacturer site to the time that they arrive at a consumer data center ready for installation.
471 Verification of these platform attributes is one aspect of securing the supply chain. Some
472 technologies include an additional feature for locking the boot process or access to these
473 platforms until a secret is provided that only the consumer and manufacturer know.

474 **3.3.1 Technology Example: Intel Transparent Supply Chain (TSC)**

475 “Intel Transparent Supply Chain (TSC) is a set of policies and procedures implemented at ODM
476 factories that enable end-users to validate where and when every component of a platform was
477 manufactured” [16]. “Intel TSC tools allow platform manufacturers to bind platform information
478 and measurements using [a TPM]. This allows customers to gain traceability and accountability
479 for platforms with component-level reporting” [17].

480 Intel TSC provides the following key features [16]:

- 481 • Digitally signed statement of conformance for every platform
- 482 • Platform certificates linked to a discrete TPM, providing system-level traceability
- 483 • Component level traceability via a direct platform data file that contains integrated
484 components, including a processor, storage, memory, and add-in cards
- 485 • Auto Verify tool that compares the snapshot of the direct platform data taken during
486 manufacturing with a snapshot of the platform components taken at first boot
- 487 • Firmware load verification
- 488 • Conformity with Defense Federal Acquisition Regulation Supplement (DFARS)
489 246.870-2 [18]

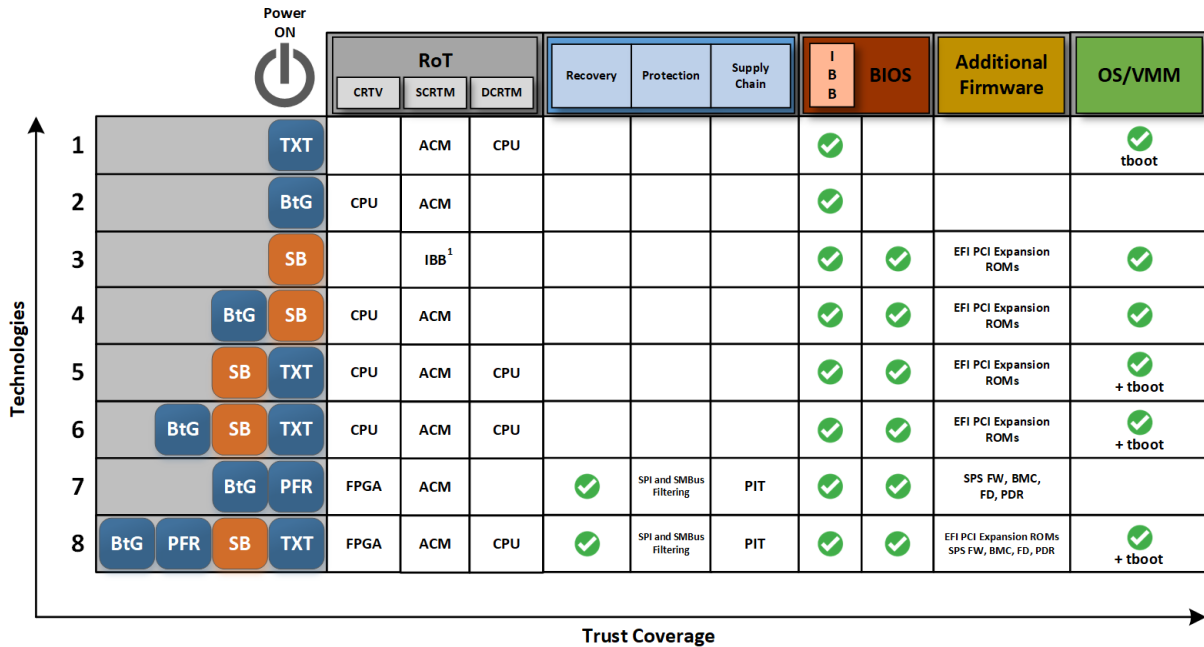
490 **3.3.2 Technology Example: PFR with Protection in Transit (PIT)**

491 In addition to the platform protection, detection, and recovery features, PFR also offers
492 protection in transit (PIT) or supply chain protection. Platform lockdown requires that a
493 password be present in the PFR FPGA as well as a radio frequency (RF) component. The
494 password is removed before platform shipment and must be replaced before the platform will be
495 allowed to power up. With platform firmware sealing, the PFR FPGA computes hashes of
496 platform firmware in the PCH and BMC attached flash devices, including static and dynamic
497 regions, and stores them in a NVRAM space before shipment. Upon delivery, the PFR FPGA

498 will recompute the hashes and report any mismatches to ensure that the firmware has not been
 499 tampered with during system transit.

500 **3.4 Technology Example Summary**

501 There are several technologies that provide different levels of platform integrity and trust.
 502 Individual technologies do not provide a complete CoT. When used in combination, they can
 503 provide comprehensive coverage all the way up to the OS and VMM layer. Figure 1 outlines the
 504 firmware and software coverage of each existing CoT technology example.



¹ IBB is meant to describe the portion of BIOS which performs the first measurement

505
 506 **Figure 1: Firmware and Software Coverage of Existing Chain of Trust Technologies**

507 Figure 1 identifies the components of each technology that make up the RoT in their own
 508 respective chains and also shows a rough outline of the firmware and software coverage of each
 509 technology.

510 Because many technologies are available, it can be difficult to decide on the right combination
 511 for deployment. Figure 1 illustrates the possible combinations of technologies that extend
 512 measurements to a TPM for platform integrity attestation. Note that each combination includes at
 513 least one hardware technology to ensure an RoT implementation. A complementary option for
 514 extending the CoT up through the OS can also be provided. Including only the hardware
 515 technologies would break the CoT by supplying integrity measurements for only pre-OS
 516 firmware. Using only UEFI Secure Boot will use firmware as the RoT that does not have
 517 hardware security protections and is much more susceptible to attack. By enabling both parts, the
 518 CoT can be extended from a hardware RoT into the OS and beyond.

519 These combinations will help ensure that appropriate measurements are extended to a TPM for
 520 integrity attestation and can prevent a server from booting if specific security modules are

521 compromised. The attestation mechanisms provided by these technologies give cryptographic
522 proof of the integrity of measured components, which can be used to provide visibility into
523 platform security configurations and prove integrity. Note the combination of UEFI SB with
524 TXT in Figure 1. This combination provides the UEFI SB signature verification capability on top
525 of the tboot integrity measurement in the OS/VMM layer.

526 In addition to attestation, PFR provides both additional verification of platform firmware and
527 adds automatic recovery of compromised firmware to known good versions. PFR works with any
528 combination of CoT technologies, providing a defense and resilience against firmware attack
529 vectors. Combining a hardware-based firmware resilience technology like PFR with a hardware-
530 based CoT configuration is part of a layered security strategy.

531 4 Data Protection and Confidential Computing

532 With the increase in adoption of consumer-based cloud services, virtualization has become a
533 necessity in cloud data center infrastructure. Virtualization simulates the hardware that multiple
534 cloud workloads run on top of. Each workload is isolated from others so that it has access to only
535 its own resources, and each workload can be completely encapsulated for portability [19] [20].
536 Traditional virtual machines (VMs) have an isolated kernel space running all aspects of a
537 workload alongside the kernel. Today, the virtualized environment has been extended to include
538 containers and full-featured workload orchestration engines. Containers offer application
539 portability by sharing an underlying kernel, which drastically reduces workload-consumed
540 resources and increases performance.

541 While containers can provide a level of convenience, vulnerabilities in the kernel space and
542 shared layers can be susceptible to widespread exploitation, making security for the underlying
543 platform even more important. With the need for additional protection in the virtualized
544 workspace, an emphasis has been placed on encrypting data both at rest and while in use. *At-rest*
545 encryption provides protection for data on disk. This typically refers to an unmounted data store
546 and protects against threats such as the physical removal of a disk drive. Protecting and securing
547 cloud data while *in use*, also referred to as *confidential computing*, utilizes hardware-enabled
548 features to isolate and process encrypted data in memory so that the data is at less risk of
549 exposure and compromise from concurrent workloads or the underlying system and platform
550 [21]. This section describes technologies that can be leveraged for providing confidential
551 computing for cloud and edge.

552 A *trusted execution environment (TEE)* is an area or enclave protected by a system processor.
553 Sensitive secrets like cryptographic keys, authentication strings, or data with intellectual property
554 and privacy concerns can be preserved within a TEE, and operations involving these secrets can
555 be performed within the TEE, thereby eliminating the need to extract the secrets outside of the
556 TEE. A TEE also helps ensure that operations performed within it and the associated data cannot
557 be viewed from outside, not even by privileged software or debuggers. Communication with the
558 TEE is designed to only be possible through designated interfaces, and it is the responsibility of
559 the TEE designer/developer to define these interfaces appropriately. A good TEE interface limits
560 access to the bare minimum required to perform the task.

561 4.1 Memory Isolation

562 There are many technologies that provide data protection via encryption. Most of these solutions
563 focus on protecting the respective data while at rest and do not cover the fact that the data is
564 decrypted and vulnerable while in use. Applications running in memory share the same platform
565 hardware and can be susceptible to attacks either from other workloads running on the same
566 hardware or from compromised cloud administrators. There is a strong desire to secure
567 intellectual property and ensure that private data is encrypted and not accessible at any point in
568 time, particularly in cloud data centers and edge computing facilities. Various hardware
569 technologies have been developed to encrypt content running in platform memory.

570 4.2 Application Isolation

571 Application isolation utilizes a TEE to help protect the memory reserved for an individual
572 application. The trust boundary associated with the application is restricted to only the CPU.
573 Future generations of these techniques will allow entire applications to be isolated in their own
574 enclaves rather than only protecting specific operations or memory. By using separate
575 application enclaves with unique per-application keys, sensitive applications can be protected
576 against data exposure, even to malicious insiders with access to the underlying platform.
577 Implementations of application isolation will typically involve developer integration of a toolkit
578 within the application layer, and it is the developer's responsibility to ensure secure TEE design.

579 The section below presents an application isolation example using a TEE.

580 4.2.1 Technology Example: Intel Software Guard Extensions (SGX)

581 Intel Software Guard Extensions (SGX) is a set of instructions that increases the security of
582 application code and data. Developers can partition security-sensitive code and data into an *SGX*
583 *enclave*, which is executed in a CPU protected region. The developer creates and runs SGX
584 enclaves on server platforms where only the CPU is trusted to provide attestations and protected
585 execution environments for enclave code and data. SGX also provides an enclave remote
586 attestation mechanism. This mechanism allows a remote provider to verify the following [22]:

- 587 1. The enclave is running on a real Intel processor inside an SGX enclave.
- 588 2. The platform is running at the latest security level (also referred to as the TCB version).
- 589 3. The enclave's identity is as claimed.
- 590 4. The enclave has not been tampered with.

591 Once all of this is verified, the remote attester can then provision secrets into the enclave. SGX
592 enclave usage is reserved for Ring-3 applications and cannot be used by an OS or BIOS
593 driver/module.

594 SGX removes the privileged software (e.g., OS, VMM, SMM, devices) and unprivileged
595 software (e.g., Ring-3 applications, VMs, containers) from the trust boundary of the code
596 running inside the enclave, enhancing security of sensitive application code and data. An SGX
597 enclave trusts the CPU for execution and memory protections. SGX encrypts memory to protect
598 against memory bus snooping and cold boot attacks for enclave code and data in host DRAM.
599 SGX includes instruction set architecture (ISA) instructions that can be used to handle Enclave
600 Page Cache (EPC) page management for creating and initializing enclaves.

601 SGX relies on the system UEFI BIOS and OS for initial provisioning, resource allocation, and
602 management. However, once an SGX enclave starts execution, it is running in a
603 cryptographically isolated environment separate from the OS and BIOS.

604 SGX can allow any application (whole or part of) to run inside an enclave and puts application
605 developers in control of their own application security. However, it is recommended that
606 developers keep the SGX code base small, validate the entire system (including software side
607 channel resistance), and follow other secure software development guidelines.

608 SGX enclaves can be used for applications ranging from protecting private keys and managing
609 security credentials to providing security services. In addition, industry security standards, like
610 European Telecommunications Standards Institute (ETSI) Network Functions Virtualization
611 (NFV) Security (ETSI NFV SEC) [23], have defined and published requirements for Hardware
612 Mediated Execution Enclaves (HMEEs) for the purposes of NFV, 5G, and edge security. SGX is
613 an HMEE.

614 **4.3 VM Isolation**

615 As new memory and execution isolation technologies become available, it is more feasible to
616 isolate entire VMs. VMs already enjoy a degree of isolation due to technologies like hardware-
617 assisted virtualization, but the memory of each VM remains in the clear. Existing memory
618 isolation technologies require implicit trust of the VMM. New isolation technologies in future
619 platform generations will remove the VMM from the trust boundary and allow full encryption of
620 VM memory with per-VM unique keys, protecting the VMs from not only malicious software
621 running on the hypervisor host but also rogue firmware.

622 VM isolation can be used to help protect workloads in multi-tenant environments like public and
623 hybrid clouds. Isolating entire VMs translates to protection against malicious insiders at the
624 cloud provider, or malware exposure and data leakage to other tenants with workloads running
625 on the same platform. Many modern cloud deployments use VMs as container worker nodes.
626 This provides a highly consistent and scalable way to deploy containers regardless of the
627 underlying physical platforms. With full VM isolation, the virtual workers hosting container
628 workloads can be effectively isolated without impacting the benefits of abstracting the container
629 from the underlying platform.

630 **4.4 Cryptographic Acceleration**

631 Encryption is quickly becoming more widespread in data center applications as industry adopts
632 more standards and guidelines regarding the sensitivity of consumer data and intellectual
633 property. Because cryptographic operations can drain system performance and consume large
634 amounts of compute resources, the industry has adopted specialized hardware interfaces called
635 *cryptographic accelerators*, which offload cryptographic tasks from the main processing unit
636 onto a separate coprocessor chip. Cryptographic accelerators often come in the form of pluggable
637 peripheral adapter cards.

638 **4.4.1 Technology Example: Intel QuickAssist Technology (QAT) with Intel Key** 639 **Protection Technology (KPT)**

640 Intel QuickAssist Technology (QAT) is a high-performance hardware accelerator for performing
641 cryptographic, security, and compression operations. Applications like VMs, containers, and
642 Function as a Service (FaaS) call Intel QAT using industry-standard OpenSSL, Transport Layer
643 Security (TLS), and Internet Protocol Security (IPsec) interfaces to offload symmetric and
644 asymmetric cryptographic operations. Cloud, multi-tenancy, NFV, edge, and 5G infrastructures
645 and applications are best suited for QAT for all types of workloads, including software-defined
646 networks (SDNs), content delivery networks (CDNs), media, and storage [24].

647 Intel Key Protection Technology (KPT) helps enable customers to secure their keys to be used
648 with QAT through a bring-your-own-key (BYOK) paradigm. KPT allows customers to deliver
649 their own cryptographic keys to the QAT device in the target platform where their workload is
650 running. KPT-protected keys are never in the clear in host DRAM or in transit. The customers
651 encrypt their workload key (e.g., RSA private key for Nginx) using KPT inside their HSMs. This
652 encrypted workload key is delivered to the target QAT platform, where it is decrypted
653 immediately prior to use. KPT provides key protection at rest, in transit, and while in use [25].

654 **4.5 Technology Example Summary**

655 Cloud infrastructure creates improvements in the efficiency, agility, and scalability of data center
656 workloads by abstracting hardware from the application layer. This introduces new security
657 concerns as workloads become multi-tenant, attack surfaces become shared, and infrastructure
658 administrators from the cloud operator gain access to underlying platforms. Isolation techniques
659 provide answers to these concerns by adding protection to VMs, applications, and data during
660 execution, and they represent a crucial layer of a layered security approach for data center
661 security architecture.

662 Various isolation techniques exist and can be leveraged for different security needs. Full memory
663 isolation defends a platform against physical memory extraction techniques, while the same
664 technology extended with multiple keys allows individual VMs or platform tenants to have
665 uniquely encrypted memory. Future generations of these technologies will allow full memory
666 isolation of VMs, protecting them against malicious infrastructure insiders, multi-tenant
667 malware, and more. Application isolation techniques allow individual applications to create
668 isolated enclaves that require implicit trust in the platform CPU and nothing else and that have
669 the ability to provide proof of the enclave to other applications before data is sent.

670 **5 Remote Attestation Services**

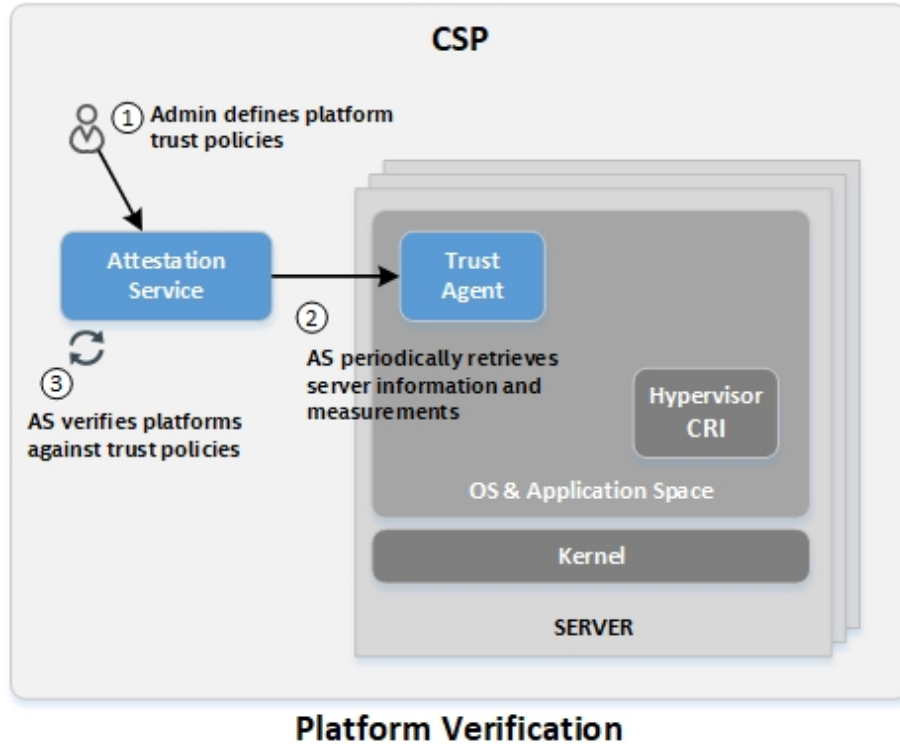
671 Measuring a server's firmware/configuration and extending these measurements to a hardware
672 interface can help keep track of which firmware is running on a platform. Some platform
673 integrity technologies can even perform local attestation and enforcement of firmware and
674 configuration on a server. However, data centers are usually made up of thousands of servers,
675 and keeping track of them and their respective firmware is an overwhelming task for an operator.
676 A remote service can address this by collating server information and measurement details.
677 Cryptographic signatures can be used to ensure the integrity of transferred measurement data.
678 Furthermore, the remote service can be used to define whitelist policies, specifying which
679 firmware versions and event measurements are acceptable for servers in a particular data center
680 environment. This service would verify or attest each server's collected data against these
681 policies, feeding the results into a policy orchestrator to report, alert, or enforce rules based on
682 the events.

683 A remote attestation service can provide additional benefits besides verifying server firmware.
684 Specifying whitelist policies for specific firmware versions can allow data center administrators
685 to easily invalidate old versions and roll out new upgrades. In some cases, certain hardware
686 technologies and associated capabilities on platforms can be discoverable by their specific event
687 log measurements recorded in an HSM. The information tracked in this remote attestation
688 service can even be exposed through the data center administration layer directly to the
689 enterprise user. This would give endpoint consumers hardware visibility and the ability to
690 specify firmware requirements or require platform features for the hardware on which their
691 services are running.

692 The key advantage to remote attestation is the enforcement of compliance across all hardware
693 systems in a data center. The ability to verify against a collective whitelist as opposed to a local
694 system enforcing a supply chain policy provides operators more flexibility and control in a
695 cryptographically secured manner. These enforcement mechanisms can even be combined to
696 provide stronger security policies.

697 **5.1 Platform Attestation**

698 Figure 2 shows a remote attestation service (AS) collecting platform configurations and integrity
699 measurements from data center servers at a cloud service provider (CSP) via a trust agent service
700 running on the platform servers. A cloud operator is responsible for defining whitelisted trust
701 policies. These policies should include information and expected measurements for desired
702 platform CoT technologies. The collected host data is compared and verified against the policies,
703 and a report is generated to record the relevant trust information in the AS database.



704

705

Figure 2: Notional Example of Remote Attestation Service

706 Platform attestation can be extended to include application integrity or the measurement and
707 verification of the hypervisor container runtime interface (CRI) and applications installed on
708 bare-metal servers. During boot time, an application agent on the server can measure operator-
709 specified files and directories that pertain to particular applications. A whitelist trust policy can
710 be defined to include these expected measurements, and this policy can be included in the overall
711 trust assessment of the platform in the remote AS. By extending measurements to a platform
712 TPM, applications running on the bare-metal server can be added to the CoT. The components of
713 the trust agent and application agent can be added to the policy and measured alongside other
714 applications to ensure that the core feature elements are not tampered with. For example, a
715 typical Linux implementation of the application agent could run inside `initrd`, and measurements
716 made on the filesystem could be extended to the platform TPM.

717 An additional feature commonly associated with platform trust is the concept of *asset tagging*.
718 *Asset tags* are simple key value attributes that are associated with a platform like location,
719 company name, division, or department. These key value attributes are tracked and recorded in a
720 central remote service, such as the AS, and can be provisioned directly to a server through the
721 trust agent. The trust agent can then secure these attribute associations with the host platform by
722 writing hash measurement data for the asset tag information to a hardware security chip, such as
723 the platform TPM NVRAM. Measurement data is then retrieved by the AS and included in the
724 platform trust report evaluation.

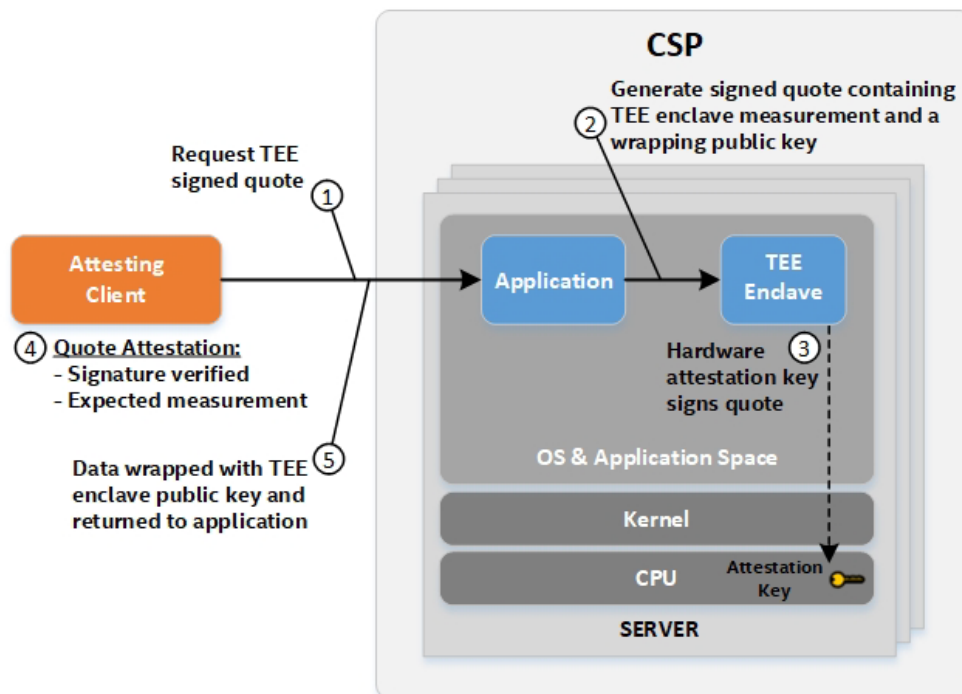
725 5.2 TEE Attestation

726 There are instances when the high assurance that the output of the processing in a TEE can be
727 trusted should be extended to an external attesting client. This is achieved thanks to a TEE
728 attestation flow. *TEE attestation* involves the generation of a verifiable cryptographic quote of
729 the enclave by the TEE. The quote is then sent to the attesting client, which can validate the
730 signature of the quote. If the signature is valid, the attesting client concludes that the remote code
731 is running in a genuine TEE enclave.

732 A quote usually contains the measurement of the TEE enclave, as well as data related to the
733 authenticity of the TEE and the compliant version of it. The measurement is a digest of the
734 content of the enclave (i.e., code, data, stack, and heap) and other information. The measurement
735 obtained at build time is typically known to the attesting client and is compared against a
736 measurement contained in the quote that is actively taken during runtime. This allows the
737 attesting client to determine that the remote code has not been tampered with. A quote may also
738 contain the enclave's developer signature and platform TCB information. The authenticity and
739 version of the TEE are verified against TEE provider certificates that are accessible to the tenant
740 or attesting client.

741 The quote may also contain the public key part of an enclave public/private key pair. The public
742 key allows the attesting client to wrap secrets that it wants to send to the enclave. This capability
743 allows the attesting client to provision secrets directly to the TEE enclave without needing to
744 trust any other software running on the server.

745 Figure 3 shows an example TEE attestation flow.



746

747

Figure 3: Notional Example of TEE Attestation Flow

748 **5.3 Technology Example: Intel Security Libraries for the Data Center (ISecL-DC)**

749 Intel Security Libraries for the Data Center (ISecL-DC) is an open-source remote attestation
750 implementation of a set of building blocks that utilize Intel Security features to discover, attest,
751 and enable critical foundation security and confidential computing use-cases. This middleware
752 technology provides a consistent set of APIs for easy integration with cloud management
753 software and security monitoring and enforcement tools. ISecL-DC applies the remote attestation
754 fundamentals described in this section and standard specifications to maintain a platform data
755 collection service and an efficient verification engine to perform comprehensive trust
756 evaluations. These trust evaluations can be used to govern different trust and security policies
757 applied to any given workload, as referenced in the workload scheduling use case in Section 6.2.
758 In future generations, the product will be extended to include TEE attestation to provide
759 assurance and validity of the TEE to enable confidential computing [26].

760 **5.4 Technology Summary**

761 Platform attestation provides auditable foundational reports for server firmware and software
762 integrity and can be extended to include the location of other asset tag information stored in a
763 TPM, as well as integrity verification for applications installed on the server. These reports
764 provide visibility into platform security configurations and can be used to control access to data
765 and workloads. Platform attestation is performed on a per-server basis and typically consumed
766 by cloud orchestration or a wide variety of infrastructure management platforms.

767 TEE attestation provides a mechanism by which a user or application can validate that a genuine
768 TEE enclave with an acceptable TCB is actually being used before releasing secrets or code to
769 the TEE. Formation of a TEE enclave is performed at the application level, and TEE attestations
770 are typically consumed by a user or application requiring evidence of enclave security before
771 passing secrets.

772 These different attestation techniques serve complementary purposes in a cloud deployment in
773 the data center or at the edge computing facility.

774 **6 Cloud Use Case Scenarios Leveraging Hardware-Based Security**

775 This section describes a number of cloud use case scenarios that take advantage of the hardware-
776 based security capability and trust attestation capability integrated with the operator orchestration
777 tool to support various security and compliance objectives.

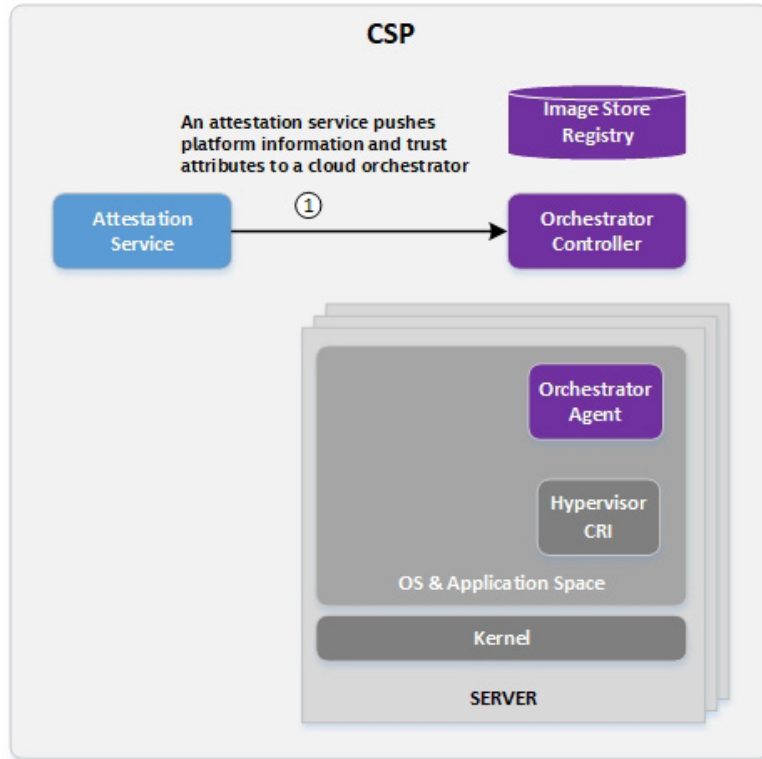
778 **6.1 Visibility to Security Infrastructure**

779 A typical attestation includes validation of the integrity of platform firmware measurements.
780 These measurements are unique to a specific BIOS/UEFI version, meaning that the attestation
781 report provides visibility into the specific firmware version currently in use, in addition to the
782 integrity of that firmware. Attestation can also include hardware configuration and feature
783 support information, both by attesting feature support directly and by resulting in different
784 measurements based on which platform integrity technologies are used.

785 Cryptographically verifiable reports of platform integrity and security configuration details (e.g.,
786 BIOS/UEFI versions, location information, application versions) are extremely useful for
787 compliance auditing. These attestation reports for the physical platform can be paired with
788 workload launch or key release policies, providing traceability to confirm that data and
789 workloads have only been accessed on compliant hardware in compliant configurations with
790 required security technologies enabled.

791 **6.2 Workload Placement on Trusted Platforms**

792 Platform information and verified firmware/configuration measurements retained within an
793 attestation service can be used for policy enforcement in countless use cases. One example is
794 orchestration scheduling. Cloud orchestrators, such as Kubernetes and OpenStack, provide the
795 ability to label server nodes in their database with key value attributes. The attestation service
796 can publish trust and informational attributes to orchestrator databases for use in workload
797 scheduling decisions. Figure 4 illustrates this.



798

799

Figure 4: Notional Example of Orchestrator Platform Labeling

800 In OpenStack, this can be accomplished by labeling nodes using custom traits. Workload images
801 can be uploaded to an image store containing metadata that specifies required trait values to be
802 associated with the node that is selected by the scheduling engine. In Kubernetes, nodes can be
803 labeled in etcd via node selector or node affinity. Custom resource definitions (CRDs) can be
804 written and plugged into Kubernetes to receive label values from the attestation service and
805 associate them with nodes in the etcd. When a deployment or container is launched, node
806 selector or node affinity attributes can be included in the configuration yaml to instruct
807 Kubernetes to only select nodes that have the specified labels. Other orchestrator engines and
808 flavors can be modified to accommodate a similar use case. Figure 5 illustrates how an
809 orchestrator can be configured to only launch workloads on trusted platforms or platforms with
810 specified asset tag attributes.

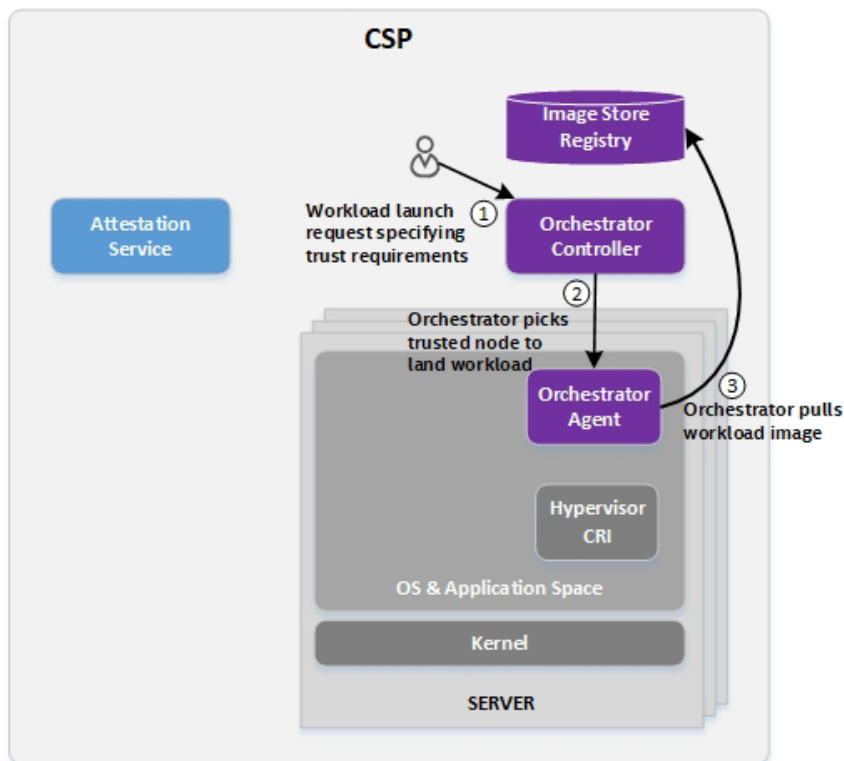


Figure 5: Notional Example of Orchestrator Scheduling

811

812

813 6.3 Asset Tagging and Trusted Location

814 Trusted geolocation is a specific implementation of the aforementioned trusted asset tag feature
 815 used with platform attestation. Key attribute values specifying location information are used as
 816 asset tags and provisioned to server hardware, such as the TPM. In this way, location information
 817 can be included in platform attestation reports and therefore consumed by cloud orchestrators,
 818 infrastructure management applications, policy engines, and other entities [27]. Orchestration
 819 using asset tags can be used to segregate workloads and data access in a wide variety of
 820 scenarios. Geolocation can be an important attribute to consider with hybrid cloud environments
 821 subject to regulatory controls like, for example, GDPR. Violating these constraints by allowing
 822 access to data outside of specific geopolitical boundaries can trigger substantial penalties.

823 In addition to location, the same principle can apply to other sorts of tag information. For
 824 example, some servers might be tagged as appropriate for storing health information subject to
 825 Health Insurance Portability and Accountability Act (HIPAA) compliance. Data and workloads
 826 requiring this level of compliance should only be accessed on platforms configured to meet those
 827 compliance requirements. Other servers may be used to store or process information and
 828 workloads not subject to HIPAA requirements. Asset tags can be used to flag which servers are
 829 appropriate for which workloads beyond a simple statement of the integrity of those platforms.
 830 The attestation mechanisms help ensure that the asset tag information is genuine, preventing easy
 831 subversion.


832 Outside of specific regulatory requirements, an organization may wish to segregate workloads by
 833 department. For example, human resources and finance information could be restricted to

834 platforms with different security profiles, and big data workloads could be required to run on
835 platforms tagged for performance capabilities. For cloud orchestration platforms that do not
836 natively support discovery or scheduling of workloads based on specific platform features, asset
837 tags can provide a mechanism for seamlessly adding such a capability. For example, workloads
838 that require Intel SGX can be orchestrated to only run on platforms that support the SGX
839 platform feature, even if the cloud platform does not natively discover support for SGX. The
840 open-ended user-configurable asset tag functionality allows virtually any level of subdivision of
841 resources for business, security, or regulatory needs.

842 **6.4 Workload Confidentiality**

843 Consumers who place their workloads in the cloud or the edge are typically forced to accept that
844 their workloads are secured by their service providers without insight or knowledge as to what
845 security mechanisms are in place. The ability for end users to encrypt their workload images can
846 provide at-rest cryptographic isolation to protect consumer data and intellectual property. Key
847 control is integral to the workload encryption process. While it is preferable to transition key
848 storage, management, and ownership to the endpoint consumer, an appropriate key release policy
849 must be defined that includes a guarantee from the service provider that the utilized hardware
850 platform and firmware are secure and uncompromised.

851 There are several key management solutions (KMSs) in production that provide services to
852 create and store keys. Many of these are compliant with the industry-standardized Key
853 Management Interoperability Protocol (KMIP) and can be deployed within consumer enterprises.
854 The concept is to provide a thin layer on top of the KMS called a *key broker*, as illustrated in
855 Figure 6, that applies and evaluates policies to requests that come into the KMS. Supported
856 requests to the key broker include key creation, key release policy association, and key request
857 by evaluating associated policies. The key release policy can be any arbitrary set of rules that
858 must be fulfilled before a key is released. The policy for key release is open-ended and meant to
859 be easily extendible, but for the purpose of this discussion, a policy associated with platform
860 trust is assumed.

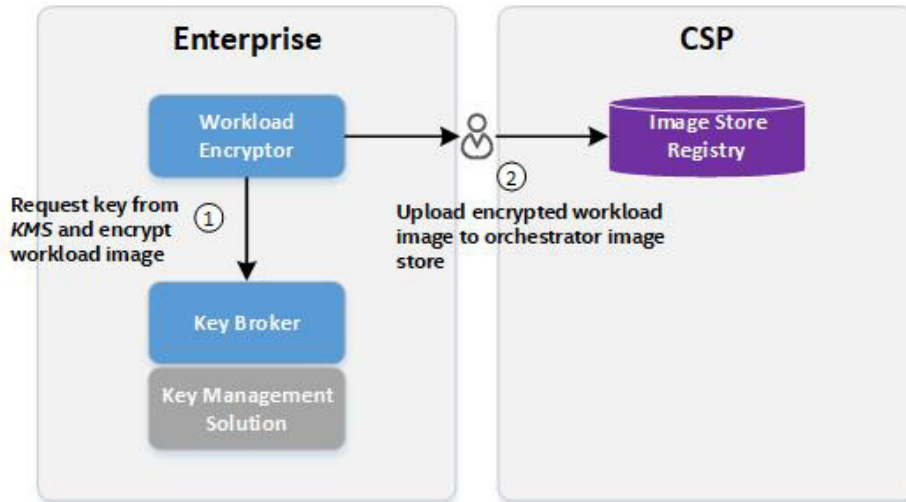


861

862

Figure 6: Notional Example of Key Brokerage

863 Once the key policy has been determined, a KMS-created and managed key can be used to
864 encrypt a workload image, as shown in Figure 7. The enterprise user may then upload the
865 encrypted image to a CSP orchestrator image store or registry.



866

867

Figure 7: Notional Example of Workload Image Encryption

868 The key retrieval and decryption process is the most complex piece of the workload
869 confidentiality story, as Figure 8 shows. It relies on a secure key transfer between the enterprise
870 and CSP with an appropriate key release policy managed by the key broker. The policy for key
871 release discussed here is based on platform trust and the valid proof thereof. The policy can also
872 dictate a requirement to wrap the key using a public wrapping key, with the private portion of the
873 wrapping key only known to the hardware platform within the CSP.

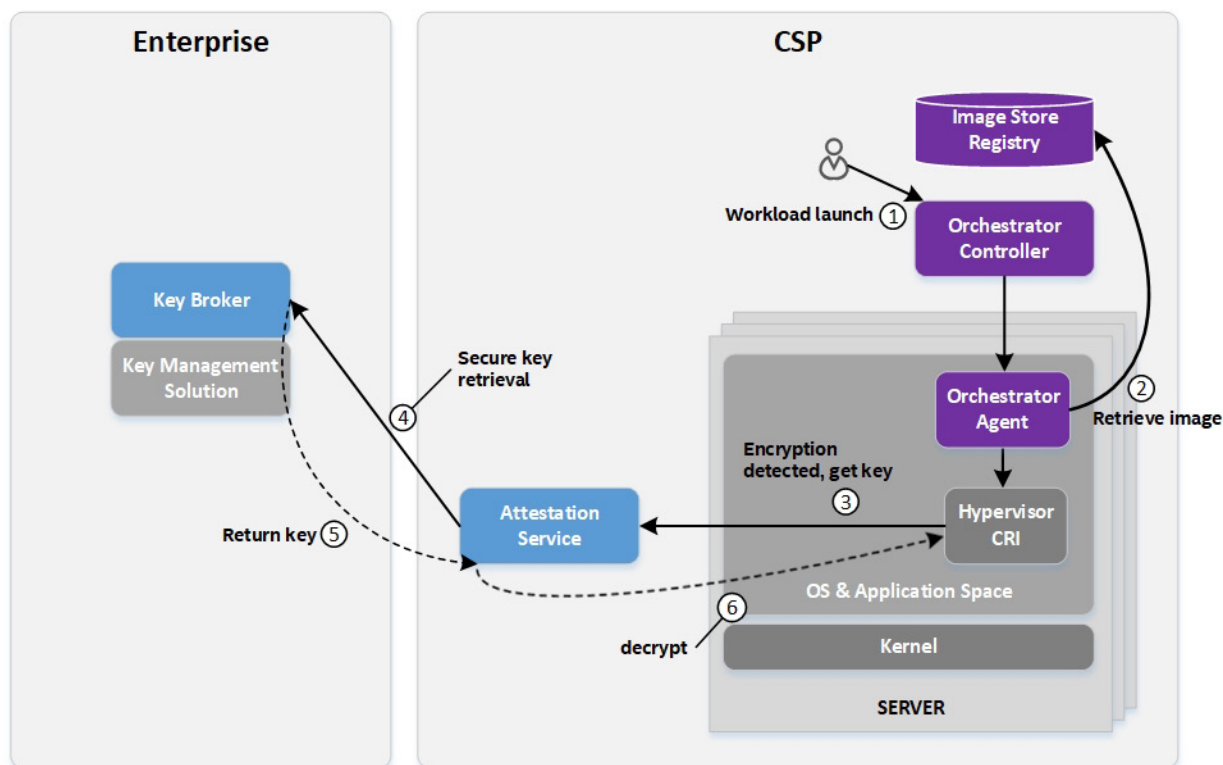


Figure 8: Notional Example of Workload Decryption

874

875

876 When the runtime node service receives the launch request, it can detect that the image is
 877 encrypted and make a request to retrieve the decryption key. This request can be passed through
 878 an attestation service so that an internal trust evaluation for the platform can be performed. The
 879 key request is forwarded to the key broker with proof that the platform has been attested. The
 880 key broker can then verify the attested platform report and release the key back to the CSP and
 881 node runtime services. At that time the node runtime can decrypt the image and proceed with the
 882 normal workload orchestration. The disk encryption kernel subsystem can provide at-rest
 883 encryption for the workload on the platform.

884 **6.5 Protecting Keys and Secrets**

885 Cryptographic keys are high-value assets in workloads, especially in environments where the
 886 owner of the keys is not in complete control of the infrastructure, such as public clouds, edge
 887 computing, and NFV deployments. In these environments, keys are typically provisioned on disk
 888 as flat files or entries in configuration files. At runtime, workloads read the keys into RAM and
 889 use them to perform cryptographic operations like data signing, encryption/decryption, or TLS
 890 termination.

891 Keys on disk and in RAM are exposed to traditional attacks like privilege escalation, remote
 892 code execution (RCE), and input buffer mismanagement. Keys can also be stolen by malicious
 893 administrators or be disclosed because of operational errors. For example, an improperly
 894 protected VM snapshot can be used by a malicious agent to extract keys.

895 An HSM can be attached to a server and used by workloads to store keys and perform
896 cryptographic operations. This results in keys being protected at rest and in use. In this model,
897 keys are never stored on disk or loaded into RAM. If attaching an HSM to a server is not an
898 option, or if keys are needed in many servers at the same time, an alternative option is to use a
899 network HSM. Workloads send the payload that needs cryptographic processing over a network
900 connection to the network HSM, which then performs the cryptographic operations locally,
901 typically in an attached HSM.

902 An HSM option is not feasible in some environments. Workload owners may not have access to
903 a cloud or edge environment in order to attach their HSM to a hardware server. Network HSMs
904 can suffer from network latency, and some workloads require an optimized response time.
905 Additionally, network HSMs are often provided as a service by the cloud, edge, or NFV
906 providers and are billed by the number of transactions. Cost is often a deciding factor for using a
907 provider network HSM.

908 **7 Next Steps**

909 NIST is seeking feedback from the community on the content of the white paper and soliciting
910 additional technology example contributions from other companies. The white paper is intended
911 to be a living document that will be frequently updated to reflect advances in technology and the
912 availability of commercial implementations and solutions. This can help raise the bar on platform
913 security and evolve the use cases.

914 Please send your feedback and comments to hwsec@nist.gov.

915 References

- [1] Wired (2018) *Russia's Elite Hackers Have a Clever New Trick That's Very Hard to Fix*. Available at <https://www.wired.com/story/fancy-bear-hackers-uefi-rootkit>
- [2] Intel Corporation (2019) *Intel® Trusted Execution Technology (Intel® TXT) – Software Development Guide – Measured Launched Environment Developer's Guide, Revision 016*. Available at <https://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>
- [3] Barker EB, Barker WC (2019) Recommendation for Key Management: Part 2 – Best Practices for Key Management Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-57 Part 2, Rev. 1. <https://doi.org/10.6028/NIST.SP.800-57pt2r1>
- [4] Regenscheid AR (2014) BIOS Protection Guidelines for Servers. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-147B. <https://doi.org/10.6028/NIST.SP.800-147B>
- [5] Trusted Computing Group (2003) *TPM 1.2 Main Specification*. Available at <https://www.trustedcomputinggroup.org/tpm-main-specification>
- [6] National Institute of Standards and Technology (2001) Security Requirements for Cryptographic Modules. (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 140-2, Change Notice 2 December 03, 2002. <https://doi.org/10.6028/NIST.FIPS.140-2>
- [7] Intel Corporation (2018) *Intel® Trusted Execution Technology (Intel® TXT) Overview*. Available at <https://www.intel.com/content/www/us/en/support/articles/000025873/technologies.html>
- [8] Futral W, Greene J (2013) *Intel® Trusted Execution Technology for Server Platforms* (Apress, Berkeley, CA). Available at <https://www.apress.com/gp/book/9781430261483>
- [9] Linux Kernel Organization (2020) *Intel® TXT Overview*. Available at https://www.kernel.org/doc/Documentation/intel_txt.txt
- [10] Debian Wiki (2019) *Secure Boot*. Available at <https://wiki.debian.org/SecureBoot>
- [11] Wilkins R, Richardson B (2013) *UEFI Secure Boot in Modern Computer Security Solutions* (Unified Extensible Firmware Interface Forum). Available at https://uefi.org/sites/default/files/resources/UEFI_Secure_Boot_in_Modern_Computer_Security_Solutions_2013.pdf
- [12] RedHat (2018) *README* (for shim). Available at <https://github.com/rhboot/shim/blob/master/README>

- [13] Regenscheid AR (2018) Platform Firmware Resiliency Guidelines. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-193. <https://doi.org/10.6028/NIST.SP.800-193>
- [14] Diamond T, Grayson N, Paulsen C, Polk T, Regenscheid A, Souppaya M, Brown C (2020) Validating the Integrity of Computing Devices: Supply Chain Assurance. (National Institute of Standards and Technology, Gaithersburg, MD). Available at <https://www.nccoe.nist.gov/projects/building-blocks/supply-chain-assurance>
- [15] National Institute of Standards and Technology (2020) *Cyber Supply Chain Risk Management*. Available at <https://csrc.nist.gov/projects/supply-chain-risk-management/>
- [16] Intel Corporation (2020) *Transparent Supply Chain*. Available at <https://www.intel.com/content/www/us/en/products/docs/servers/transparent-supply-chain.html>
- [17] Intel Corporation (2020) *Intel Highlights Latest Security Investments at RSA 2020*. Available at <https://newsroom.intel.com/news-releases/intel-highlights-latest-security-investments-rsa-2020/>
- [18] Department of Defense (2018) Subpart 246.870, Contractors' Counterfeit Electronic Part Detection and Avoidance Systems, *Defense Federal Acquisition Regulation Supplement*. https://www.acq.osd.mil/dpap/dars/dfars/html/current/246_8.htm#246.870-2
- [19] Scarfone KA, Souppaya MP, Hoffman P (2011) Guide to Security for Full Virtualization Technologies. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-125. <https://doi.org/10.6028/NIST.SP.800-125>
- [20] Intel Corporation (2020) *Intel® Virtualization Technology (Intel® VT)*. Available at <https://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html>
- [21] Linux Foundation (2020) *Confidential Computing Consortium*. Available at <https://confidentialcomputing.io>
- [22] Intel Corporation (2020) *Strengthen Enclave Trust with Attestation*. Available at <https://software.intel.com/en-us/sgx/attestation-services>
- [23] European Telecommunications Standards Institute (2020) *Network Functions Virtualisation (NFV)*. Available at <http://www.etsi.org/technologies-clusters/technologies/nfv>
- [24] Intel Corporation (2020) *Flexible Workload Acceleration on Intel Architecture Lowers Equipment Cost*. Available at

<https://www.intel.fr/content/dam/www/public/us/en/documents/white-papers/communications-quick-assist-paper.pdf>

- [25] Tadepalli, H (2017) Intel QuickAssist Technology with Intel Key Protection Technology in Intel Server Platforms Based on Intel Xeon processor Scalable Family. (Intel Corporation). Available at <https://www.aspsys.com/images/solutions/hpc-processors/intel-xeon/Intel-Key-Protection-Technology.pdf>
- [26] Intel Corporation (2020) *Intel® Security Libraries for Data Center (Intel® SecL-DC)*. Available at <https://01.org/intel-secl>
- [27] Bartock MJ, Souppaya MP, Yeluri R, Shetty U, Greene J, Orrin S, Prafullchandra H, McLeese J, Scarfone KA (2015) Trusted Geolocation in the Cloud: Proof of Concept Implementation. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 7904. <https://doi.org/10.6028/NIST.IR.7904>

917 **Appendix A – Acronyms**

918 Selected acronyms used in this paper are defined below.

AC RAM	Authenticated Code Random Access Memory
ACM	Authenticated Code Module
AS	Attestation Service
BIOS	Basic Input Output System
BKC	Best-Known Configuration
BMC	Board Management Controller
BtG	Boot Guard
BYOK	Bring Your Own Key
CA	Certificate Authority
CDN	Content Delivery Network
CoT	Chain of Trust
CPLD	Complex Programmable Logic Device
CPU	Central Processing Unit
CRD	Custom Resource Definition
CRI	Container Runtime Interface
CRTM	Core Root of Trust for Measurement
CRTV	Core Root of Trust for Verification
CSK	Code Signing Key
CSP	Cloud Service Provider
DCRTM	Dynamic Core Root of Trust for Measurement
DFARS	Defense Federal Acquisition Regulation Supplement
DIMM	Dual In-Line Memory Module
DoS	Denial of Service
DRAM	Dynamic Random-Access Memory
EPC	Enclave Page Cache
ETSI	European Telecommunications Standards Institute
ETSI NFV	European Telecommunications Standards Institute Network Functions
SEC	Virtualization Security
FaaS	Function as a Service
FIPS	Federal Information Processing Standard
FPGA	Field Programmable Gate Array
GDPR	General Data Protection Regulation
HIPAA	Health Insurance Portability and Accountability Act

HMEE	Hardware Mediated Execution Enclave
HSM	Hardware Security Module
IBB	Initial Boot Block
IEC	International Electrotechnical Commission
IPsec	Internet Protocol Security
ISA	Instruction Set Architecture
ISecL-DC	Intel Security Libraries for the Data Center
ISO	International Organization for Standardization
IT	Information Technology
ITL	Information Technology Laboratory
KBS	Key Broker Service
KEK	Key Exchange Key
KMIP	Key Management Interoperability Protocol
KMS	Key Management Service
KPT	Key Protection Technology
LCP	Launch Control Policy
ME	Manageability Engine
MLE	Measured Launch Environment
MOK	Machine Owner Key
NFV	Network Functions Virtualization
NIST	National Institute of Standards and Technology
NVRAM	Non-Volatile Random-Access Memory
ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
OS	Operating System
PCH	Platform Controller Hub
PCIE	Peripheral Component Interconnect Express
PCONF	Platform Configuration
PCR	Platform Configuration Register
PFM	Platform Firmware Manifest
PFR	Platform Firmware Resilience
PIT	Protection in Transit
PK	Platform Key
QAT	QuickAssist Technology
RAM	Random Access Memory
RCE	Remote Code Execution

RF	Radio Frequency
RK	Root Key
RoT	Root of Trust
RTD	Root of Trust for Detection
RTM	Root of Trust for Measurement
RTRec	Root of Trust for Recovery
RTU	Root of Trust for Update
SB	UEFI Secure Boot
SCRMTM	Static Core Root of Trust for Measurement
SDN	Software Defined Network
SGX	Software Guard Extensions
SINIT ACM	Secure Initialization Authenticated Code Module
SMBus	System Management Bus
SMM	System Management Mode
SP	Special Publication
SPI	Serial Peripheral Interface
SPS FW	Server Platform Services Firmware
TCB	Trusted Computing Base
TCG	Trusted Computing Group
TEE	Trusted Execution Environment
TLS	Transport Layer Security
TPM	Trusted Platform Module
TXT	Trusted Execution Technology
UEFI	Unified Extensible Firmware Interface
USB	Universal Serial Bus
VM	Virtual Machine
VMM	Virtual Machine Manager