

Draft NISTIR 8320A

1  
2  
  
3  
4  
  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19

# Hardware-Enabled Security:

*Container Platform Security Prototype*

Michael Bartock  
Murugiah Souppaya  
Jerry Wheeler  
Tim Knoll  
Uttam Shetty  
Ryan Savino  
Joseprabu Inbaraj  
Stefano Righi  
Karen Scarfone

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8320A-draft>

20  
21  
  
22  
23  
  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43

# Hardware-Enabled Security:

## *Container Platform Security Prototype*

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

Joseprabu Inbaraj  
Stefano Righi  
*AMI  
Atlanta, GA*

Jerry Wheeler  
Tim Knoll  
Uttam Shetty  
Ryan Savino  
*Intel Corporation  
Santa Clara, California*

Karen Scarfone  
*Scarfone Cybersecurity  
Clifton, VA*

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8320A-draft>

December 2020



44  
45  
46  
47  
48  
49  
50

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*

51 National Institute of Standards and Technology Interagency or Internal Report 8320A  
52 49 pages (December 2020)

53 This publication is available free of charge from:  
54 <https://doi.org/10.6028/NIST.IR.8320A-draft>

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

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

65 Organizations are encouraged to review all draft publications during public comment periods and provide feedback to  
66 NIST. All NIST Computer Security Division publications, other than the ones noted above, are available at  
67 <http://csrc.nist.gov/publications>.

68

69 **Public comment period: *December 7, 2020 through January 29, 2021***

70 National Institute of Standards and Technology  
71 Attn: Applied Cybersecurity Division, Information Technology Laboratory  
72 100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-2000  
73 Email: [hwsec@nist.gov](mailto:hwsec@nist.gov)

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

76

77

## Reports on Computer Systems Technology

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

86

### Abstract

87 In today’s cloud data centers and edge computing, attack surfaces have significantly increased,  
88 hacking has become industrialized, and most security control implementations are not coherent  
89 or consistent. The foundation of any data center or edge computing security strategy should be  
90 securing the platform on which data and workloads will be executed and accessed. The physical  
91 platform represents the first layer for any layered security approach and provides the initial  
92 protections to help ensure that higher-layer security controls can be trusted. This report explains  
93 an approach based on hardware-enabled security techniques and technologies for safeguarding  
94 container deployments in multi-tenant cloud environments. It also describes a proof-of-concept  
95 implementation of the approach—a prototype—that is intended to be a blueprint or template for  
96 the general security community.

97

### Keywords

98 container; hardware-enabled security; hardware root of trust; platform security; trusted compute  
99 pool; virtualization.

100

### Acknowledgments

101 The authors thank everyone who contributed their time and expertise to the development of this  
102 report, in particular Eric Johnson, Muthukkumaran Ramalingam, and Madhan B. Santharam  
103 from AMI.

104

### Audience

105 The primary audiences for this report are security professionals, such as security engineers and  
106 architects; system administrators and other information technology (IT) professionals for cloud  
107 service providers; and hardware, firmware, and software developers who may be able to leverage  
108 hardware-enabled security techniques and technologies to improve platform security for  
109 container deployments in multi-tenant cloud environments.

110

### Trademark Information

111 All registered trademarks or other trademarks belong to their respective organizations.

112

113

## Call for Patent Claims

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

120

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

123

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

126

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

130

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

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

135

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

141

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

144

145 Such statements should be addressed to: [hwsec@nist.gov](mailto:hwsec@nist.gov)

146

147 **Table of Contents**

148 **1 Introduction ..... 1**

149 1.1 Purpose and Scope .....1

150 1.2 Document Structure .....1

151 **2 Prototype Implementation ..... 2**

152 2.1 Objective .....2

153 2.2 Goals.....2

154 2.2.1 Stage 0: Platform attestation and measured worker node launch ..... 3

155 2.2.2 Stage 1: Trusted workload placement ..... 4

156 2.2.3 Stage 2: Asset tagging and trusted location ..... 5

157 **3 Prototyping Stage 0 ..... 7**

158 3.1 Solution Overview .....7

159 3.2 Solution Architecture .....8

160 **4 Prototyping Stage 1 .....11**

161 4.1 Solution Overview ..... 11

162 4.2 Solution Architecture ..... 12

163 **5 Prototyping Stage 2 .....13**

164 5.1 Solution Overview ..... 13

165 **References .....14**

166 **Appendix A— Hardware Architecture and Prerequisites .....15**

167 A.1 High-Level Implementation Architecture ..... 15

168 A.2 Intel Trusted Execution Technology (Intel TXT) & Trusted Platform Module (TPM) .15

169 A.3 Attestation ..... 17

170 **Appendix B— Platform Implementation: AMI TruE .....19**

171 B.1 Solution Architecture ..... 19

172 B.2 Hardware Description ..... 19

173 B.3 AMI TruE Installation and Configuration .....20

174 B.3.1 Installing AMI TruE Trust Manager .....20

175 B.3.2 Installing AMI TruE Attestation Server .....21

176 B.3.3 Configuring Firewall for AMI True .....21

177 B.3.4 Configuring Device Access Keys .....22

178 B.3.5 Configuring Discovery Range and Manageability Range .....22

179 B.4 Trusted Cloud Cluster Installation and Configuration .....23

180 B.4.1 Provisioning TruE Agent Remotely .....23

181 B.4.2 Provisioning TruE Agent Manually ..... 24

182 B.5 Using AMI TruE .....26

183 B.5.1 Monitoring Trust Status using AMI TruE .....26

184 B.5.2 Generating Trust Reports .....28

185 B.5.3 Tagging Platforms Using AMI TruE .....29

186 B.5.4 Receiving Trust Event Alert Notification .....30

187 B.5.5 Using AMI TruE for Remediation .....31

188 **Appendix C— Platform Implementation: Kubernetes .....34**

189 C.1 Prototype Architecture ..... 34

190 C.2 Hardware Description .....34  
 191 C.3 Kubernetes Installation and Configuration .....34  
 192 C.3.1 Kubernetes Control Node Configuration .....35  
 193 C.3.2 Kubernetes Worker Configuration .....36  
 194 C.3.3 Kubernetes Orchestration .....36  
 195 **Appendix D— Supporting NIST SP 800-53 Security Controls.....37**  
 196 **Appendix E— Cybersecurity Framework Subcategory Mappings .....39**  
 197 **Appendix F— Acronyms and Other Abbreviations .....40**

198

199

**List of Tables**

200 Table 1: Trusted Boot Compliance View .....13  
 201 Table 2: System Hostnames and IP Configurations .....15  
 202 Table 3: Hardware for Implemented Architecture .....19  
 203 Table 4: Minimum System Requirements to Install AMI TruE .....20  
 204 Table 5: Network Ports for Trust Manager.....21  
 205 Table 6: Network Ports for Attestation Server .....22  
 206 Table 7: Security Capabilities Provided by the Trusted Asset Tag Prototype .....37  
 207 Table 8: Mapping of Security Capabilities to NIST SP 800-53 Controls.....38  
 208 Table 9: Mapping of Security Capabilities to NIST Cybersecurity Framework Subcategories....39

209

210

**List of Figures**

211 Figure 1: Concept of Trusted Compute Pools ..... 7  
 212 Figure 2: Stage 0 Solution System Architecture ..... 9  
 213 Figure 3: Stage 1 Solution Overview .....11  
 214 Figure 4: Single Server Overview.....13  
 215 Figure 5: Logical Network Architecture of the Prototype Implementation.....15  
 216 Figure 6: Remote Attestation Protocol.....17  
 217 Figure 7: AMI TruE Prototype Implementation .....19  
 218 Figure 8: Examples of Manageability Ranges .....23  
 219 Figure 9: Example of Adding an Image .....24  
 220 Figure 10: Example of Job List.....24  
 221 Figure 11: Example of AMI TruE Report Dashboard .....27  
 222 Figure 12: Example of Summary Page for Server Trust Information.....28  
 223 Figure 13: Example of Trust Report .....29

224 Figure 14: Example of Creating and Deploying an Asset Tag .....30

225 Figure 15: Example of the "Add New Subscription" Page.....31

226 Figure 16: Example of Remote Power Control .....32

227 Figure 17: Kubelet and Docker Versions .....35

228

## 229 **1 Introduction**

### 230 **1.1 Purpose and Scope**

231 The purpose of this publication is to describe an approach for safeguarding application container  
232 deployments in multi-tenant cloud environments. This publication first explains selected security  
233 challenges involving Infrastructure as a Service (IaaS) cloud computing technologies and geolocation in  
234 the form of resource asset tags. It then describes a proof-of-concept implementation—a prototype—that  
235 was designed to address those challenges. The publication provides sufficient details about the prototype  
236 implementation so that organizations can reproduce it if desired. The publication is intended to be a  
237 blueprint or template that can be used by the general security community to validate and implement the  
238 described implementation.

239 It is important to note that the prototype implementation presented in this publication is only one possible  
240 way to solve the security challenges. It is not intended to preclude the use of other products, services,  
241 techniques, etc. that can also solve the problem adequately, nor is it intended to preclude the use of any  
242 cloud products or services not specifically mentioned in this publication.

243 This publication builds upon the terminology and concepts described in the NIST white paper draft,  
244 *Hardware-Enabled Security for Server Platforms: Enabling a Layered Approach to Platform Security for*  
245 *Cloud and Edge Computing Use Cases* [1]. Reading that white paper is a prerequisite for reading this  
246 publication because it explains the concepts and defines key terminology used in this publication.

### 247 **1.2 Document Structure**

248 This document is organized into the following sections and appendices:

- 249 • Section 2 defines the problem (use case) to be solved.
- 250 • Sections 3, 4, and 5 describe the three stages of the prototype implementation:
  - 251 ○ Stage 0: Platform attestation and measured worker node launch
  - 252 ○ Stage 1: Trusted workload placement
  - 253 ○ Stage 2: Asset tagging and trusted location
- 254 • The References sections provides references for the document.
- 255 • Appendix A provides an overview of the high-level hardware architecture of the prototype  
256 implementation, as well as details on how Intel platforms implement hardware modules and  
257 enhanced hardware-based security functions.
- 258 • Appendix B contains supplementary information provided by AMI describing all the required  
259 components and steps required to set up the prototype implementation.
- 260 • Appendix C contains supplementary information describing all the required components and  
261 steps required to set up the prototype implementation for Kubernetes.
- 262 • Appendix D lists the major controls from NIST Special Publication (SP) 800-53 Revision  
263 5, *Security and Privacy Controls for Information Systems and Organizations* that affect container  
264 platform security prototype implementation.
- 265 • Appendix E maps the major security features from the prototype implementation to the  
266 corresponding subcategories from the Cybersecurity Framework.
- 267 • Appendix F lists and defines acronyms and other abbreviations used in the document.

## 268 **2 Prototype Implementation**

269 This section defines the prototype implementation. Section 2.1 explains the basics of the objective.  
270 Section 2.2 provides more details, outlining all of the intermediate goals that must be met in order to  
271 achieve the desired prototype implementation. These requirements are grouped into three stages of the use  
272 case, each of which is examined more closely in Sections 2.2.1 through 2.2.3, respectively.

### 273 **2.1 Objective**

274 Shared cloud computing technologies are designed to be highly agile and flexible, transparently using  
275 whatever resources are available to process container deployments for their customers [2]. However, there  
276 are security and privacy concerns with allowing unrestricted container deployment orchestration.  
277 Whenever multiple container deployments are present on a single cloud server, there is a need to  
278 segregate those deployments from each other so that they do not interfere with each other, gain access to  
279 each other's sensitive data, or otherwise compromise the security or privacy of the containers. Imagine  
280 two rival companies with container deployments on the same server; each company would want to ensure  
281 that the server can be trusted to protect their information from the other company. Similarly, a single  
282 organization might have multiple container deployments that need to be kept separate because of differing  
283 security requirements and needs for each app.

284 Another concern with shared cloud computing is that workloads could move from cloud servers located in  
285 one country to servers located in another country. Each country has its own laws for data security,  
286 privacy, and other aspects of information technology (IT). Because the requirements of these laws may  
287 conflict with an organization's policies or mandates (e.g., laws, regulations), an organization may decide  
288 that it needs to restrict which cloud servers it uses based on their location. A common desire is to only use  
289 cloud servers physically located within the same country as the organization, or physically located in the  
290 same country as the origin of the information. Determining the approximate physical location of an  
291 object, such as a cloud computing server, is generally known as *geolocation*. Geolocation can be  
292 accomplished in many ways, with varying degrees of accuracy, but traditional geolocation methods are  
293 not secured and are enforced through management and operational controls that cannot be automated or  
294 scaled. Therefore, traditional geolocation methods cannot be trusted to meet cloud security needs.

295 The motivation behind this use case is to improve the security and accelerate the adoption of cloud  
296 computing technologies by establishing an automated, hardware root-of-trust method for enforcing and  
297 monitoring platform integrity and geolocation restrictions for cloud servers. A *hardware root-of-trust* is  
298 an inherently trusted combination of hardware and firmware responsible for measuring the integrity of the  
299 platform and geolocation information in the form of an asset tag. The measurements taken by the  
300 hardware root-of-trust are stored in tamper-resistant hardware, and are transmitted using cryptographic  
301 keys unique to that tamper-resistant hardware. This information is accessed by management and security  
302 tools using cryptographic protocols to assert the integrity of the platform and confirm the location of the  
303 host.

304 This use case builds on earlier work documented in NIST IR 7904, *Trusted Geolocation in the Cloud:  
305 Proof of Concept Implementation* [3].

### 306 **2.2 Goals**

307 Using trusted compute pools (described in Section 3) is a leading approach for aggregating trusted  
308 systems and segregating them from untrusted resources, which results in the separation of higher-value,  
309 more sensitive workloads from commodity application and data workloads. The principles of operation  
310 are to:

- 311 1. Create a part of the cloud to meet the specific and varying security requirements of users.  
 312 2. Control access to that portion of the cloud so that the right applications (workloads) get  
 313 deployed there.  
 314 3. Enable audits of that portion of the cloud so that users can verify compliance.

315 These trusted compute pools allow IT to gain the benefits of the dynamic cloud environment while still  
 316 enforcing higher levels of protections for their more critical workloads.

317 The ultimate goal is to be able to use “trust” as a logical boundary for deploying cloud workloads on  
 318 server platforms within a cloud. This goal is dependent on smaller prerequisite goals described as stages,  
 319 which can be thought of as requirements that the solution must meet. Because of the number of  
 320 prerequisites, they have been grouped into three stages:

- 321 0. Platform attestation and measured worker node launch. This ensures that the integrity of the  
 322 cloud server platform is measured and available for the subsequent stages.
- 323 1. Trusted workload placement. This stage allows container deployments to be orchestrated to  
 324 launch only on trusted server platforms within a cloud.
- 325 2. Asset tagging and trusted location. This stage allows container deployments to be launched only  
 326 on trusted server platforms within a cloud, taking into consideration qualitative asset tag  
 327 restrictions (for example, location information).

328 The prerequisite goals for each stage, along with more general information on each stage, are explained  
 329 below.

### 330 **2.2.1 Stage 0: Platform attestation and measured worker node launch**

331 A fundamental component of a solution is having some assurance that the platform the container  
 332 deployment is running on can be trusted. If the platform is not trustworthy, then not only is it putting the  
 333 tenant’s application and data at greater risk of compromise, but also there is no assurance that the claimed  
 334 asset tag of the cloud server is accurate. Having basic assurance of trustworthiness is the initial stage in  
 335 the solution.

336 Stage 0 includes the following prerequisite goals:

- 337 1. **Configure a cloud server platform as being trusted.** The “cloud server platform” includes the  
 338 hardware configuration (e.g., BIOS integrity), operating system (OS) configuration (boot loader  
 339 and OS kernel configuration and integrity), and the integrity of the container runtime. This also  
 340 includes the varying hardware security technologies enabled on the server. These chain of trust  
 341 (CoT) technologies provide platform integrity verification. Additional technologies and details  
 342 can be found in the aforementioned NIST white paper draft [1] and are discussed in Section 3.2.
- 343 2. **Before each container worker node launch, verify (measure) the trustworthiness of the**  
 344 **cloud server platform.** The items configured in goal 1 (BIOS, OS, container runtime) need to  
 345 have their configurations verified before launching the container runtime to ensure that the  
 346 assumed level of trust is still in place.
- 347 3. **During container runtime execution, periodically audit the trustworthiness of the cloud**  
 348 **server platform.** This periodic audit is essentially the same check as that performed as goal 2,  
 349 except that it is performed frequently while the container runtime is executing. Ideally this  
 350 checking would be part of continuous monitoring.

351 Achieving all of these goals will not prevent attacks from succeeding, but will cause unauthorized  
352 changes to the cloud platform to be detected more rapidly than they otherwise would have been. This  
353 prevents new container deployments with trust requirements from being launched on the compromised  
354 platform.

355 For more information on the technical topics being addressed by these goals, see the following NIST  
356 publications:

- 357 • NIST SP 800-128, *Guide for Security-Focused Configuration Management of Information*  
358 *Systems*  
359 <https://doi.org/10.6028/NIST.SP.800-128>
- 360 • NIST SP 800-137, *Information Security Continuous Monitoring for Federal Information Systems*  
361 *and Organizations*  
362 <https://doi.org/10.6028/NIST.SP.800-137>
- 363 • NIST SP 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*  
364 <https://doi.org/10.6028/NIST.SP.800-144>
- 365 • NIST SP 800-147B, *BIOS Protection Guidelines for Servers*  
366 <https://doi.org/10.6028/NIST.SP.800-147B>
- 367 • Draft NIST SP 800-155, *BIOS Integrity Measurement Guidelines*  
368 <https://csrc.nist.gov/publications/detail/sp/800-155/draft>
- 369 • NIST SP 800-190, *Application Container Security Guide*  
370 <https://doi.org/10.6028/NIST.SP.800-190>

## 371 **2.2.2 Stage 1: Trusted workload placement**

372 Once stage 0 has been successfully completed, the next objective is to be able to orchestrate the  
373 placement of workloads to launch only on trusted platforms. Workload placement is a key attribute of  
374 cloud computing, improving scalability and reliability. The purpose of this stage is to ensure that any  
375 server that a workload is launched on will meet the required level of security assurance based on the  
376 workload security policy.

377 Stage 1 includes the following prerequisite goal:

- 378 1. **Deploy workloads only to cloud servers with trusted platforms.** This basically means that you  
379 perform stage 0, goal 3 (auditing platform trustworthiness) and only deploy a workload to the  
380 cloud server if the audit demonstrates that the platform is trustworthy.

381 Achieving this goal ensures that the workloads are deployed to trusted platforms, thus reducing the  
382 chance of workload compromise.

383 For more information on the technical topics being addressed by these goals, see the following NIST  
384 publications:

- 385 • Draft NIST Cybersecurity White Paper, *Hardware-Enabled Security for Server Platforms:  
386 Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases*  
387 <https://doi.org/10.6028/NIST.CSWP.04282020-draft>
- 388 • NIST SP 800-137, *Information Security Continuous Monitoring for Federal Information Systems  
389 and Organizations*  
390 <https://doi.org/10.6028/NIST.SP.800-137>
- 391 • NIST SP 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*  
392 <https://doi.org/10.6028/NIST.SP.800-144>
- 393 • NIST SP 800-147B, *BIOS Protection Guidelines for Servers*  
394 <https://doi.org/10.6028/NIST.SP.800-147B>
- 395 • Draft NIST SP 800-155, *BIOS Integrity Measurement Guidelines*  
396 <https://csrc.nist.gov/publications/detail/sp/800-155/draft>

### 397 **2.2.3 Stage 2: Asset tagging and trusted location**

398 The next stage builds upon stage 1 by adding the ability to continuously monitor and enforce asset tag  
399 restrictions.

400 Stage 2 includes the following prerequisite goals:

- 401 1. **Have trusted asset tag information for each trusted platform instance.** This information  
402 would be stored within the cloud server's cryptographic module (as a cryptographic hash within  
403 the hardware cryptographic module) so that it could be verified and audited readily.
- 404 2. **Provide configuration management and policy enforcement mechanisms for trusted  
405 platforms that include enforcement of asset tag restrictions.** This goal builds upon stage 1,  
406 goal 1 (deploy workloads only to cloud servers with trusted platforms); it enhances stage 2, goal 1  
407 by adding an asset tag check to the server to launch the workload on.
- 408 3. **During workload orchestration, periodically audit the asset tag of the cloud server platform  
409 against asset tag policy restrictions.** This goal is built upon stage 0, goal 3 (auditing platform  
410 trustworthiness), but it is specifically auditing the asset tag information against the policies for  
411 asset tag to ensure that the server's asset tagging does not violate the policies.

412 Achieving these goals ensures that the workloads are not launched on a server in an unsuitable boundary  
413 location. This avoids issues caused by clouds spanning different physical locations (e.g., regulations,  
414 sensitivity levels, countries or states with different data security and privacy laws).

415 For more information on the technical topics being addressed by these goals, see the following NIST  
416 publications:

- 417 • NIST SP 800-128, *Guide for Security-Focused Configuration Management of Information  
418 Systems*  
419 <https://doi.org/10.6028/NIST.SP.800-128>
- 420 • NIST SP 800-137, *Information Security Continuous Monitoring for Federal Information Systems  
421 and Organizations*  
422 <https://doi.org/10.6028/NIST.SP.800-137>
- 423 • NIST SP 800-147B, *BIOS Protection Guidelines for Servers*  
424 <https://doi.org/10.6028/NIST.SP.800-147B>

- 425
- Draft NIST SP 800-155, *BIOS Integrity Measurement Guidelines*
- 426
- <https://csrc.nist.gov/publications/detail/sp/800-155/draft>

427

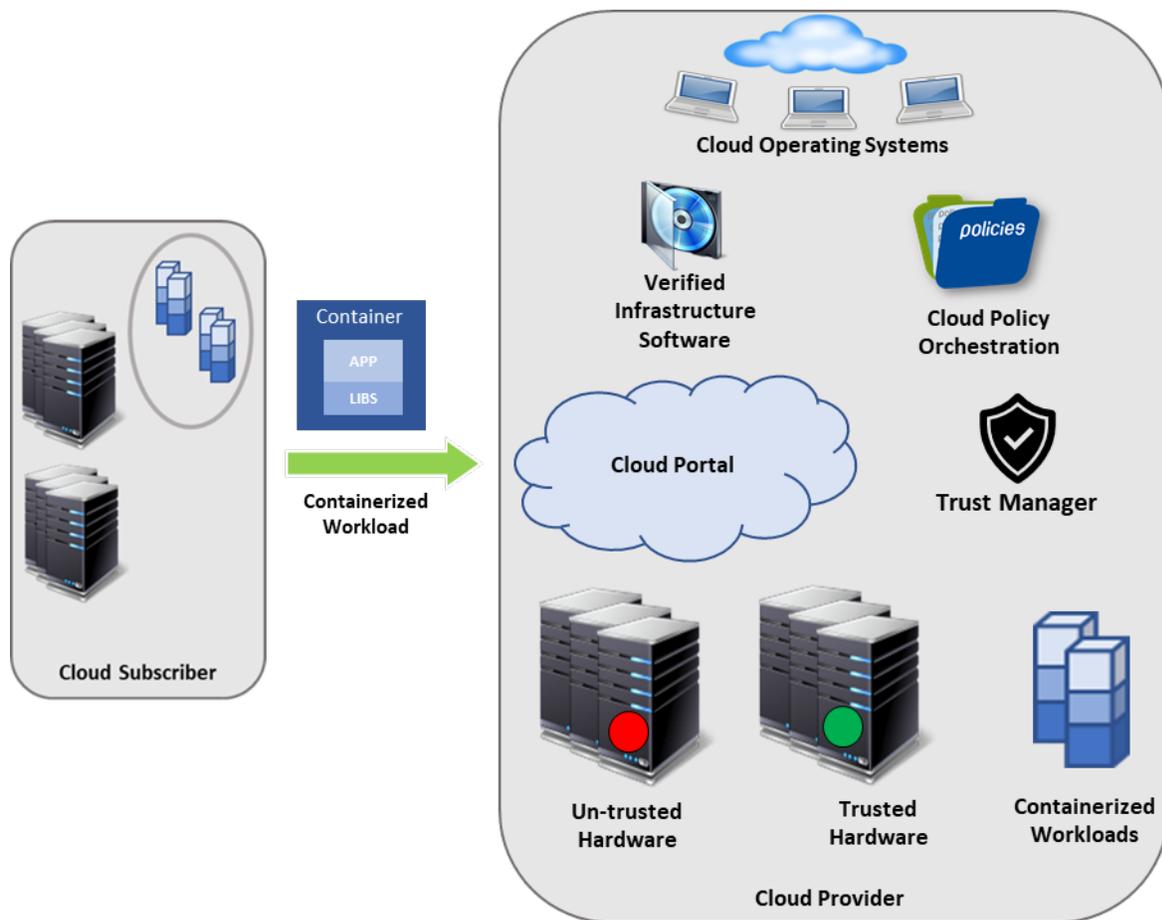
428 **3 Prototyping Stage 0**

429 This section describes stage 0 of the prototype implementation (platform attestation and measured worker  
430 node launch).

431 **3.1 Solution Overview**

432 This stage of the use case enables the creation of what are called *trusted compute pools*. Also known as  
433 *trusted pools*, they are physical or logical groupings of computing hardware in a data center that are  
434 tagged with specific and varying security policies, and the access and execution of apps and workloads  
435 are monitored, controlled, audited, etc. In this phase of the solution, an attested launch of the platform is  
436 deemed as a trusted node and is added to the trusted pool.

437 Figure 1 depicts the concept of trusted pools. The resources tagged green indicate trusted ones. Critical  
438 policies can be defined such that security-sensitive cloud services can only be launched on these trusted  
439 resources.



440  
441

Figure 1: Concept of Trusted Compute Pools

442 In order to have a trusted launch of the platform, the two key questions that should be answered are:

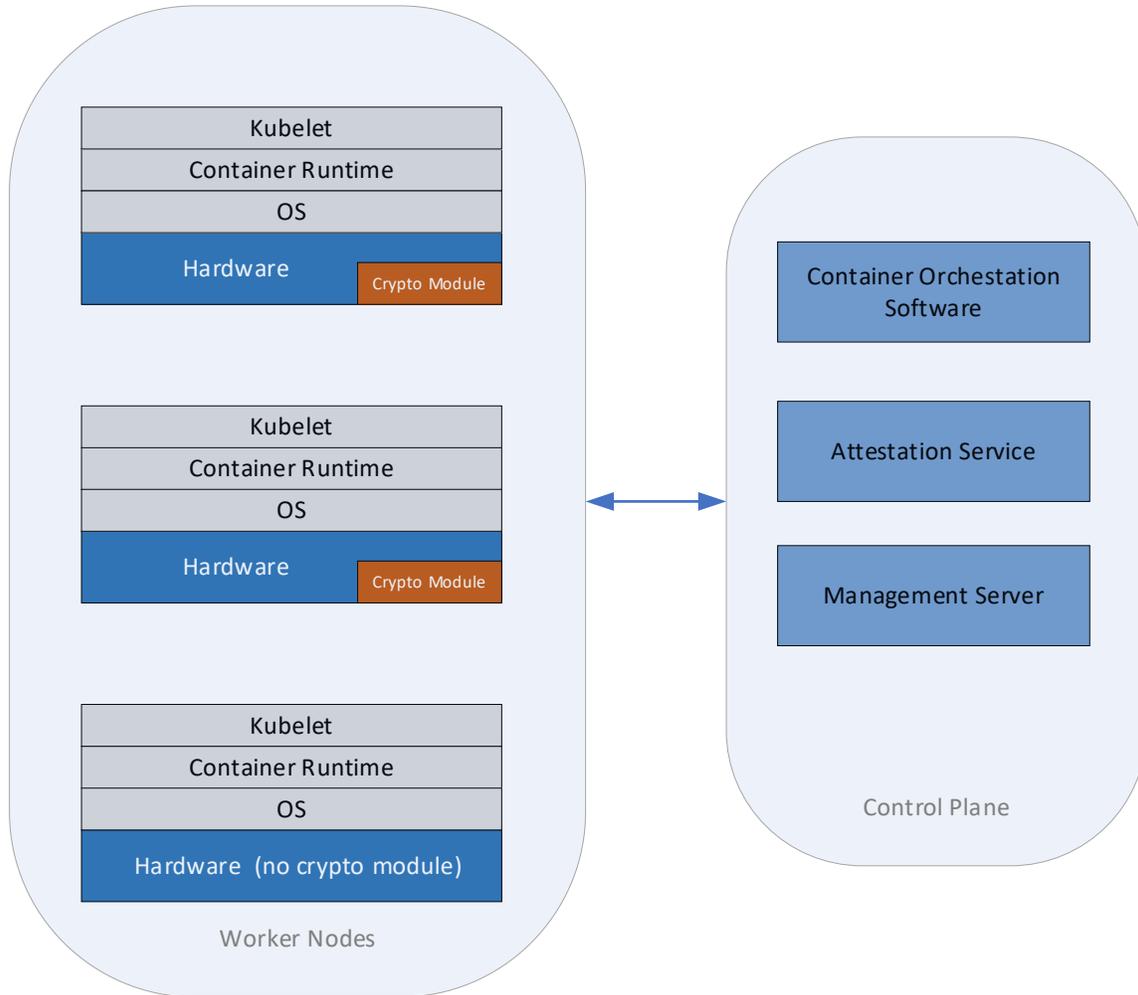
- 443 1. How would the entity needing this information know if a specific platform has the necessary  
444 enhanced hardware-based security features enabled and if a specific platform has a  
445 defined/compliant OS/platform firmware and container runtime running on it?
- 446 2. Why should the entity requesting this information, which in a cloud environment would be a  
447 scheduler/orchestrator trying to schedule a workload on a set of available nodes/servers, believe  
448 the response from the platform?

449 Attestation provides the definitive answers to these questions. *Attestation* is the process of providing a  
450 digital signature of a set of measurements securely stored in hardware, then having the requestor validate  
451 the signature and the set of measurements. Attestation requires roots of trust. The platform has to have a  
452 Root-of-Trust for Measurement (RTM) that is implicitly trusted to provide an accurate measurement, and  
453 enhanced hardware-based security features provide the RTM. The platform also has to have a Root-of-  
454 Trust for Reporting (RTR) and a Root-of-Trust for Storage (RTS), and the same enhanced hardware-  
455 based security features provide these.

456 The entity that challenged the platform for this information now can make a determination about the trust  
457 of the launched platform by comparing the provided set of measurements with “known good/golden”  
458 measurements. Managing the “known good” for different platforms and operating systems, and various  
459 BIOS software, and ensuring they are protected from tampering and spoofing is a critical IT operations  
460 challenge. This capability can be internal to a service provider, or it could be delivered as a service by a  
461 trusted third party for service providers and enterprises to use.

### 462 **3.2 Solution Architecture**

463 Figure 2 provides a layered view of the solution system architecture. The indicated servers in the resource  
464 pool include a hardware module for storing sensitive keys and measurements. All the servers are  
465 configured by the cloud management server.



466

467

**Figure 2: Stage 0 Solution System Architecture**

468  
469  
470  
471  
472  
473

The initial step in instantiating the architecture requires provisioning the server for enhanced hardware-based security features. This requires either physical or remote access to the server to access the BIOS, enable a set of configuration options to use the hardware module (including taking ownership of the module), and activate the enhanced hardware-based security features. This process is highly BIOS and original equipment manufacturer (OEM) dependent. This step is mandatory for a measured launch of the platform.

474  
475  
476  
477  
478  
479  
480

The management console provides remote monitoring and management of all servers in this solution architecture. It allows remote configuration of BIOS that are required for a server to be measured and secured. It periodically checks the measurements of all monitored servers and compares them against golden measurements that were taken in pristine condition. When any such measurements do not match, indicating a platform security compromise, it notifies the administrator through email and/or text message. The administrator can then use the management console to take remediation actions, which could include powering down the server or reconfiguring or updating the firmware of the server.

481  
482  
483

The platform undergoes a measured launch, and the BIOS and OS components are measured (cryptographically) and placed into the server hardware security module. These measurement values are accessible through the cloud management server via the application programming interface (API). When

484 the hosts are initially configured with the cloud management server, the relevant measurement values are  
485 cached in the cloud management database.

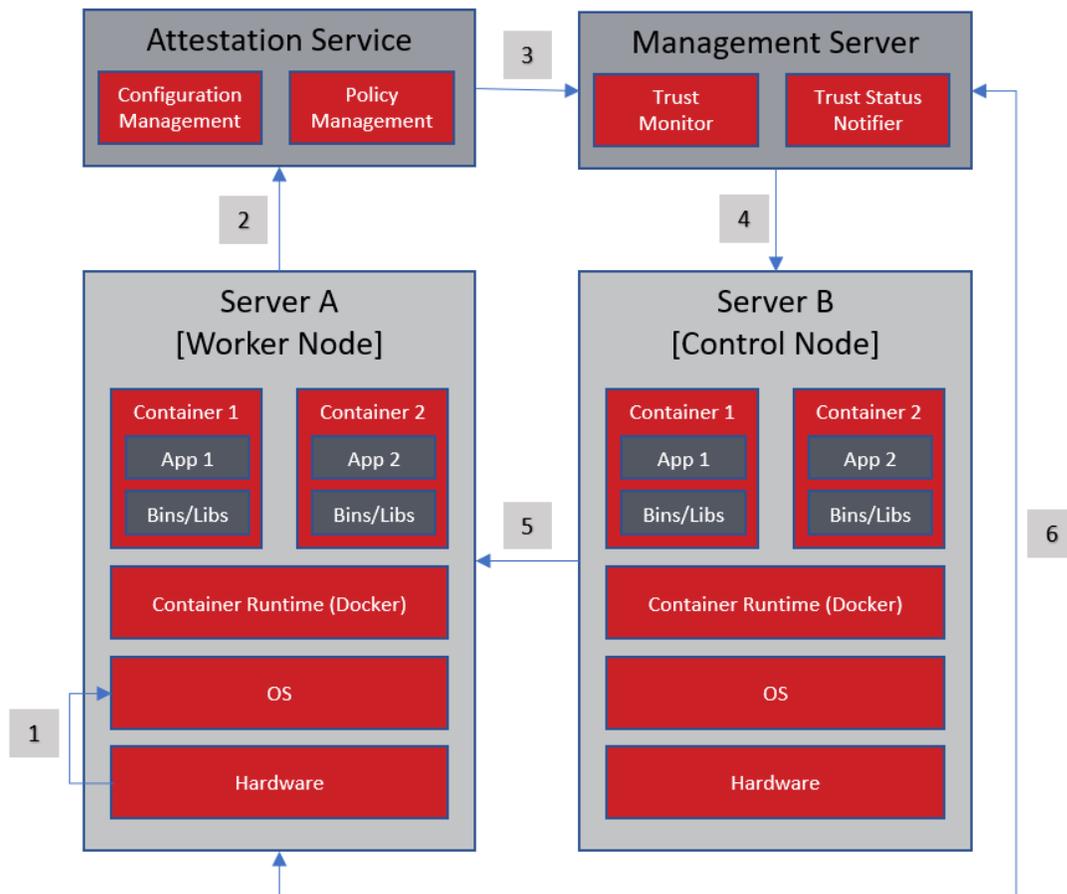
486 In addition to the measured launch, this solution architecture also provides provisions to assign a secure  
487 asset tag to each of the servers during the provisioning process. The asset tag is provisioned to a non-  
488 volatile index in the hardware module via an out-of-band mechanism, and on a platform launch the  
489 contents of the index are inserted/extended into the hardware module. Enhanced hardware-based security  
490 features provide the interface and attestation to the asset tag information, including the asset tag lookup  
491 and user-readable/presentable string/description.

492 **4 Prototyping Stage 1**

493 This section discusses stage 1 of the prototype implementation (trusted workload placement), which is  
 494 based on the stage 0 work and adds components that orchestrate the placement of workloads to launch on  
 495 trusted platforms.

496 **4.1 Solution Overview**

497 Figure 3 shows the operation of the stage 1 solution. It assumes that Server A and Server B are two  
 498 servers within the same cloud.



499  
500

**Figure 3: Stage 1 Solution Overview**

501 There are six generic steps performed in the operation of the stage 1 prototype, as outlined below and  
 502 reflected by the numbers in Figure 3:

- 503 1. Server A performs a measured launch, with the enhanced hardware-based security features  
 504 populating the measurements in the hardware module.
- 505 2. Server A sends a quote to the Trust Authority. The quote includes signed hashes of various  
 506 platform firmware and OS components.
- 507 3. The Trust Authority verifies the signature and hash values, and sends the attestation of the  
 508 platform’s integrity state to the management server.

- 509        4. The management server enforces workload policy requirements on Server B based on user  
510            requirements.
- 511        5. Server B launches workloads that require trusted infrastructure only on server platforms that have  
512            been attested to be trusted.
- 513        6. Each server platform gets audited periodically based on its measurement values.

## 514    **4.2 Solution Architecture**

515    The stage 1 architecture is identical to the stage 0 architecture (see Figure 2), with additional  
516    measurement occurring related to the orchestration of workload placement among trusted hosts.

517

518 **5 Prototyping Stage 2**

519 This section discusses stage 2 of the prototype implementation (trust-based and asset tag-based secure  
520 workload placement), which is based on the stage 1 work and adds components that take into account  
521 asset tag restrictions.

522 **5.1 Solution Overview**

523 Stage 2 adds the monitoring of measurements in a governance, risk, and compliance dashboard. One chart  
524 that might appear in such a dashboard could reflect the relative size of the pools of trusted and untrusted  
525 cloud servers. This could be displayed by percentage and/or count.

526 Table 1 is a drill-down page from the high-level dashboard view. It provides more details on all the  
527 servers within the cloud. In this example, there are three servers. Information listed for each server  
528 includes the server’s IP address and universally unique identifier (UUID), and the status of the  
529 measurements (trusted boot validation and system health validation).

530 **Table 1: Trusted Boot Compliance View**

Cloud Host IP	Hardware UUID	Trusted Boot Validation	System Health Validation
<Host 1>	<UUID 1>	Yes/No	Yes/No
<Host 2>	<UUID 2>	Yes/No	Yes/No
<Host 3>	<UUID 3>	Yes/No	Yes/No

531

532 Figure 4 shows a drill-down from Table 1 for an individual server. It includes a detailed measurement  
533 data for the trusted boot validation, alongside the connection status and asset tag list which may include  
534 asset tag value. It also shows when the server was measured and when the validity of this measurement  
535 expires. Measuring each server’s characteristics frequently (such as every five minutes) helps to achieve a  
536 continuous monitoring solution for the servers.

General Information			
<b>Host Info:</b>	<IP Address or Host Name>	<b>UUID:</b>	<Unique ID>
<b>Trust Report Created On:</b>	<Time Stamp>	<b>Trust Report Expires On:</b>	<Time Stamp>
<b>Asset Tag Status:</b>	<Deployed/Not Deployed>	<b>Asset Tag List:</b>	<Name-1/Value-1> <Name-N/Value-N>
<b>Flavor Group Name:</b>	<Name>	<b>Connection Status:</b>	<Connected / Not connected>

Trust Information			
<b>Overall System Trust:</b>	<Trusted/Untrusted>		
<b>Software Trust:</b>	<Trusted/Untrusted>	<b>Platform Trust:</b>	<Trusted/Untrusted>
<b>Asset Tag Trust:</b>	<Trusted/Untrusted>	<b>Host Unique Trust:</b>	<Trusted/Untrusted>

537

538

**Figure 4: Single Server Overview**

539

540 **References**

541 References for this publication are listed below.

- [1] Bartock MJ, Souppaya MP, Savino R, Knoll T, Shetty U, Cherfaoui M, Yeluri R, Scarfone KA (2020) Hardware-Enabled Security for Server Platforms: Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases. (National Institute of Standards and Technology, Gaithersburg, MD), Draft NIST Cybersecurity White Paper. <https://doi.org/10.6028/NIST.CSWP.04282020-draft>
- [2] Souppaya MP, Scarfone KA, Morello J (2017) Application Container Security Guide. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-190. <https://doi.org/10.6028/NIST.SP.800-190>
- [3] 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>
- [4] Greene J (2012) Intel Trusted Execution Technology: Hardware-based Technology for Enhancing Server Platform Security. (Intel Corporation). Available at <http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-security-paper.html>
- [5] AMI (2020) AMI TruE Trusted Environment. Available at <https://ami.com/en/products/security-services-and-solutions/ami-true-trusted-environment/>
- [6] Joint Task Force (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-53, Rev. 5. <https://doi.org/10.6028/NIST.SP.800-53r5>
- [7] National Institute of Standards and Technology (2018) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (National Institute of Standards and Technology, Gaithersburg, MD). <https://doi.org/10.6028/NIST.CSWP.04162018>

542

543 **Appendix A—Hardware Architecture and Prerequisites**

544 This appendix provides an overview of the high-level hardware architecture of the prototype  
545 implementation, as well as details on how Intel platforms implement hardware modules and enhanced  
546 hardware-based security functions [4].

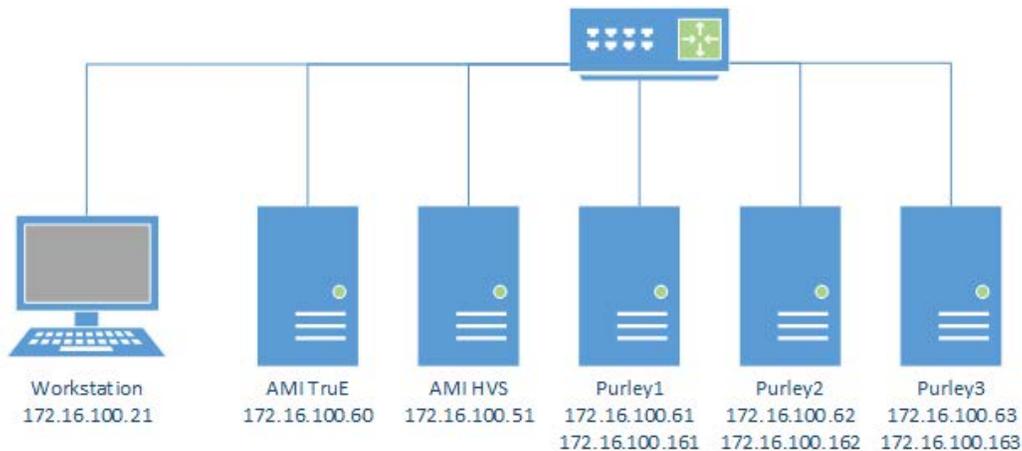
547 **A.1 High-Level Implementation Architecture**

548 The prototype implementation network is a flat management network for the AMI components and Intel  
549 compute servers with a generic laptop to use as a workstation. The Intel servers utilize an overlay network  
550 for communication between the containers that run on top of them. Each Intel server has a socketed BIOS  
551 and Baseboard Management Controller (BMC) slot so that a BIOS and BMC chip could be flashed with  
552 AMI firmware and installed on the server. There are more technical details regarding the AMI  
553 components and Intel server configurations and installation in Appendices B and C. Table 2 displays the  
554 hostname and IP configuration of each system, and Figure 5 shows the logical network architecture of the  
555 implementation.

556 **Table 2: System Hostnames and IP Configurations**

Hostname	Management IP Address	BMC IP Address
AMI-TruE	172.16.100.50	N/A
AMI-HVS	172.16.100.51	N/A
Purley1	172.16.100.61	172.16.100.161
Purley2	172.16.100.62	172.16.100.162
Purley3	172.16.100.63	172.16.100.163
Workstation	172.16.100.21	N/A

557



558

559 **Figure 5: Logical Network Architecture of the Prototype Implementation**

560 **A.2 Intel Trusted Execution Technology (Intel TXT) & Trusted Platform Module (TPM)**

561 Hardware-based root-of-trust, when coupled with an enabled BIOS, OS, and components, constitutes the  
562 foundation for a more secure computing platform. This secure platform ensures BIOS and OS integrity at

563 boot from rootkits and other low-level attacks. It establishes the trustworthiness of the server and host  
564 platforms.

565 There are three roots of trust in a trusted platform: RTM, RTR, and RTS. They are the foundational  
566 elements of a single platform. These are the system elements that must be trusted because misbehavior in  
567 these normally would not be detectable in the higher layers. In an Intel TXT-enabled platform the RTM is  
568 the Intel microcode: the Core-RTM (CRTM). An RTM is the first component to send integrity-relevant  
569 information (measurements) to the RTS. Trust in this component is the basis for trust in all the other  
570 measurements. RTS contains the component identities (measurements) and other sensitive information. A  
571 trusted platform module (TPM) provides the RTS and RTR capabilities in a trusted computing platform.

572 Intel® Trusted Execution Technology (Intel® TXT) is the RTM, and it is a mechanism to enable  
573 visibility, trust, and control in the cloud. Intel TXT is a set of enhanced hardware components designed to  
574 protect sensitive information from software-based attacks. Intel TXT features include capabilities in the  
575 microprocessor, chipset, I/O subsystems, and other platform components. When coupled with an enabled  
576 OS and enabled applications, these capabilities provide confidentiality and integrity of data in the face of  
577 increasingly hostile environments.

578 Intel TXT incorporates a number of secure processing innovations, including:

- 579 • **Protected execution:** Lets applications run in isolated environments so that no unauthorized  
580 software on the platform can observe or tamper with the operational information. Each of these  
581 isolated environments executes with the use of dedicated resources managed by the platform.
- 582 • **Sealed storage:** Provides the ability to encrypt and store keys, data, and other sensitive  
583 information within the hardware. This can only be decrypted by the same environment that  
584 encrypted it.
- 585 • **Attestation:** Enables a system to provide assurance that the protected environment has been  
586 correctly invoked and to take a measurement of the software running in the protected space. The  
587 information exchanged during this process is known as the attestation identity key credential and  
588 is used to establish mutual trust between parties.
- 589 • **Protected launch:** Provides the controlled launch and registration of critical system software  
590 components in a protected execution environment.

591 Intel® Xeon® Platinum Scalable processor series and the previous generation Xeon Processor E3, Xeon  
592 Processor E5, and Xeon Processor E7 series processors support Intel TXT.

593 Intel TXT works through the creation of a measured launch environment (MLE) enabling an accurate  
594 comparison of all the critical elements of the launch environment against a known good source. Intel TXT  
595 creates a cryptographically unique identifier for each approved launch-enabled component and then  
596 provides a hardware-based enforcement mechanism to block the launch of any code that does not match  
597 or, alternately, indicate when an expected trusted launch has not happened through a process of secure  
598 remote attestation. In the latter case, when an attestation indicates that one or more measured components  
599 in the MLE do not match expectations, orchestration of workloads can be prevented on the suspect  
600 platform, even though the platform itself still launches. This hardware-based solution provides the  
601 foundation on which IT administrators can build trusted platform solutions to protect against aggressive  
602 software-based attacks and to better control their virtualized or cloud environments. For additional  
603 information on TXT and other RTM technologies, see the NIST white paper draft, *Hardware-Enabled  
604 Security for Server Platforms: Enabling a Layered Approach to Platform Security for Cloud and Edge  
605 Computing Use Cases* [1].

606 **A.3 Attestation**

607 There are two main considerations for use cases to be instantiated and delivered in a cloud:

- 608 • How would the entity needing this information know if a specific platform has Intel TXT enabled
- 609 or if a specific server has a defined or compliant BIOS or OS running on it (i.e., can it be
- 610 trusted)?
- 611 • Why should the entity requesting this information (which, in a cloud environment, could be a
- 612 resource scheduler or orchestrator trying to schedule a service on a set of available nodes or
- 613 servers) trust the response from the platform?

614 An attestation authority provides the definitive answers to these questions. Attestation provides

615 cryptographic proof of compliance, utilizing the root of trust concept to provide actionable security

616 controls by making the information from various roots of trust visible and usable by other entities. Figure

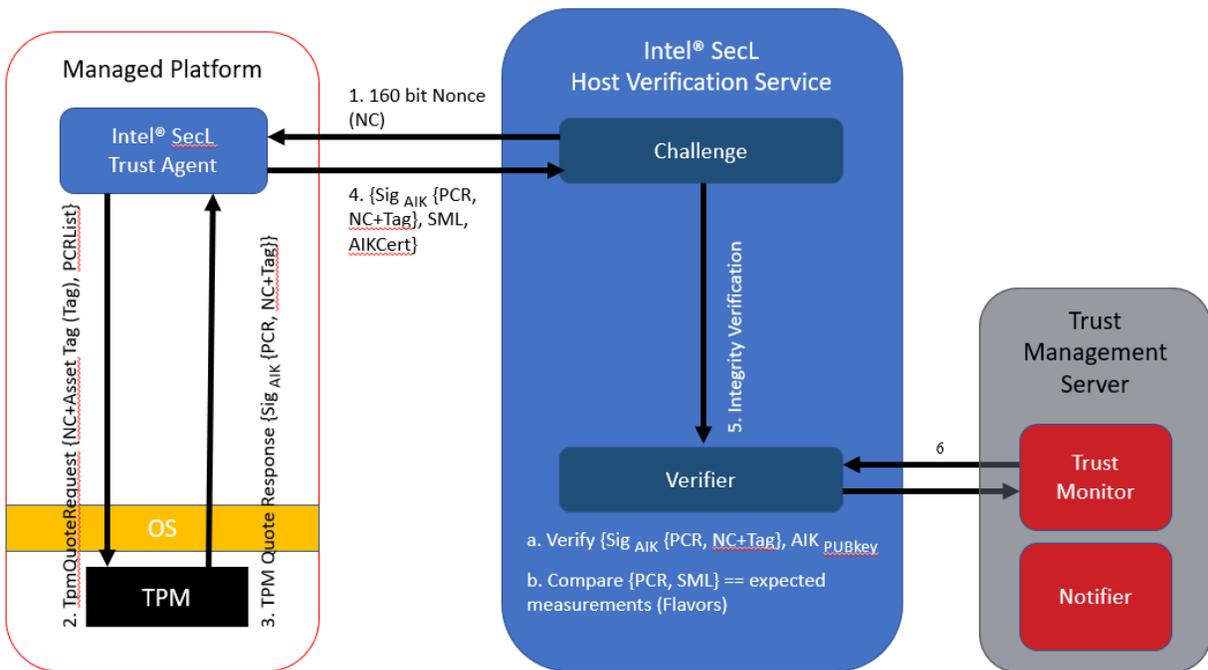
617 6 illustrates the attestation protocol providing the means for conveying measurements to the challenger.

618 The endpoint attesting device must have a means of measuring the BIOS firmware, low-level device

619 drivers, and OS and other measured components, and forwarding those measurements to the attestation

620 authority. The attesting device must do this while protecting the integrity, authenticity, nonrepudiation,

621 and in some cases, confidentiality of those measurements.



622 **Figure 6: Remote Attestation Protocol**

623 Here are the steps shown in Figure 6 for the remote attestation protocol:

- 625 1. The challenger, at the request of a requester, creates a non-predictable nonce (NC) and sends it to
- 626 the attestation agent on the attesting node, along with the selected list of Platform Configuration
- 627 Registers (PCRs).
- 628 2. The attestation agent sends that request to the TPM as a TPMQuoteRequest with the nonce and
- 629 the PCR List.

- 630 3. In response to the TPMQuote request, the TPM loads the attestation identity key from protected  
631 storage in the TPM by using the storage root key (SRK), and performs a *TPM Quote* command,  
632 which is used to sign the selected PCRs and the NC with the private key *AIKpriv*. Additionally,  
633 the attesting agent retrieves the stored measurement log (SML).
- 634 4. In the *integrity response* step, the attesting agent sends the response consisting of the signed  
635 quote, signed NC, and the SML to the challenger. The attesting agent also delivers the Attestation  
636 Identity Key (AIK) credential, which consists of the AIKpub that was signed by a privacy  
637 certificate authority (CA).
- 638 5. For the *integrity verification* step:
- 639 a. The challenger validates if the AIK credential was signed by a trusted Privacy-CA, thus  
640 belonging to a genuine TPM. The challenger also verifies whether AIKpub is still valid  
641 by checking the certificate revocation list of the trusted issuing party.
- 642 b. The challenger verifies the signature of the quote and checks the freshness of the quote.
- 643 c. Based on the received SML and the PCR values, the challenger processes the SML,  
644 compares the individual module hashes that are extended to the PCRs against the “good  
645 known or golden values,” and recomputes the received PCR values. If the individual  
646 values match the golden values and if the computed values match the signed aggregate,  
647 the remote node is asserted to be in a trusted state.
- 648 6. The verifier informs the trust state of the remote node to the manager. The manager records the  
649 trust state in its management database and uses it for any individual or aggregated device status  
650 requests. If an administrator subscribed to trust-related events, the manager will also send email  
651 notifications when a managed remote node is detected as being untrusted.

652 This protocol can help mitigate replay attacks, tampering, and masquerading.

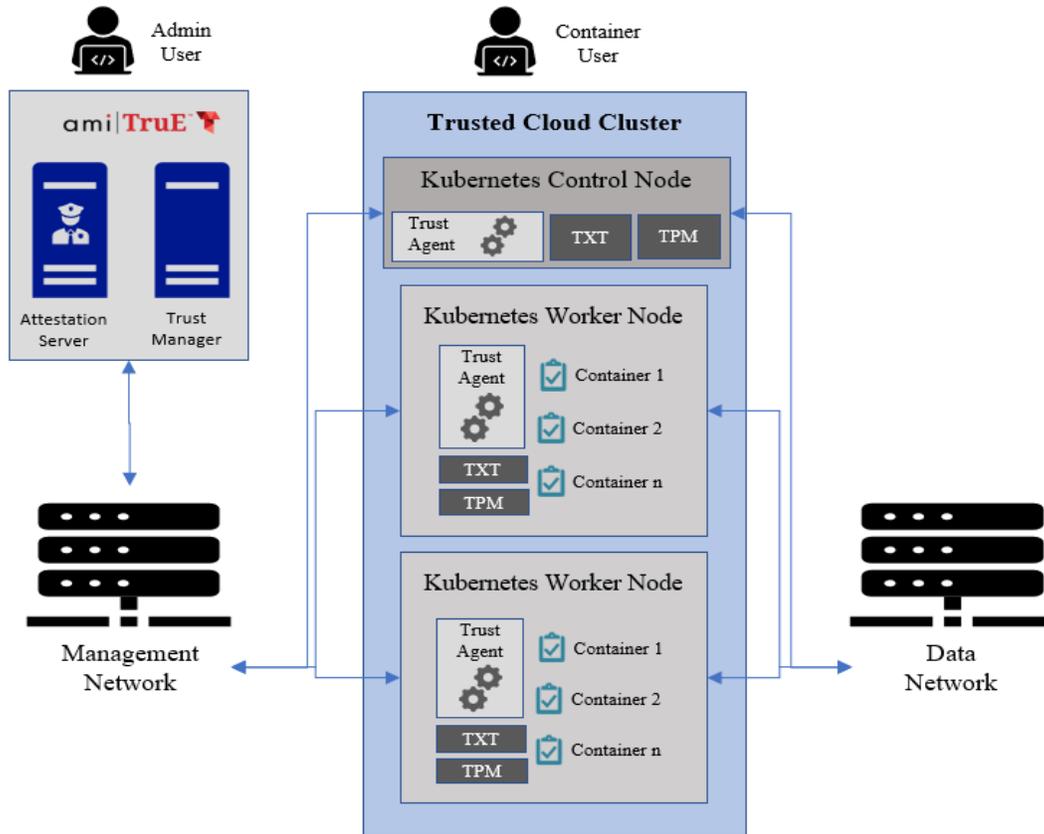
653

654 **Appendix B—Platform Implementation: AMI TruE**

655 This section contains supplementary information provided by AMI Trusted Environment (AMI TruE)  
656 describing all the components and steps required to set up the prototype implementation [5].

657 **B.1 Solution Architecture**

658 Figure 7 shows the architecture depicted in Appendix A, but with the specific products used in the AMI  
659 TruE platform implementation.



660  
661 **Figure 7: AMI TruE Prototype Implementation**

662 **B.2 Hardware Description**

663 The implemented architecture is composed of two Intel Next Units of Computing (NUCs) acting as  
664 Management Nodes and one Intel NUC together with two Intel TXT-enabled Server Platforms serving as  
665 the Trusted Cloud Cluster. Their hardware is detailed in Table 3.

666 **Table 3: Hardware for Implemented Architecture**

Hardware	Processor	Memory	Disk Space	OS
One Intel NUC acting as 'Kubernetes Control Node'	Intel i5-7300U @ 2.60 gigahertz (GHz)	8 gigabytes (GB)	250 GB	Red Hat Enterprise Linux (RHEL) 8.1

Hardware	Processor	Memory	Disk Space	OS
Two Server Platforms (with Intel TXT enabled) acting as 'Kubernetes Worker'	Intel Xeon Platinum 8260L @ 2.40 GHz	96 GB	250 GB	RHEL 8.1
Intel NUC acting as Trust Manager	Intel i5-7300U @ 2.60 GHz	8 GB	250 GB	RHEL 7.6
Intel NUC acting as Attestation Server	Intel i5-7300U @ 2.60 GHz	8 GB	250 GB	RHEL 8.1

667

668 **B.3 AMI TruE Installation and Configuration**

669 AMI TruE provides datacenter security solutions using Intel security technologies and libraries. With  
 670 AMI TruE, datacenters can achieve a level of trust and security. The following is a list of high-level  
 671 features offered by AMI TruE to manage and secure servers. Some of these features are discussed in more  
 672 detail later in this section.

- 673 • Automatic discovery of servers
- 674 • Asset inventory information collection
- 675 • Server health monitoring
- 676 • Trust status monitoring for all discovered servers
- 677 • TruE agent provisioning
- 678 • Remediation of untrusted servers
- 679 • Alert emails for health or trust events
- 680 • Remote power control
- 681 • Remote console (keyboard, video, mouse [KVM] redirection)
- 682 • BIOS/BMC firmware configuration and update
- 683 • OS and software provisioning
- 684 • Hypertext Markup Language version 5 (HTML5) based web interface
- 685 • Representational State Transfer (REST) APIs for automation or integration

686 AMI TruE has two components, Trust Manager and Attestation Server, so it requires two physical or  
 687 virtual systems to deploy AMI TruE. Table 4 specifies the minimum requirements for those systems.

688 **Table 4: Minimum System Requirements to Install AMI TruE**

System Element	Trust Manager	Attestation Server
<b>Processor</b>	4-core 2.66 GHz central processing unit (CPU)	4-core 2.66 GHz CPU
<b>Memory</b>	8 GB	8 GB
<b>Disk Space</b>	100 GB	100 GB
<b>OS</b>	RHEL 7.6, 64-bit	RHEL 8.1, 64-bit

689

690 **B.3.1 Installing AMI TruE Trust Manager**

691 To install the Trust Manager onto its system, perform the following steps:

- 692 1. Log into the Trust Manager system as a root user.

693 2. Download and extract the “amitrue\_trustmanager\_artifacts.zip” file into the “/root” folder.

694 3. Run the commands below as root user:

695 a. Set execution permission for the install script:

```
696 # chmod +x ./install.sh
```

697 b. Install AMI TruE Trust Manager by running the following install script:

```
698 #./install.sh
```

### 699 **B.3.2 Installing AMI TruE Attestation Server**

700 To install the Attestation Server onto its system, perform the following steps:

701 1. Log into the Attestation Server system as a root user.

702 2. Download and extract the “amitrue\_attestationserver\_artifacts.zip” file into the “/root” folder.

703 3. Edit the “amitrue\_security.env” file to configure the following:

```
704 HOSTNAME
```

```
705 IP_HOSTNAME_ARRAY
```

706 4. Run the commands below as root user:

707 a. Set execution permission for the install script:

```
708 # chmod +x ./install.sh
```

709 b. Install AMI TruE Attestation Server by running the following install script:

```
710 #./install.sh
```

### 711 **B.3.3 Configuring Firewall for AMI True**

712 AMI True uses several network ports for managing and securing platforms. The install script  
713 automatically configures the firewall to allow these ports. Ensure that no other software or utility disables  
714 any of the ports listed in Table 5 and Table 6.

715

**Table 5: Network Ports for Trust Manager**

Port	Purpose	Direction
9090	HTTP port for NGINX	Inbound/outbound
9443	HTTPS port for NGINX	Inbound/outbound
6379	Redis Database	Internal
5432	PostgreSQL Database	Internal
1900	Simple Service Discovery Protocol (SSDP) Discovery Module	Inbound/outbound
25/625	Core Notification Service (Simple Mail Transfer Protocol [SMTP])	Outbound
9089	Core Service Manager	Internal
9075	Core Discovery	Internal
9065	Core Platform Security	Internal
9055	Core API Server	Internal
9080	Redfish Server	Internal
9091	Server Manager – KVM	Inbound/outbound

716

**Table 6: Network Ports for Attestation Server**

Port	Purpose	Direction
5432	PostgreSQL Database	Internal
8443	Host Verification Service (HVS)	Inbound/outbound
8444	Authentication and Authorization Service (AAS)	Inbound/outbound
8445	Certificate Management Service (CMS)	Inbound/outbound
5000	Workload Service (WLS)	Inbound/outbound
9443	Key Management Service (KMS)	Inbound/outbound
1443	Trust Agent (TA)	Inbound/outbound
19082	Integration Hub (HUB)	Inbound/outbound
19445	Integration Hub (HUB)	Inbound/outbound

717

**718 B.3.4 Configuring Device Access Keys**

719 AMI TruE needs credentials in order to securely communicate with discovered and manageable devices.  
720 To configure these access keys, follow these steps:

- 721 1. Under the “Settings” submenu in the main menu, choose “Authentication Keys.”
- 722 2. On the Keys page, use “Add” or “Edit” to add access credentials for different types of resources.

**723 B.3.5 Configuring Discovery Range and Manageability Range**

724 To enable AMI TruE to scout and discover devices in a network, it needs to be configured with IP address  
725 ranges. Use the following steps to configure the discovery range:

- 726 1. Click on the hamburger menu icon  in the top left corner.
- 727 2. Under the “Settings” menu group, click “Discovery Settings.”
- 728 3. Select the “Global” tab under the “Discovery Ranges” section.
- 729 4. Click the “Add” button on the right side of the page to add a new discovery range.

730 You may not want to manage all discovered devices. A manageable device range can be configured so  
731 that AMI TruE will manage only devices that fall within that range. Use the following steps to configure a  
732 manageability range:

- 733 1. On the “Discovery Settings” page, under the “Discovery Ranges” section, select the  
734 “Manageability” tab.
- 735 2. Click the “Add” button to add a manageability range. Figure 8 shows a sample set of ranges.

Range	↕	Range Type	↕	Effective Range	↕	Manageability
10.2.0.0/15		CIDR		10.2.0.0 - 10.3.255.255		rmm
10.2.0.0/15		CIDR		10.2.0.0 - 10.3.255.255		rss
10.2.1.0/15		CIDR		10.2.0.0 - 10.3.255.255		fpx
10.2.1.0/15		CIDR		10.2.0.0 - 10.3.255.255		psme
10.2.1.0/24		CIDR		10.2.1.0 - 10.2.1.255		osm
10.2.3.3		Static		10.2.3.3		redfish
10.2.3.4		Static		10.2.3.4		redfish

736

737

Figure 8: Examples of Manageability Ranges

## 738 B.4 Trusted Cloud Cluster Installation and Configuration

739 All three nodes in the Trusted Cloud Cluster need to be configured to be managed and secured by AMI  
 740 TruE. This includes configuring BIOS settings, installing the TruE agent, and registering those agents  
 741 with AMI TruE. You can do these operations either remotely using AMI TruE or manually. See Appendix  
 742 B.4.1 for the remote instructions and Appendix B.4.2 for the manual instructions.

### 743 Prerequisites for being secured by AMI TruE

744 To be attested and be monitored for trust status, managed platforms should have:

- 745 • Intel® Xeon® or Intel® Xeon® Scalable Family processor that supports Intel TXT
- 746 • TPM (version 1.2 or 2.0) installed and provisioned for use with Intel TXT, according to Trusted  
 747 Compute Group specifications. If a version 2.0 TPM will be used, the Secure Hash Algorithm  
 748 256-bit (SHA256) PCR bank must be enabled.
- 749 • TPM and Intel TXT enabled in the BIOS
- 750 • TPM ownership cleared before installation

751 To be remediated or recovered from trust compromises, managed platforms should have:

- 752 • BMC with Redfish support
- 753 • BIOS with Redfish Host Interface support
- 754 • Secure shell (SSH) enabled in the host OS

### 755 B.4.1 Provisioning TruE Agent Remotely

756 To provision the TruE agent remotely using AMI TruE, follow these steps:

- 757 1. Log into the web interface of AMI TruE.
- 758 2. Under the “Provisions” menu, use the “Images” option to go to the “Images List” page.
- 759 3. Click on the “Add Image” option in the top menu. Enter the details about the image as depicted in  
 760 Figure 9. When finished, click the “Save” button.

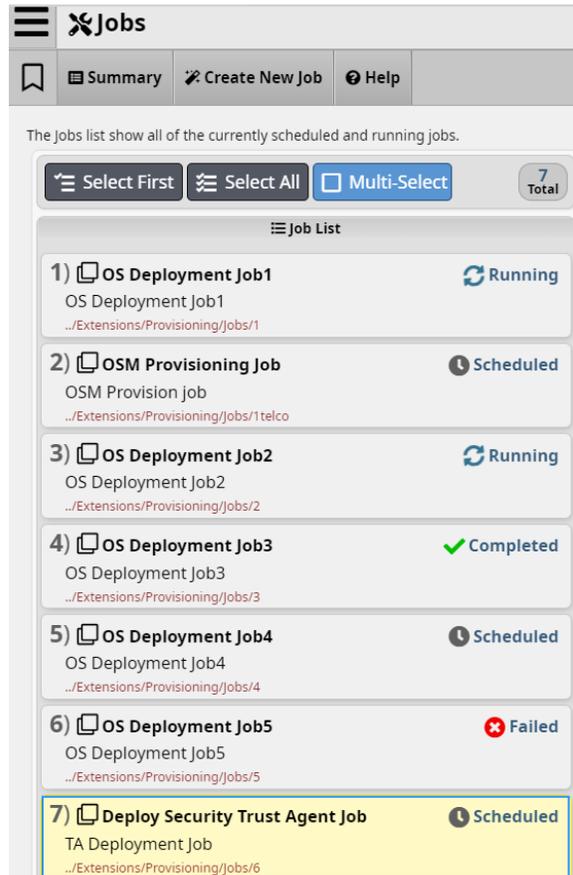
Name	TA-RHEL-1	Required
Description	Image with Trust Agent for RHEL platforms	
Version	1.0	Required
Server IP Address	1.2.3.4	Required
Image Type(required)	Trust Agent Deployment	Required
Path	(Not Selected)	Required
File Name	BIOS Update	Required
Share Type	BMC Update	Required
	NSD Deployment	Required
	NST Deployment	
	OS Deployment	
	OSM Deployment	
	Trust Agent Deployment	

761  
762

Figure 9: Example of Adding an Image

763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788

- Go to the “Create Job Wizard” page by choosing the “Create Job” option under the “Provisions” menu. Once there, enter a name and description for the job that is getting created, and choose “TA Deployment” as the job type.
- Click “Next” to move to the “Select Image” page. This page lists all images that were previously uploaded by the administrator. Select the image that needs to be used for deployment.
- Click “Next” to go to the “Targets” page. Select one or more platforms from the list of target platforms where the trust agent needs to be deployed.
- Go to the “Schedule” page by clicking “Next”. You can either opt to perform the deployment now or schedule it to be performed at a future date and time. Set this to the desired option for this deployment.
- Click the “Next” button to go to the “Review Job Settings” page. This page summarizes the information entered for this specific job. Use the “Previous” button to go back to any of the pages in the wizard and make changes if required. When you are satisfied with the settings, click “Start” to initiate the deployment process.
- To view the status of the job, click the “Jobs” option in the “Provisioning” menu. This page lists all scheduled jobs with their current statuses, as the example in Figure 10 shows. If needed, a scheduled job can be canceled by using the “Cancel Job” button.



**B.4.2 Provisioning TruE Agent Manually**

To provision the TruE agent manually, follow these steps:

- It is mandatory to register the RedHat subscription manager. Use the following instructions to register:

792  
793  
794

Figure 10: Example of Job List

- 795           a. Create a RedHat account.
- 796           b. Make sure the system has access to the internet.
- 797           c. Start the terminal access by logging in as root user.
- 798           d. Type “subscription-manager-gui” and the Subscription Manager window will pop open.
- 799           e. Click on the “Register” button and provide access credentials to register.
- 800       2. Disable the firewall on the target system by running the following commands:
- 801           \$ sudo systemctl stop firewalld
- 802           \$ sudo systemctl disable firewalld
- 803       3. Go to the BIOS settings. Enable TPM2 and clear ownership.
- 804       4. If you want to use the Unified Extensible Firmware Interface (UEFI) SecureBoot option for
- 805           trusted boot, enable it in the BIOS settings and skip the next step, otherwise you will use tboot
- 806           and should not enable SecureBoot.
- 807       5. If you want to use tboot, perform these steps:
- 808           a. Run this command to install tboot:
- 809               # yum install tboot
- 810           b. Make a backup of the current “grub.cfg” file:
- 811               # cp /boot/grub2/grub.cfg /boot/grub2/redhat/grub.bak
- 812           c. Generate a new “grub.cfg” file with the tboot boot option:
- 813               # grub2-mkconfig -o /boot/grub2/grub.cfg
- 814           d. Update the default boot option. Ensure that the GRUB\_DEFAULT value is set to the
- 815               tboot option. The tboot boot option can be found by looking in the
- 816               “/boot/redhat/grub.cfg” file.
- 817       6. Reboot the system. Because measurement happens at system boot, a reboot is needed to boot to
- 818           the tboot boot option and populate measurements in the TPM.
- 819       7. To verify a successful trusted boot with tboot, run the `txt-stat` command to show the tboot
- 820           log. Verify if the output shows “TRUE” for both “TXT measured launch” and “secrets flag set.”
- 821               \*\*\*\*\*
- 822               TXT measured launch: TRUE
- 823               secrets flag set: TRUE
- 824               \*\*\*\*\*
- 825       8. Install AMI TruE agents by following these steps:
- 826           a. Log in as a root user, and run all the commands below as root user.
- 827           b. Copy the “pkgs/platform-security-agent-artifacts.zip” file into the “/root” folder.
- 828           c. Extract the artifacts into the “/root” folder.
- 829           d. From “platform-security-agent-artifacts,” copy the “install\_agent.sh” and
- 830               “install\_agent.env” files into the “/root” folder.
- 831           e. Configure the “install\_agent.env” file as follows:

- 832                   i. HOSTNAME should not be empty. This variable is to set up the hostname for the  
833                   system (e.g., demo, demo-name-one, demo2).
- 834                   **Note:** Do not use an underscore as part of HOSTNAME.
- 835                   ii. IP\_HOSTNAME\_ARRAY should not be empty. Users need to provide the IP  
836                   and HOSTNAME pairs with space separation. This variable will edit the  
837                   “/etc/hosts” file with given IPs and hostnames.
- 838                   For example: IP\_HOSTNAME\_ARRAY=(10.0.0.0 demo 10.0.0.1  
839                   demo-name-one 10.0.0.2 demo2)
- 840                   **Note:** First give the IP and then give the hostname for that IP. Be careful not to  
841                   mismatch the IPs and hostnames.
- 842                   iii. Replace all the instances of “127.0.0.1” with the system IP, all instances of  
843                   “localhost” in URLs by the system IP, and all other instances of “localhost” by  
844                   the system hostname.
- 845                   iv. Generate BEARER\_TOKEN and CMS\_TLS\_CERT\_SHA384 by using the  
846                   service installation script with the “populate-users” option on the machine where  
847                   all the services are running. Then copy and paste the values in the env file.
- 848                   v. For “NO\_PROXY” and “HTTPS\_PROXY”, provide registry\_ip and proxy\_url,  
849                   respectively.
- 850                   **Note:** Do not use unnecessary spaces in the env file.
- 851                   **Note:** These are basic guidelines. If users have the proper knowledge about these  
852                   env variables, then they can modify the env variables according to their need.
- 853                   f. Set execution permission for the install script.  
854                   # chmod +x install\_agent.sh
- 855                   g. Run the install script.  
856                   # ./install\_agent.sh install
- 857                   h. The script will show some dependency messages and ask whether or not to continue  
858                   installation. Enter “yes” to continue installation.
- 859                   i. After TA installation, the script will restart the system. After restart, log in as a root user  
860                   and run the script again.  
861                   # ./install\_agent.sh install
- 862                   j. After the Workload Agent installation, the script will restart the system again to complete  
863                   the installation.

## 864 **B.5 Using AMI TruE**

865 Once installed and configured, AMI TruE starts discovering and monitoring the health and trust status of  
866 all managed platforms. This section explains the features offered by AMI TruE to monitor and secure the  
867 platforms.

### 868 **B.5.1 Monitoring Trust Status using AMI TruE**

869 Connect to AMI TruE from any standard web browser and use the credentials to log into its web  
870 interface. You will see a dashboard with several widgets. The dashboard can also be reached from any

871 other page in the web interface by clicking on the hamburger menu icon in the top left corner and  
872 selecting the “Dashboard” menu.

873 Figure 11 shows a dashboard with a chart reflecting the relative size of the pools of trusted (green),  
874 untrusted (red), and unknown (gray) cloud servers. In this example, there are eight servers in the trusted  
875 pool and two servers in the untrusted pool. More detailed trust information, including the security state of  
876 nodes, flavor-wise trust status, etc., is also depicted in this dashboard.



877

878

**Figure 11: Example of AMI TruE Report Dashboard**

879 Clicking on the “Go to Hosts Collection” link in the top right corner of the dashboard switches to a more  
880 detailed management page. Clicking on one of the servers listed in the “Hosts” page displays trust  
881 information for that selected server, as Figure 12 depicts. Details provided include the hostname,  
882 hardware UUID, flavor details, policy tags, connection URL, and component trust status.

The screenshot displays a web interface for server trust information. At the top, the host name is 10.2.0.0 with a UUID of b92c44c1-2c9d-4a7c-a425-0b9c5f87acbd. The interface is divided into several sections:

- Summary Information:** Lists host details such as ID (a1e563a8-ebb8-4d75-89b8-cb1b73481b9e), Hardware UUID (b92c44c1-2c9d-4a7c-a425-0b9c5f87acbd), TLS Policy ID (b10a6f24-2a89-456d-b43c-0ce3f59f505b), Flavor Group Name (automatic), and Connection String (https://10.2.0.0:1443).
- Host Trust Status:** Shows the host is CONNECTED with an Overall Trust Status of Trusted. Individual components like Platform, Host Unique, Software, Asset Tag, and OS are also marked as Trusted.
- I-Hub Host Information:** Provides details like I-Hub Host ID (99933eee-2888-457e-8e87-3e07402c2449), Created By (admin), Created Date (2018-06-07T13:37:19-0700), Modified by (admin), Modified Date (2018-06-08T09:11:50-0700), Valid To (2018-06-08T17:11:15.079Z), and a REST Browser Link.
- Asset Tags:** Shows the Deployment Status as DEPLOYED with Flavor ID 8d982427-ee29-461d-a83f-f81160aebef1. Tags include Country (USA), Department (Finance), and Compliance (PCI).
- Flavor Group:** Name is automatic (ID: 826501bd-3c75-4839-a08f-db5f744f8498). It contains four Flavor Group Info tables:
 

Flavor Group Info		Flavor Group Info		Flavor Group Info		Flavor Group Info	
Flavor Part	PLATFORM	Flavor Part	OS	Flavor Part	ASSET_TAG	Flavor Part	HOST_UNIQUE
Match Type	ANY_OF	Match Type	ANY_OF	Match Type	ANY_OF	Match Type	ANY_OF
Required	REQUIRED	Required	REQUIRED	Required	REQUIRED_IF_DEFINED	Required	REQUIRED_IF_DEFINED
- Report:** Shows a report ID (0a3cd9df-a713-4b3f-8b54-c6baea9d16e0) and Host ID (a1e563a8-ebb8-4d75-89b8-cb1b73481b9e).

883

884

Figure 12: Example of Summary Page for Server Trust Information

885

### B.5.2 Generating Trust Reports

886

To generate a new report for a platform, select that platform from the host list and click the “Create

887

Report” option. The retrieved report is presented as part of the host information, as shown in Figure 13.

The screenshot displays a 'Report' window with the following details:

- Id:** 0a3cd9df-a713-4b3f-8b54-c6baea9d16e0
- Host Id:** a1e563a8-ebb8-4d75-89b8-cb1b73481b9e
- Overall:** 🚫 false
- Created:** 2020-06-18T14:56:46-0400
- Expiration:** 2020-06-19T14:56:46-0400

Below the report details is the 'Flavor Trust Information' section, which includes tabs for 'Host Unique', 'Software', 'OS', 'Platform' (selected), and 'Asset Tag'. Three rules are listed:

- Rule 1:**
  - Name: com.intel.mtwilson.core.verifier.policy.rule.PcrMatchesConstant
  - Digest Type: com.intel.mtwilson.core.common.model.PcrSha256
  - Pcr Index: pcr\_3
  - Pcr Bank: SHA256
  - Flavor Id: e22dc729-e1dc-4f53-a386-8c26e89ab1ad
  - Trusted: 🔒 true
  - Markers: PLATFORM
- Rule 2:**
  - Name: com.intel.mtwilson.core.verifier.policy.rule.PcrMatchesConstant
  - Digest Type: com.intel.mtwilson.core.common.model.PcrSha1
  - Pcr Index: pcr\_6
  - Pcr Bank: SHA1
  - Flavor Id: e22dc729-e1dc-4f53-a386-8c26e89ab1ad
  - Trusted: 🔒 true
  - Markers: PLATFORM
- Rule 3:**
  - Name: com.intel.mtwilson.core.verifier.policy.rule.AikCertificateTrusted
  - Trusted: 🔒 true
  - Markers: PLATFORM

888

889

**Figure 13: Example of Trust Report**

890 To view the raw JavaScript Object Notation (JSON) data for any analytic needs, click the “Data” option  
891 in the top right corner of the “Report” window.

### 892 **B.5.3 Tagging Platforms Using AMI TruE**

893 Asset tags are used to tag managed platforms with one or more user-defined attributes, such as asset tag,  
894 compliance information, or customer type. This enables policy-based workload placement and  
895 orchestration.

896 AMI TruE provides options to create and deploy new asset tags to one or more managed platforms. It also  
897 has provisions to revoke any previously created asset tag. To create and deploy an asset tag:

- 898 1. Go to the “Hosts” page by opening the main menu and choosing the “Hosts” menu item under the  
899 “Security” menu group.
- 900 2. Select one or more servers in the host list to choose the platforms for which the asset tag needs to  
901 be deployed.
- 902 3. Click on the “Asset Tag” menu on the right side of the page, and choose the “Create and Deploy”  
903 option as shown in Figure 14.

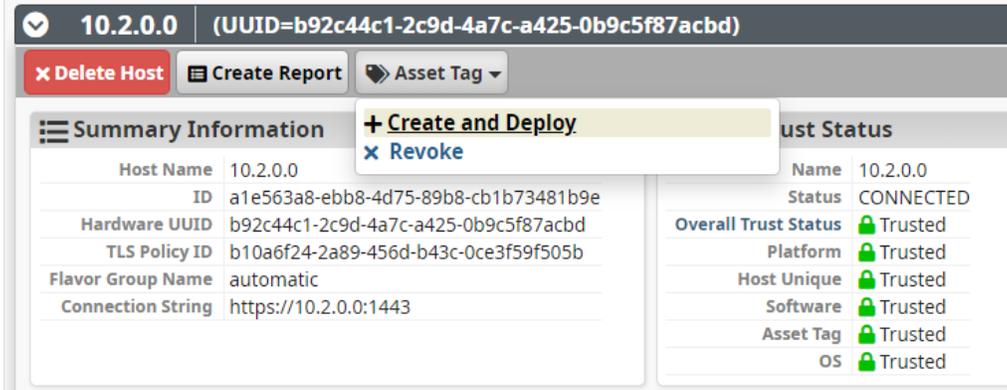


Figure 14: Example of Creating and Deploying an Asset Tag

904  
905

4. An “Asset Tag Creation” window should open.
  - a. To add an asset tag to the list, enter the asset tag name and value in the “Key Name” and “Key Value” fields, then click the “Add” button.
  - b. To remove a specific tag from the list, click on the tag’s “X” button in the list of asset tags.
  - c. After adding one or more asset tags, click the “Deploy” button to deploy all listed asset tags to the selected platforms.

906  
907  
908  
909  
910  
911  
912

Alternatively, you can click on the “Asset Tag” menu item in the top menu while being on the “Hosts” page to launch the “Asset Tag” page. That page has options to filter platforms using the “Search” option, and you could then use the filtered list of platforms to deploy an asset tag, as explained above.

## 916 B.5.4 Receiving Trust Event Alert Notification

917 Being notified when a platform turns untrusted allows administrators to quickly take remediation actions.  
918 AMI TruE provides two modes of notifications for any security events: email alerts and event log entries.

### 919 B.5.4.1 Email Alerts

920 Administrators can opt to receive email alerts on security events. This starts with adding information  
921 about one or more email servers that can be used to email alert information.

- 922 1. To add an email server, select the main menu (hamburger icon) icon in the top left corner and  
923 choose “Email Notifications” under the “Notifications” menu group.
- 924 2. Select “Configure Email Servers” in the resulting menu. This presents a list of configured email  
925 servers and provides options to add a new email server configuration, modify an existing email  
926 server configuration, or delete a previously configured email server.
- 927 3. Click “Add” to add a new entry, or select a row and click “Edit” to edit an existing entry. Enter  
928 the details of the email server and choose “Save.”
- 929 4. Once at least one email server is configured, the next step is to add email addresses of  
930 administrators or support engineers who need to be notified on any trust events. To view the  
931 email address book, click the main menu (hamburger icon) icon in the top left corner and select  
932 “Email Notifications” under the “Notifications” menu group. Choose “Email Address Book” in  
933 the resulting menu. Use the “Add,” “Edit,” and “Delete” options in the “Email Addresses” tab to

- 934 manage email addresses. Also, email groups can be used to notify a group of administrators on  
 935 any specific event. Use the “Email Groups” tab in the “Email Address Book” page to manage  
 936 email groups.
- 937 5. After adding email addresses or groups, the next step is to configure notification subscriptions.  
 938 From the main menu, select “Notification Subscriptions” under the “Notifications” menu group.
- 939 6. Use the “Add New Subscriptions” option on the “Notification Subscriptions” page to configure  
 940 event subscriptions. From the “Add New Subscription” page, as shown in Figure 15, choose the  
 941 types of events and resources (event sources) for which notifications need to be sent. You can  
 942 choose “Security” as the resource to receive any trust-related events.

943

944

**Figure 15: Example of the "Add New Subscription" Page**

- 945 7. Next, click “Select Recipients” to add one or more email addresses or groups that need to be  
 946 notified. When done, click the “Save” button to add the subscription.

947 **B.5.4.2 Event Logs**

948 AMI TruE records all platform-related events, including security events, into an event log. Administrators  
 949 can view those events through a web interface.

- 950 1. To view event logs, select the main menu (hamburger icon) icon in the top left corner and choose  
 951 “Global Event Log” under the “Logs” menu group.
- 952 2. Once event logs are viewed and acted upon, administrators can delete the events using the “Clear  
 953 Log” option on the “Global Event Log” page.

954 **B.5.5 Using AMI TruE for Remediation**

955 Being able to remediate and recover completely is one of the key needs for platform resilience. AMI TruE  
 956 offers multiple options to recover a compromised platform.

957 **B.5.5.1 Remote Power Control**

958 AMI TruE provides remote power management features to either shut down/power off or reset/restart the  
 959 platform as part of remediation efforts.

1. Click the main menu icon in the top left corner and select “Server Summary” under the “Server Manager” menu. This page lists all managed servers in the left pane with icons depicting their trust, health, and power state. Select a server that needs to shut down or powered off. On the right pane, click on the “Actions” button and select the “Power Reset” option, as shown in Figure 16.

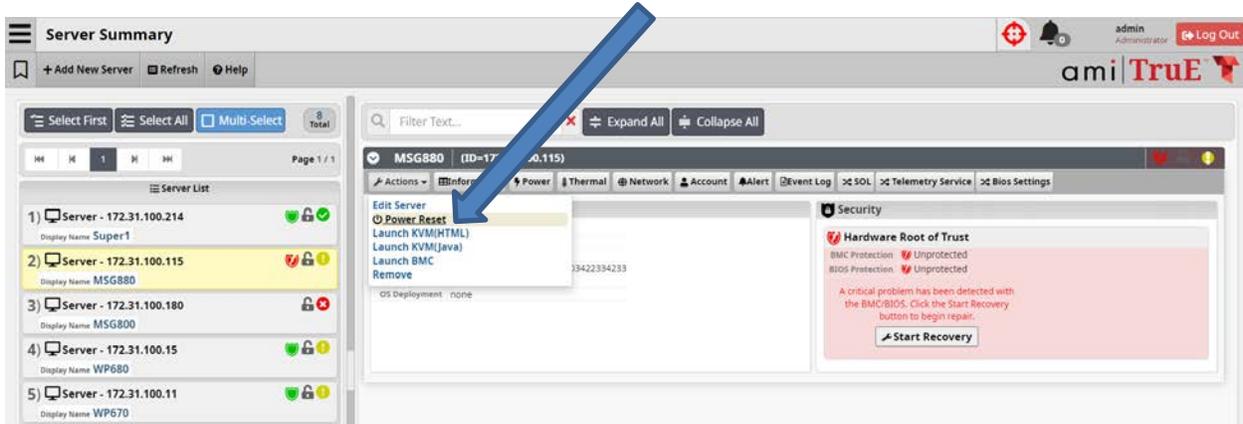


Figure 16: Example of Remote Power Control

2. A popup window with options to choose the type of power operation to be performed is presented. Select the appropriate power control operation and click the “OK” button to proceed.

### B.5.5.2 Remote Firmware Update

Using AMI TruE, BIOS/BMC configurations can be made, or the entire BIOS/BMC firmware can be updated if the firmware layer becomes untrusted.

1. To update either BIOS or BMC firmware, the first step is to upload the firmware images. Select the “Images” option under the “Provisions” menu view.
2. Click on the “Add Image” button in the top menu to add a new image for any provisioning need. Enter details on where the firmware image is located, version details, etc. to allow AMI TruE to filter and present matching images for a specific provisioning task.
3. Once a BMC or BIOS firmware image is uploaded, create a job to update the firmware either immediately or at a scheduled time. AMI TruE provides a wizard for creating a job to choose images, select target nodes, and either update them immediately or schedule the update for a future date and time. To create a job, click on “Create New Job” on the “Jobs” page.
4. Enter a name for the job, add a description, and choose either “BIOS Update” or “BMC Update” as the job type. Click “Next” to go the “Image” tab.
5. Select an image that needs to be used and click “Next” to go to the “Targets” page. This page lists all target platforms with a BIOS or BMC that can be updated, depending on the Job Type selected.
6. Choose one or more target platforms that need to be updated and select “Next” to go to the “Schedule” tab. This tab provides options to either run the job immediately or schedule it to be run at a future date and time. After configuring this, click “Next” to go to the “Review Job Settings” tab.
7. Review the information entered for this update job. Use the “Previous” button to navigate to other tabs in this wizard to change any data, if needed. When ready, click “Start” to start or schedule the update job.

- 992 8. To know the status of a scheduled task, choose the “Jobs” option under the “Provisioning” menu.  
993 This lists all running or scheduled jobs and provides the capability to cancel them if needed.

994 **B.5.5.3 Remote OS Installation**

995 If an OS is compromised, AMI TruE provides options to remotely re-install a version of the OS that is  
996 trusted by your enterprise or datacenter. Follow the steps in Appendix B.5.5.2, but choose “OS  
997 Deployment” in the “Job Type” field.

998

## 999 **Appendix C—Platform Implementation: Kubernetes**

1000 This section contains supplementary information describing all the required components and steps  
1001 required to set up the prototype implementation for Kubernetes.

### 1002 **C.1 Prototype Architecture**

1003 Refer to Figure 7 from Appendix B.1 for the relevant architecture diagram.

### 1004 **C.2 Hardware Description**

1005 Refer to the hardware descriptions from Appendix B.2 that are used for the Kubernetes setup.

### 1006 **C.3 Kubernetes Installation and Configuration**

1007 Kubernetes deployments minimally consist of master node(s) and worker node(s), which utilize a specific  
1008 container runtime. There is a common set of prerequisites that both types of nodes need, and there are  
1009 unique configurations for each type of node as well. This implementation installs Docker 19.3.5 as the  
1010 container runtime, and is running kubelet version 1.17.0 as its prerequisite.

#### 1011 **Prerequisite installation:**

1012 The following commands enable the network traffic overlays for communications between Kubernetes  
1013 pods within the cluster.

```
1014     # cat > /etc/modules-load.d/containerd.conf <<EOF
1015     overlay
1016     br_netfilter
1017     EOF
1018
1019     # modprobe overlay
1020     # modprobe br_netfilter
1021     # cat > /etc/sysctl.d/99-kubernetes-cri.conf <<EOF
1022     net.bridge.bridge-nf-call-iptables = 1
1023     net.ipv4.ip_forward = 1
1024     net.bridge.bridge-nf-call-ip6tables = 1
1025     EOF
1026
1027     # sysctl --system
1028     # systemctl enable containerd
```

1030 The following commands add the Kubernetes repository needed for the software package installations.

```
1031     # cat <<EOF > /etc/yum.repos.d/kubernetes.repo
1032     [kubernetes]
```

```

1033     name=Kubernetes
1034     baseurl=https://packages.cloud.google.com/yum/repos/kubernetes-el7-
1035     x86_64
1036     enabled=1
1037     gpgcheck=1
1038     repo_gpgcheck=1
1039     gpgkey=https://packages.cloud.google.com/yum/doc/yum-key.gpg
1040     https://packages.cloud.google.com/yum/doc/rpm-package-key.gpg
1041     EOF
1042

```

1043 The following commands install and start the necessary Kubernetes software packages.

```

1044     # dnf install -y kubeadm-1.17.0 kubelet-1.17.0 kubect1-1.17.0
1045     # systemctl enable kubelet
1046     # echo 'KUBELET_EXTRA_ARGS="--fail-swap-on=false"' >
1047     /etc/sysconfig/kubelet
1048     #systemctl start kubelet
1049

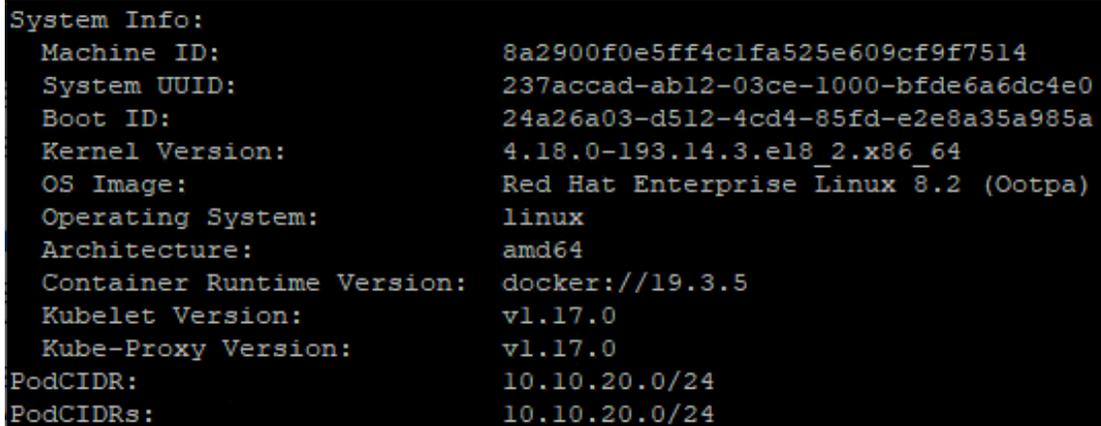
```

1050 Running the following command will produce the output shown in Figure 17:

```

1051     # kubect1 describe nodes

```



```

System Info:
Machine ID:      8a2900f0e5ff4c1fa525e609cf9f7514
System UUID:    237accad-ab12-03ce-1000-bfde6a6dc4e0
Boot ID:        24a26a03-d512-4cd4-85fd-e2e8a35a985a
Kernel Version: 4.18.0-193.14.3.el8_2.x86_64
OS Image:       Red Hat Enterprise Linux 8.2 (Ootpa)
Operating System: linux
Architecture:  amd64
Container Runtime Version: docker://19.3.5
Kubelet Version: v1.17.0
Kube-Proxy Version: v1.17.0
PodCIDR:        10.10.20.0/24
PodCIDRs:       10.10.20.0/24

```

1052

1053

Figure 17: Kubelet and Docker Versions

1054

### 1055 C.3.1 Kubernetes Control Node Configuration

1056 The following command is executed to complete the configuration of the Kubernetes master node in this  
 1057 implementation by establishing the control-plane overlay network.

```

1058     # kubeadm init --pod-network-cidr=10.10.20.0/24

```

1059 Run the following commands to enable the current user, root in this case, to use the cluster.

```
1060     # mkdir -p $HOME/.kube
1061     # sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
1062     # sudo chown $(id -u):$(id -g) $HOME/.kube/config
```

1063 Run the following command to enable Calico as the Container Network Interface (CNI).

```
1064     # kubectl apply -f https://docs.projectcalico.org/manifests/calico.yaml
```

### 1065 **C.3.2 Kubernetes Worker Configuration**

1066 Once the master node is up and running, worker nodes can be joined to the Kubernetes deployment. A  
1067 token is required to join the worker nodes to the Kubernetes cluster. Run the following on the master node  
1068 to obtain the token.

```
1069     # kubeadm token create --print-join-command
```

1070 The following command is executed to join worker nodes to the Kubernetes deployment, which is  
1071 obtained from the previous command.

```
1072     # kubeadm join 172.16.100.61:6443 --token <token redacted> --discovery-
1073     token-ca-cert-hash \
1074         sha256:<hash redacted>
```

1075 After the worker node is joined to the Kubernetes cluster, run the following command on the master to  
1076 verify the nodes in the cluster are ready:

```
1077     # kubectl get nodes
```

### 1078 **C.3.3 Kubernetes Orchestration**

1079 In order for the Kubernetes cluster to use the trust measurements and asset tags of hosts in its scheduling  
1080 policies, the Kubernetes master must be configured to communicate with the attestation hub service in the  
1081 AMI installation. There is an installation binary on the AMI host verification server at  
1082 /root/binaries/k8s/isecl-k8s-extensions-v2.0.0.bin that needs to be copied to and run on the Kubernetes  
1083 control node. This will create Custom Resource Definitions (CRDs) that allow the ISeCL trust  
1084 measurements and asset tags to be leveraged by the Kubernetes scheduler.

1085 A tenant must be created for the Kubernetes hosts, which they must be added to, in the AMI TruE web  
1086 client. When the tenant is created, the user can choose to associate hosts that are already in the host  
1087 verification service. After the hosts have been added into the tenant, their trust measurements and asset  
1088 tags can be used for Kubernetes scheduling policies. In order to enable the trust measurements being used  
1089 by the Kubernetes scheduler, a few modifications need to be made to the Kubernetes scheduler  
1090 configurations, and access keys need to be shared between the Kubernetes control node and the AMI  
1091 TruE host verification server. The AMI TruE host verification server is built from the ISeCL code base;  
1092 the full steps can be found in Section 6.13.3.2 in the ISeCL product documentation at  
1093 [https://01.org/sites/default/files/documentation/intelr\\_secl-dc\\_v1.6\\_ga\\_product\\_guide\\_1.pdf](https://01.org/sites/default/files/documentation/intelr_secl-dc_v1.6_ga_product_guide_1.pdf).

1094 **Appendix D—Supporting NIST SP 800-53 Security Controls**

1095 The major controls in the NIST SP 800-53 Revision 5, *Security and Privacy Controls for Information*  
 1096 *Systems and Organizations* [6] control catalog that affect the container platform security prototype  
 1097 implementation are:

- 1098 • AU-2, Event Logging
- 1099 • CA-2, Control Assessments
- 1100 • CA-7, Continuous Monitoring
- 1101 • CM-2, Baseline Configuration
- 1102 • CM-3, Configuration Change Control
- 1103 • CM-8, System Component Inventory
- 1104 • IR-4, Incident Handling
- 1105 • SA-9, External System Services
- 1106 • SC-1, Policy and Procedures [for System and Communications Protection Family]
- 1107 • SC-7, Boundary Protection
- 1108 • SC-29, Heterogeneity
- 1109 • SC-32, System Partitioning
- 1110 • SC-36, Distributed Processing and Storage
- 1111 • SI-2, Flaw Remediation
- 1112 • SI-3, Malicious Code Protection
- 1113 • SI-4, System Monitoring
- 1114 • SI-6, Security and Privacy Function Verification
- 1115 • SI-7, Software, Firmware, and Information Integrity

1116 Table 7 lists the security capabilities provided by the trusted asset tag prototype:

1117 **Table 7: Security Capabilities Provided by the Trusted Asset Tag Prototype**

Capability Category	Capability Number	Capability Name
IC1 – Measurements	IC1.1	Measured Boot of BIOS
	IC1.2	Baseline for BIOS measurement (allowed list)
	IC1.3	Remote Attestation of Boot Measurements
	IC1.4	Security Capability & Config Discovery
IC2 – Tag Verification	IC2.1	Asset Tag Verification
IC3 – Policy Enforcement	IC3.1	Policy-Based Workload Provisioning
	IC3.2	Policy-Based Workload Migration
IC4 – Reporting	IC4.1	Support for Continuous Monitoring
	IC4.2	Support for On-Demand Reports

Capability Category	Capability Number	Capability Name
	IC4.3	Support for Notification of Trust Events
IC5 – Remediation	IC5.1	Remotely Powering Down a Compromised Platform
	IC5.2	Updating a Compromised BIOS Firmware
	IC5.3	Reinstalling a Compromised OS

1118

1119 Table 8 maps the security capabilities from Table 7 to the NIST SP 800-53 controls in the list at the  
1120 beginning of this appendix.

1121

**Table 8: Mapping of Security Capabilities to NIST SP 800-53 Controls**

NIST SP 800-53 Control	Measurements				Tag Veri- fication	Policy Enforcement		Reporting			Remediation		
	IC1.1	IC1.2	IC1.3	IC1.4		IC2.1	IC3.1	IC3.2	IC4.1	IC4.2	IC4.3	IC5.1	IC5.2
AU-2								X	X	X			
CA-2				X				X	X				
CA-7								X	X				
CM-2		X		X	X								
CM-3	X		X		X								
CM-8				X	X								
IR-4										X	X	X	X
SA-9						X	X						
SC-1						X	X						
SC-7	X			X		X	X						
SC-29						X	X						
SC-32					X	X	X						
SC-36					X	X	X						
SI-2											X	X	X
SI-3	X	X		X				X	X				
SI-4		X	X	X				X	X				
SI-6	X	X	X	X									
SI-7	X	X	X			X	X				X	X	X

1122

**Appendix E—Cybersecurity Framework Subcategory Mappings**

This appendix maps the major security features of the container platform security prototype implementation to the following Subcategories from the Cybersecurity Framework [7]:

- DE.CM-4: Malicious code is detected
- DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events
- ID.GV-1: Organizational cybersecurity policy is established and communicated
- ID.GV-3: Legal and regulatory requirements regarding cybersecurity, including privacy and civil liberties obligations, are understood and managed
- ID.RA-1: Asset vulnerabilities are identified and documented
- PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity
- PR.IP-3: Configuration change control processes are in place
- PR.IP-5: Policy and regulations regarding the physical operating environment for organizational assets are met
- PR.IP-12: A vulnerability management plan is developed and implemented
- PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy
- RS.MI-1: Incidents are contained
- RS.MI-2: Incidents are mitigated

Table 9 indicates the mappings from the security capabilities in Table 7 in the previous appendix to the Cybersecurity Framework Subcategories listed above.

**Table 9: Mapping of Security Capabilities to NIST Cybersecurity Framework Subcategories**

Cybersecurity Framework Subcategory	Measurements				Tag Verification	Policy Enforcement		Reporting			Remediation		
	IC1.1	IC1.2	IC1.3	IC1.4		IC2.1	IC3.1	IC3.2	IC4.1	IC4.2	IC4.3	IC5.1	IC5.2
DE.CM-4	X							X					
DE.CM-6	X		X			X	X	X	X	X			
ID.GV-1						X	X						
ID.GV-3					X	X	X						
ID.RA-1											X	X	X
PR.DS-6	X	X	X			X	X				X	X	X
PR.IP-3	X		X	X	X								
PR.IP-5					X	X	X						
PR.IP-12											X	X	X
PR.PT-1	X		X					X	X	X			
RS.MI-1											X	X	X
RS.MI-2												X	X

1145

## 1146 **Appendix F—Acronyms and Other Abbreviations**

1147 Selected acronyms and abbreviations used in the report are defined below.

1148	<b>AAS</b>	Authentication and Authorization Service
1149	<b>AIK</b>	Attestation Identity Key
1150	<b>AMI TruE</b>	AMI Trusted Environment
1151	<b>API</b>	Application Programming Interface
1152	<b>BIOS</b>	Basic Input/Output System
1153	<b>BMC</b>	Baseboard Management Controller
1154	<b>CA</b>	Certificate Authority
1155	<b>CMS</b>	Certificate Management Service
1156	<b>CNI</b>	Container Network Interface
1157	<b>CoT</b>	Chain of Trust
1158	<b>CPU</b>	Central Processing Unit
1159	<b>CRD</b>	Custom Resource Definition
1160	<b>CRTM</b>	Core Root of Trust for Measurement
1161	<b>FOIA</b>	Freedom of Information Act
1162	<b>GB</b>	Gigabyte
1163	<b>GHz</b>	Gigahertz
1164	<b>HTML5</b>	Hypertext Markup Language (version 5)
1165	<b>HVS</b>	Host Verification Service
1166	<b>IaaS</b>	Infrastructure as a Service
1167	<b>Intel TXT</b>	Intel Trusted Execution Technology
1168	<b>I/O</b>	Input/Output
1169	<b>IP</b>	Internet Protocol
1170	<b>IR</b>	Interagency or Internal Report
1171	<b>IT</b>	Information Technology
1172	<b>ITL</b>	Information Technology Laboratory
1173	<b>JSON</b>	JavaScript Object Notation
1174	<b>KMS</b>	Key Management Service
1175	<b>KVM</b>	Keyboard, Video, Mouse
1176	<b>MLE</b>	Measured Launch Environment
1177	<b>NC</b>	Nonce
1178	<b>NIST</b>	National Institute of Standards and Technology
1179	<b>NISTIR</b>	National Institute of Standards and Technology Interagency or Internal Report
1180	<b>NUC</b>	Next Unit of Computing
1181	<b>OEM</b>	Original Equipment Manufacturer
1182	<b>OS</b>	Operating System
1183	<b>PCR</b>	Platform Configuration Register
1184	<b>REST</b>	Representational State Transfer
1185	<b>RHEL</b>	Red Hat Enterprise Linux

1186	<b>RTM</b>	Root of Trust for Measurement
1187	<b>RTR</b>	Root of Trust for Reporting
1188	<b>RTS</b>	Root of Trust for Storage
1189	<b>SHA256</b>	Secure Hash Algorithm 256-bit
1190	<b>SML</b>	Stored Measurement Log
1191	<b>SMTP</b>	Simple Mail Transfer Protocol
1192	<b>SP</b>	Special Publication
1193	<b>SRK</b>	Storage Root Key
1194	<b>SSDP</b>	Simple Service Discovery Protocol
1195	<b>SSH</b>	Secure Shell
1196	<b>TA</b>	Trust Agent
1197	<b>TPM</b>	Trusted Platform Module
1198	<b>UEFI</b>	Unified Extensible Firmware Interface
1199	<b>URL</b>	Uniform Resource Locator
1200	<b>UUID</b>	Universally Unique Identifier
1201	<b>WLS</b>	Workload Service
1202		