

1
2

3
4

5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23

Hardware-Enabled Security:

Policy-Based Governance in Trusted Container Platforms

Michael Bartock
Murugiah Souppaya
Haidong Xia
Raghu Yeluri
Uttam Shetty
Brandon Lum
Mariusz Sabath
Harmeet Singh
Alaa Youssef
Gosia Steinder
Yu Cao
Jayashree Ramanathan

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

24
25

26
27
28
29

Hardware-Enabled Security:

Policy-Based Governance in Trusted Container Platforms

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

Brandon Lum
Mariusz Sabath
Harmeet Singh
Alaa Youssef
Gosia Steinder
*IBM
Armonk, New York*

Haidong Xia
Raghu Yeluri
Uttam Shetty
*Intel Corporation
Santa Clara, California*

Yu Cao
Jayashree Ramanathan
*Red Hat
Raleigh, North Carolina*

30
31
32
33
34
35
36

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

October 2021



U.S. Department of Commerce
Gina M. Raimondo, Secretary

37
38
39
40
41
42
43
44

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

45 National Institute of Standards and Technology Interagency or Internal Report 8320B
46 41 pages (October 2021)

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

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

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

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

62 **Public comment period: *October 27, 2021 through December 6, 2021***

63 National Institute of Standards and Technology
64 Attn: Applied Cybersecurity Division, Information Technology Laboratory
65 100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-2000
66 Email: hwsec@nist.gov

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

68

Reports on Computer Systems Technology

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

77

Abstract

78 In today’s cloud data centers and edge computing, attack surfaces have significantly increased,
79 cyber attacks are industrialized, and most security control implementations are not coherent or
80 consistent. The foundation of any data center or edge computing security strategy should be
81 securing the platform on which data and workloads will be executed and accessed. The physical
82 platform represents the foundation for any layered security approach and provides the initial
83 protections to help ensure that higher-layer security controls can be trusted. This report explains
84 an approach based on hardware-enabled security techniques and technologies for safeguarding
85 container deployments in multi-tenant cloud environments. It also describes a prototype
86 implementation of the approach intended to be a blueprint or template for the general security
87 community.

88

Keywords

89 cloud; container; hardware-enabled security; hardware root of trust; platform security; trusted
90 compute pool; virtualization.

91

Audience

92 The primary audiences for this report are security professionals, such as security engineers and
93 architects; system administrators and other information technology (IT) professionals for cloud
94 service providers; and hardware, firmware, and software developers who may be able to leverage
95 hardware-enabled security techniques and technologies to improve the containers deployment in
96 in multi-tenant cloud environments.

97

Trademark Information

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

99

Call for Patent Claims

100 This public review includes a call for information on essential patent claims (claims whose use
101 would be required for compliance with the guidance or requirements in this Information
102 Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be
103 directly stated in this ITL Publication or by reference to another publication. This call also
104 includes disclosure, where known, of the existence of pending U.S. or foreign patent applications
105 relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.
106

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

109 a) assurance in the form of a general disclaimer to the effect that such party does not hold
110 and does not currently intend holding any essential patent claim(s); or
111

112 b) assurance that a license to such essential patent claim(s) will be made available to
113 applicants desiring to utilize the license for the purpose of complying with the guidance
114 or requirements in this ITL draft publication either:
115

116 i. under reasonable terms and conditions that are demonstrably free of any unfair
117 discrimination; or
118

119 ii. without compensation and under reasonable terms and conditions that are
120 demonstrably free of any unfair discrimination.
121

122 Such assurance shall indicate that the patent holder (or third party authorized to make assurances
123 on its behalf) will include in any documents transferring ownership of patents subject to the
124 assurance, provisions sufficient to ensure that the commitments in the assurance are binding on
125 the transferee, and that the transferee will similarly include appropriate provisions in the event of
126 future transfers with the goal of binding each successor-in-interest.
127

128 The assurance shall also indicate that it is intended to be binding on successors-in-interest
129 regardless of whether such provisions are included in the relevant transfer documents.
130

131 Such statements should be addressed to: hwsec@nist.gov

132

133

Table of Contents

134 **1 Introduction 1**

135 1.1 Purpose and Scope 1

136 1.2 Terminology 1

137 1.3 Document Structure 2

138 **2 Prototype Implementation 3**

139 2.1 Objective 3

140 2.2 Goals 3

141 2.2.1 Stage 0: Platform attestation and measured worker node launch 4

142 2.2.2 Stage 1: Trusted placement of workloads 4

143 2.2.3 Stage 2: Asset tagging and trusted location 4

144 2.2.4 Stage 3: Trust-based workload encryption 4

145 2.2.5 Stage 4: Trust-based workload access to information 5

146 2.3 Additional Resources 5

147 **3 Prototyping Stage 0 7**

148 **4 Prototyping Stage 1 8**

149 **5 Prototyping Stage 2 9**

150 **6 Prototyping Stage 3 10**

151 6.1 Solution Overview 10

152 6.2 Solution Architecture 10

153 **7 Prototyping Stage 4 12**

154 7.1 Solution Overview 12

155 7.2 Solution Architecture 12

156 **References 14**

157 **Appendix A— Hardware Root of Trust Implementation 15**

158 A.1 High-Level Implementation Architecture 15

159 A.2 Hardware Root of Trust: Intel TXT and Trusted Platform Module (TPM) 16

160 A.3 Attestation: Intel Security Libraries (ISecL) 17

161 **Appendix B— Workload Orchestration Implementation: OpenShift 19**

162 B.1 Prototype Architecture 19

163 B.2 OpenShift Installation and Configuration 20

164 B.2.1 VMware-Based Management Cluster (Cluster A) 20

196 **1 Introduction**

197 **1.1 Purpose and Scope**

198 The purpose of this publication is to describe an approach for safeguarding application container
199 deployments in multi-tenant cloud environments. This publication builds upon selected security
200 challenges involving Infrastructure as a Service (IaaS) that are discussed in NIST Interagency or
201 Internal Report (IR) 8320A [1], which addresses cloud computing technologies and geolocation
202 in the form of resource asset tags. Specifically, it uses the three stages of deployment described
203 in Sections 3, 4, and 5 of NIST IR 8320A, and additionally describes two additional stages for
204 encrypting container images and creating data access policies for containers. It then describes a
205 prototype implementation that was designed to address those challenges. The publication
206 provides sufficient details about the prototype implementation so that organizations can
207 reproduce it if desired. The publication is intended to be a blueprint or template that can be used
208 by the general security community to validate and implement the described implementation.

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

214 This publication builds upon the terminology and concepts described in NIST IR 8320,
215 *Hardware-Enabled Security: Enabling a Layered Approach to Platform Security for Cloud and*
216 *Edge Computing Use Cases* [2]. Reading that NIST IR is a prerequisite for reading this
217 publication because it explains the concepts and defines key terminology used in this publication.

218 **1.2 Terminology**

219 For consistency with related NIST reports, this report uses the following definitions for trust-
220 related terms:

- 221 • **Trust:** “The confidence one element has in another that the second element will behave
222 as expected.” [3]
- 223 • **Trusted:** An element that another element relies upon to fulfill critical requirements on
224 its behalf.
- 225 • **Trusted boot:** A system boot where aspects of the hardware and firmware are measured
226 and compared against known good values to verify their integrity and thus their
227 trustworthiness.
- 228 • **Trustworthy:** Worthy of being trusted to fulfill whatever critical requirements may be
229 needed.¹

¹ Based on the definition from NIST Special Publication (SP) 800-160 Volume 2, <https://doi.org/10.6028/NIST.SP.800-160v2>

230 1.3 Document Structure

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

- 232 • Section 2 defines the objective for the prototype implementation and the intermediate
233 goals to be met in order to achieve the objective.
- 234 • Sections 3 through 7 describe the five stages of the prototype implementation:
 - 235 ○ Stage 0: Have assurance that the platform the container deployment is running on
236 can be trusted
 - 237 ○ Stage 1: Orchestrate the placement of workloads to launch only on trusted
238 platforms
 - 239 ○ Stage 2: Be able to continuously monitor and enforce asset tag restrictions
 - 240 ○ Stage 3: Enable end users to encrypt their workload images
 - 241 ○ Stage 4: Be able to grant workloads access to sensitive information via
242 authentication mechanisms
- 243 • The References section lists the references cited throughout the document.
- 244 • Appendix A provides an overview of the high-level hardware architecture of the
245 prototype implementation, as well as details on how Intel platforms implement hardware
246 modules and enhanced hardware-based security functions.
- 247 • Appendix B contains supplementary information provided by IBM and Red Hat
248 describing the components and steps required to set up the OpenShift and Multi-Cloud
249 Manager solutions.
- 250 • Appendix C contains supplementary information describing all the required components
251 and steps required to set up the workload encryption implementation.
- 252 • Appendix D contains supplementary information describing all the required components
253 and steps required to set up the prototype implementation for using the Trusted Service
254 Identity.
- 255 • Appendix E lists the major controls from NIST Special Publication (SP) 800-53 Revision
256 5 that affect the prototype implementation, as well as the security capabilities the
257 prototype provides, and then maps the prototype's capabilities to the NIST SP 800-53
258 controls.
- 259 • Appendix F maps the major security features of the prototype to Cybersecurity
260 Framework Subcategories.
- 261 • Appendix G contains a list of acronyms for the report.

262 **2 Prototype Implementation**

263 This section defines the prototype implementation. Section 2.1 presents the objective. Section 2.2
264 provides more details, outlining all of the intermediate goals that must be met in order to achieve
265 the desired prototype implementation. These requirements are grouped into five stages of the use
266 case, each of which is examined more closely in Sections 2.2.1 through 2.2.5, respectively.

267 **2.1 Objective**

268 There are security and privacy concerns with allowing unrestricted container deployment
269 orchestration. A common desire is to only use cloud servers physically located within the same
270 country as the organization, or physically located in the same country as the origin of the
271 information. Whenever multiple container deployments are present on a single cloud server,
272 there is a need to segregate those deployments from each other so that they do not interfere with
273 each other, gain access to each other's sensitive data, or otherwise compromise the security or
274 privacy of the containers. NIST IR 8320A, *Hardware-Enabled Security: Container Platform*
275 *Security Prototype* [1] provides an overview of challenges organizations may face when using
276 cloud-based container workloads, as well as techniques to improve the security of cloud
277 computing and accelerate the adoption of cloud computing technologies by establishing an
278 automated hardware root-of-trust method for enforcing and monitoring platform integrity and
279 geolocation restrictions for cloud servers.

280 The motivation behind this use case is to build upon the stages of NIST IR 8320A and
281 implement additional techniques that leverage hardware roots of trust in server platforms. The
282 integrity and location data of each host are leveraged in the orchestration and protection of
283 workloads, as well as providing workloads access to specific data. Workload orchestration can
284 ensure that containers are instantiated only on server platforms that meet trustworthiness
285 requirements and are in acceptable locations. Orchestration can also involve initial encryption of
286 container images and releasing the decryption keys only to trusted servers. Additionally, the
287 workloads themselves can be assigned identities based on these trusted attributes of the physical
288 servers they reside on and be granted access to sensitive information based on their identities.

289 **2.2 Goals**

290 Using trusted compute pools, described in NIST IR 8320A Sections 3 through 5, is a leading
291 approach to aggregate trusted systems and segregate them from untrusted resources, which
292 results in the separation of higher-value, more sensitive workloads from commodity application
293 and data workloads. The principles of operation are to:

- 294 1. Create a part of the cloud to meet the specific and varying security requirements of users.
- 295 2. Control access to that portion of the cloud so that the correct applications (workloads) get
296 deployed there.
- 297 3. Enable audits of that portion of the cloud so that users can verify compliance.

298 Once the trusted compute pools are created, additional techniques can be used to protect the
299 workloads that run on them, or the information that the workloads can access. These additional
300 principles are to:

- 301 4. Encrypt workload images and ensure only specific servers can decrypt them.
- 302 5. Ensure that only specific applications with location-based restriction enforcement can
- 303 access sensitive data.

304 These trusted compute pools allow IT to gain the benefits of the dynamic cloud environment

305 while still enforcing higher levels of protections for their more critical workloads.

306 The ultimate goal is to be able to use “trust” as a logical boundary for deploying cloud workloads

307 on server platforms within a cloud. This goal is dependent on smaller prerequisite goals

308 described as stages, which can be thought of as requirements that the solution must meet.

309 **2.2.1 Stage 0: Platform attestation and measured worker node launch**

310 A fundamental component of a solution is having some assurance that the platform the container

311 deployment is running on can be trusted. If the platform is not trustworthy, then not only is it

312 putting the tenant’s application and data at greater risk of compromise, but also there is no

313 assurance that the claimed asset tag of the cloud server is accurate. Having basic assurance of

314 trustworthiness is the initial stage in the solution.

315 NIST IR 8320A Section 2.2.1 describes the specific goals of Stage 0.

316 **2.2.2 Stage 1: Trusted placement of workloads**

317 Once stage 0 has been successfully completed, the next objective is to be able to orchestrate the

318 placement of workloads to launch only on trusted platforms. Workload placement is a key

319 attribute of cloud computing, improving scalability and reliability. The purpose of this stage is to

320 ensure that any server that a workload is launched on will meet the required level of security

321 assurance based on the workload security policy.

322 NIST IR 8320A Section 2.2.2 describes the specific goals of Stage 1.

323 **2.2.3 Stage 2: Asset tagging and trusted location**

324 This next stage builds upon stage 1 by adding the ability to continuously monitor and enforce

325 asset tag restrictions.

326 NIST IR 8320A Section 2.2.3 describes the specific goals of Stage 2.

327 **2.2.4 Stage 3: Trust-based workload encryption**

328 This next stage builds upon stage 2 and adds the ability for end users to encrypt their workload

329 images, which provides at-rest cryptographic isolation to help protect consumer data and

330 intellectual property. In order for a compute node to launch a workload instance from an

331 encrypted image, it will need to retrieve the image decryption key. The purpose of this stage is to

332 ensure that only compute nodes with acceptable platform trustworthiness and specific asset tags

333 will be provided the decryption keys for specific workload images.

334 Stage 3 includes the following prerequisite goals:

- 335 1. **Have trusted asset tag information for each trusted platform instance.** Essentially,
336 stage 2 has been completed, and the platform trust measurements and asset tag
337 information can be leveraged during workload deployment.
- 338 2. **Encrypt workload images and protect decryption keys in a key manager.** Decryption
339 keys are kept in a key manager so that authorized nodes in the trusted compute pool can
340 retrieve and launch appropriate instances of workload images.
- 341 3. **Release decryption keys for workload images only to servers with trusted platforms
342 and in trusted locations.** Decryption keys are only released to servers that have the
343 appropriate platform trustworthiness and are in allowed locations based on their asset
344 tags.

345 **2.2.5 Stage 4: Trust-based workload access to information**

346 The last stage builds upon stage 3 and adds the ability to grant workloads access to sensitive
347 information. The majority of workloads running in the cloud need some access to data sources or
348 other services, authenticating themselves using a password, application programming interface
349 (API) key, or certificate. Today, this is typically done through secrets which are designed to be
350 stored securely. The purpose of this stage is to ensure that only specific workloads running
351 within a trusted compute pool can use these authentication mechanisms to access sensitive
352 information.

353 Stage 4 includes the following prerequisite goals:

- 354 1. **Deploy workloads only to cloud servers with trusted platforms and in trusted
355 locations.** Essentially, stage 2 has been completed and workloads are running on an
356 appropriate host in the trusted compute pool.
- 357 2. **Create an identity for the workload which is signed by its compute node's root of
358 trust.** Each instance of the workload that is instantiated on a compute node will have a
359 unique identity created for it, which is signed by the root of trust on the compute node to
360 prove where it is running.
- 361 3. **Grant workloads appropriate access to sensitivity information based on their
362 identity.** When accessing sensitive information, the workload will present its identity
363 from goal 2 and will be granted the appropriate level of access to sensitive information.
364 The level of access is predefined and is determined by the function of the workload, the
365 platform trustworthiness, and the location of the compute node it is running on.

366 **2.3 Additional Resources**

367 For more information on the technical topics being addressed by these stages, see the following
368 NIST publications:

- 369 • NIST SP 800-128, *Guide for Security-Focused Configuration Management of*
370 *Information Systems*
371 <https://doi.org/10.6028/NIST.SP.800-128>

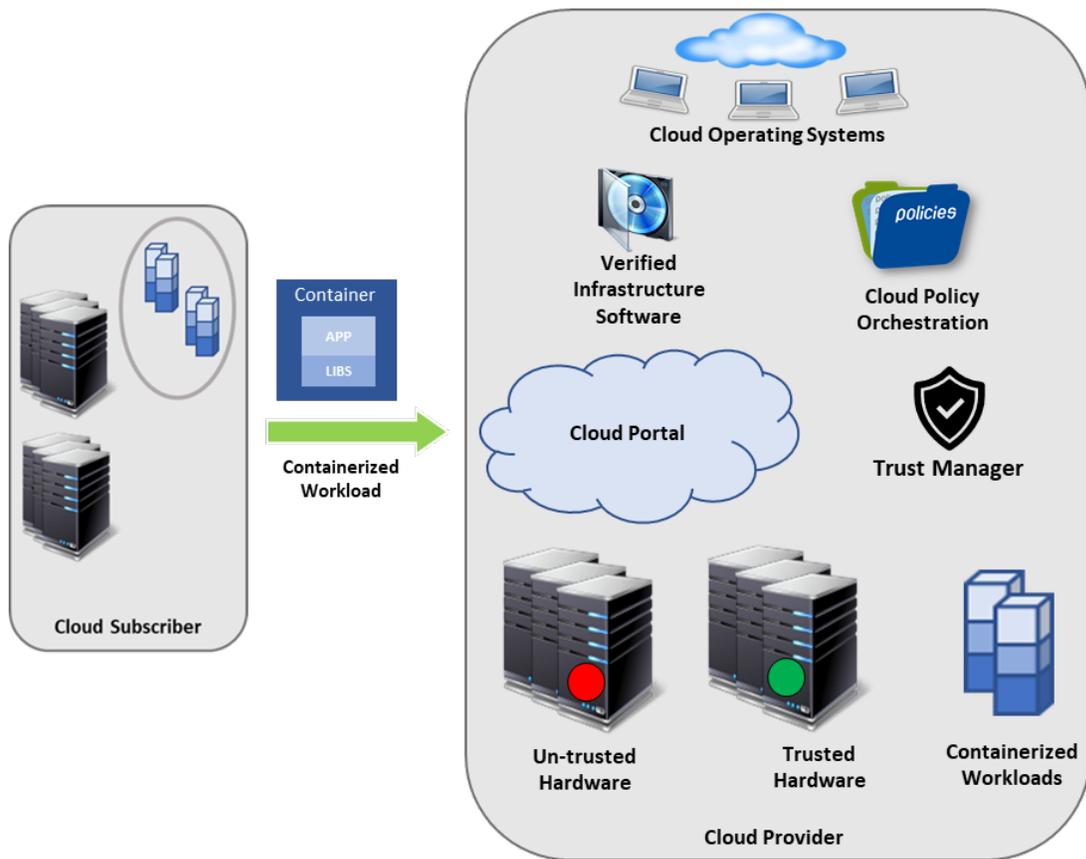
- 372 • NIST SP 800-137, *Information Security Continuous Monitoring for Federal Information*
373 *Systems and Organizations*
374 <https://doi.org/10.6028/NIST.SP.800-137>
- 375 • NIST SP 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*
376 <https://doi.org/10.6028/NIST.SP.800-144>
- 377 • NIST SP 800-147B, *BIOS Protection Guidelines for Servers*
378 <https://doi.org/10.6028/NIST.SP.800-147B>
- 379 • Draft NIST SP 800-155, *BIOS Integrity Measurement Guidelines*
380 <https://csrc.nist.gov/publications/detail/sp/800-155/draft>
- 381 • Draft NIST IR 8320, *Hardware-Enabled Security: Enabling a Layered Approach to*
382 *Platform Security for Cloud and Edge Computing Use Cases*
383 <https://csrc.nist.gov/publications/detail/nistir/8320/draft>
- 384 • NIST IR 8320A, *Hardware-Enabled Security: Container Platform Security Prototype*
385 <https://doi.org/10.6028/NIST.IR.8320A>

386 **3 Prototyping Stage 0**

387 This section provides an overview of stage 0 of the prototype implementation (platform
388 attestation and measured worker node launch).

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

395 Figure 1 depicts the concept of trusted pools. The resources tagged green indicate trusted ones.
396 Critical policies can be defined such that security-sensitive cloud services can only be launched
397 on these trusted resources. For more detailed information and the solution architecture, please
398 refer to Section 3 of NIST IR 8320A.



399

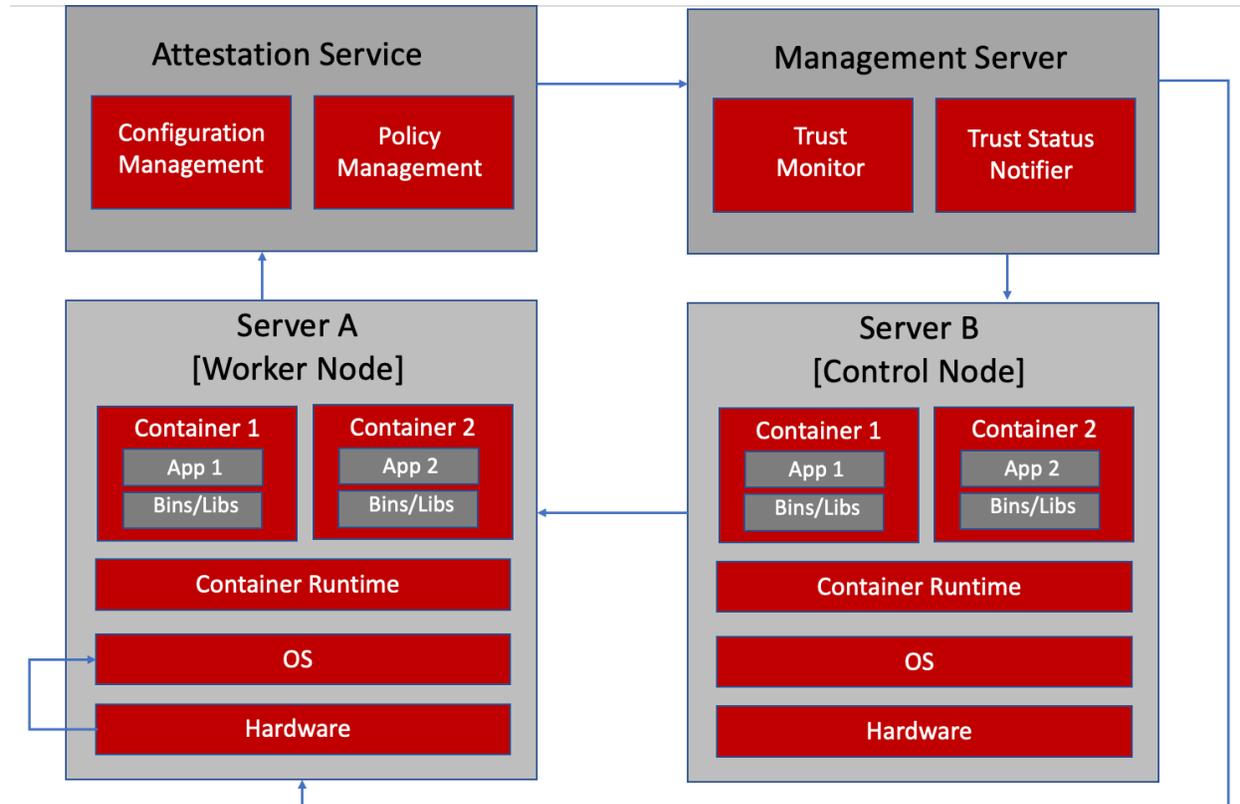
400

Figure 1: Concept of Trusted Pools

401 **4 Prototyping Stage 1**

402 This section provides an overview of stage 1 of the prototype implementation (trusted placement
403 of workloads), which is based on the stage 0 work and adds components that orchestrate the
404 placement of workloads to launch on trusted platforms.

405 Figure 2 shows the components of the stage 1 solution. It assumes that Server A and Server B are
406 two servers within the same cloud.



407
408 **Figure 2: Stage 1 Solution Overview**

409 The solution is comprised of four main components: a control node, a worker node, an attestation
410 service, and a management server. They all work together to deploy container workloads only to
411 worker nodes in a trusted compute pool. For detailed information about the solution overview
412 and the interaction of its components, please refer to Section 4 of NIST IR 8320A.

413 5 Prototyping Stage 2

414 This section discusses stage 2 of the prototype implementation (trust-based and asset tag-based
415 secure workload placement), which is based on the stage 1 work and adds components that take
416 into account asset tag restrictions. The solution architecture is the same as stage 1; however, the
417 additional asset tag measurement is leveraged in the server hardware roots of trust and is taken
418 into account during deployment of workloads.

419 Additionally, the capability of monitoring measurements in a governance, risk, and compliance
420 dashboard are introduced in stage 2. For detailed information about the solution overview and a
421 high-level example of what the dashboard may look like, please refer to Section 4 of NIST IR
422 8320A.

423

6 Prototyping Stage 3

424 This section discusses stage 3 of the prototype implementation (trust-based workload
425 decryption), which is based on the stage 2 work and adds components that allow encrypted
426 container images to be decrypted by servers with trusted platform measurements and asset tags.

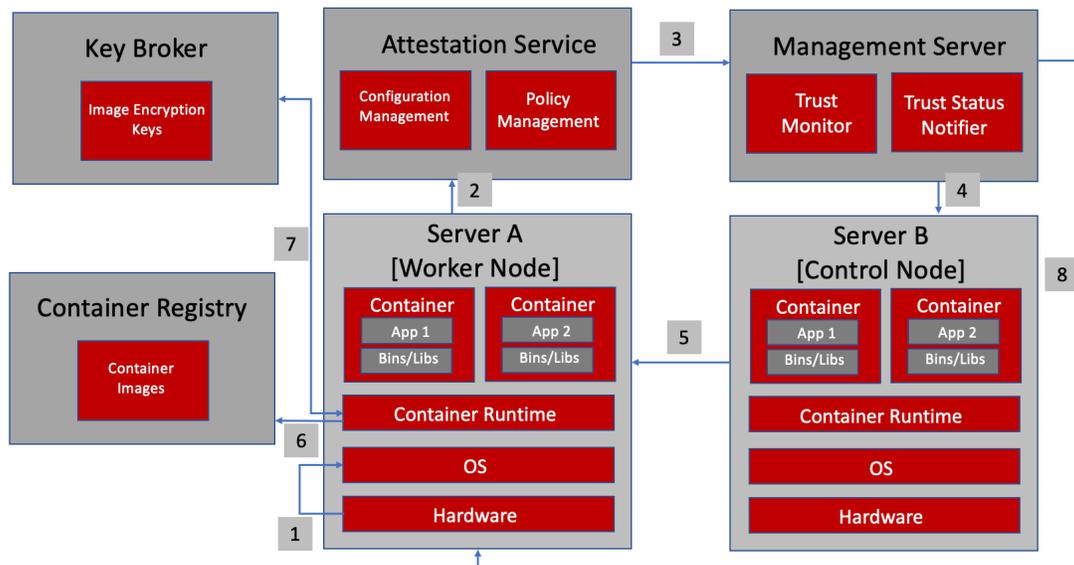
427 **6.1 Solution Overview**

428 Consumers who place their workloads in the cloud or the edge are typically forced to accept that
429 their workloads are secured by their service providers without insight or knowledge as to what
430 security mechanisms are in place. The ability for end users to encrypt their workload images can
431 provide at-rest cryptographic isolation to help protect consumer data and intellectual property.

432 When the runtime node service receives the launch request, it can detect that the image is
433 encrypted and make a request to retrieve the decryption key. This request can be passed through
434 an attestation service so that an internal trust evaluation for the platform can be performed. The
435 key request is forwarded to the key broker with proof that the platform has been attested. The
436 key broker can then verify the attested platform report and release the key back to the cloud
437 service provider and node runtime services. At that time the node runtime can decrypt the image
438 and proceed with the normal workload orchestration. The disk encryption kernel subsystem can
439 provide at-rest encryption for the workload on the platform.

440 **6.2 Solution Architecture**

441 Figure 3 shows the operation of the stage 3 solution. It assumes that Server A and Server B are
442 two servers within the same cloud. It uses the same base architecture as stages 1 and 2, but it
443 introduces two additional components: container registry and key broker. The container registry
444 is where encrypted container images are stored, and the key broker stores their decryption keys
445 and can provide trusted servers with access to the keys.



446

447

Figure 3: Stage 3 Solution Architecture

448 There are eight generic steps performed in the operation of the stage 3 prototype, as outlined
449 below and reflected by the numbers in Figure 3:

- 450 1. Server A performs a measured launch, with the enhanced hardware-based security
451 features populating the measurements in the hardware module.
- 452 2. Server A sends a quote to the Attestation Service. The quote includes signed hashes of
453 various platform firmware and OS components.
- 454 3. The Trust Authority verifies the signature and hash values, and sends the attestation of
455 the platform's integrity state to the Management Server.
- 456 4. The Management Server enforces workload policy requirements on Server B based on
457 user requirements.
- 458 5. Server B launches workloads that require trusted infrastructure only on server platforms
459 that have been attested to be trusted.
- 460 6. Server A pulls the encrypted workload image from the Container Registry so that it can
461 launch an instance of the workload.
- 462 7. The Key Broker releases the workload decryption key to Server A only if it has a trusted
463 attestation report, and Server A launches an instance of the workload.
- 464 8. Each server platform gets audited periodically based on its measurement values.

465 **7 Prototyping Stage 4**

466 This section discusses stage 4 of the prototype implementation (trust-based workload access to
467 information), which is based on the stage 3 work and adds components that create identities for
468 individual workloads so that they can be granted appropriate access to sensitivity information.

469 **7.1 Solution Overview**

470 The majority of workloads running in the cloud need some access to data sources or other
471 services. To do this, they must authenticate themselves using a password, API key, or certificate.
472 Today, this is typically done through secrets. Even though secrets are designed to be stored
473 securely (for example, encrypted at rest) by the orchestration, they can simply be mounted to an
474 arbitrary container and read by anyone who has access to the namespace, including cloud
475 administrators. Those knowing the secret can also access the sensitive data that needs to be
476 protected. The problem with the secrets is that once they are stored, they are also available to
477 administrators, cloud operators, or anyone else with access to the namespace, whether or not they
478 are authorized to access the data that the secrets protect.

479 Trust-based workload access to information protects sensitive data access by ensuring only
480 attested services with specific location-based restrictions can access the secrets. This is done
481 through the use of workload identity, which is composed of the trusted hardware identity data
482 that has been fully attested, including the data center location and region, and various runtime
483 measurements to identify the application. These measurements are securely signed by a service
484 running on each worker node, using a chain of trust created during the secure bootstrapping of
485 the environment, then continuously attested and validated. The bootstrapping of the environment
486 requires configuring a secret store that runs a root Certificate Authority (CA), and installing an
487 intermediate CA and token signing service on each worker node. Each worker node with the
488 token signing service uses its hardware root of trust to protect its individual private key.

489 **7.2 Solution Architecture**

490 It is assumed that all the steps from stage 3 have been completed, and that a workload is
491 successfully deployed to a worker node before any steps of this stage begin. Figure 4 shows the
492 operation of the stage 4 solution, with the assumptions that the workload deployed is on top of a
493 trusted worker node and policies have been defined for the measurements of the application that
494 can access secrets. These measurements represent the identity of the application.

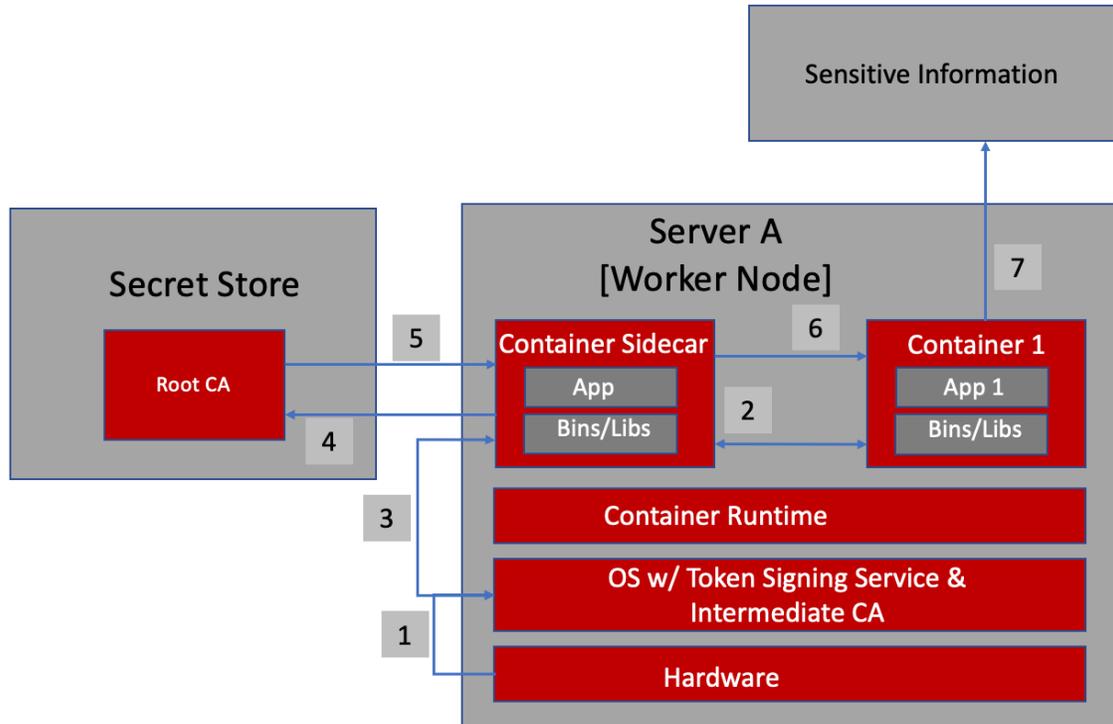


Figure 4: Stage 4 Solution Architecture

495

496

497 There are seven generic steps performed in the operation of the stage 4 prototype, as outlined
498 below and reflected by the numbers in Figure 4:

- 499
- 500 1. Server A has been bootstrapped by installing an intermediate CA that contains a signing authority which uses the hardware root of trust to protect its private key.
 - 501 2. When Server A instantiates a workload instance with its sidecar, the sidecar collects the
502 measurements called *claims* that define the identity of this application.
 - 503 3. The workload sidecar sends the measurements securely to the token signing service on
504 Server A, and the signing service signs these claims using the intermediate CA, then
505 returns the token to the sidecar.
 - 506 4. The sidecar requests the annotated secrets from the secret store by passing the signed
507 token along with the request.
 - 508 5. The secret store validates the signature and the expiration date on the token, and if
509 everything is valid, it uses the provided claims against the policies to retrieve the secret.
510 If the measurements match the policy, the secret is released to the application.
 - 511 6. The sidecar injects the secret into the running workload instance.
 - 512 7. The running workload instance can easily access the secret locally and use it to obtain
513 sensitive data.

514 **References**

- [1] Bartock M, Souppaya M, Wheeler J, Knoll T, Shetty U, Savino R, Inbaraj J, Righi S, Scarfone K (2021) Hardware-Enabled Security: Container Platform Security Prototype. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 8320A. <https://doi.org/10.6028/NIST.IR.8320A>
- [2] Bartock M, Souppaya M, Savino R, Knoll T, Shetty U, Cherfaoui M, Yeluri R, Malhotra A, Banks D, Jordan M, Pendarakis D, Rao JR, Romness P, Scarfone KA (2021) Hardware-Enabled Security: Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases. (National Institute of Standards and Technology, Gaithersburg, MD), 2nd Draft NIST Interagency or Internal Report (IR) 8320. <https://doi.org/10.6028/NIST.IR.8320-draft2>
- [3] Polydys ML, Wisseman S (2009) Software Assurance in Acquisition: Mitigating Risks to the Enterprise. <https://apps.dtic.mil/dtic/tr/fulltext/u2/a495389.pdf>
- [4] 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. Includes updates as of December 10, 2020. <https://doi.org/10.6028/NIST.SP.800-53r5>
- [5] 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>

515 **Appendix A—Hardware Root of Trust Implementation**

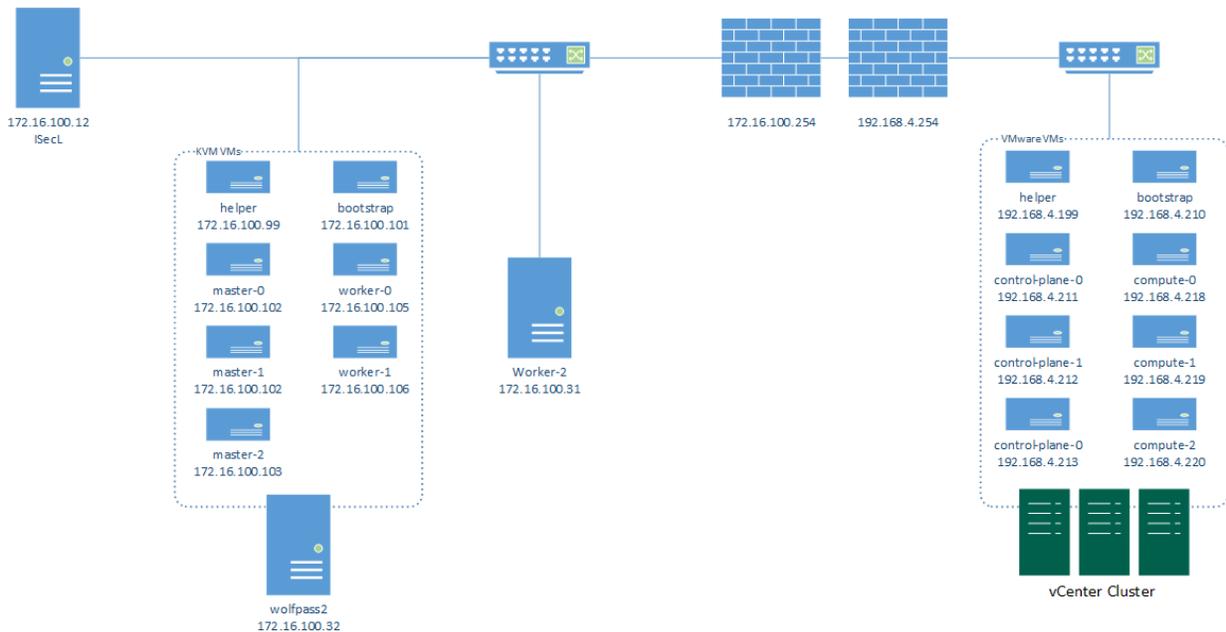
516 This appendix provides an overview of the high-level hardware architecture of the prototype
 517 implementation, as well as details on how Intel platforms implement hardware modules and
 518 enhanced hardware-based security functions.

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

520 Figure 5 shows the high-level implementation architecture. The Intel Security Libraries for Data
 521 Center (ISecL-DC) server (in the upper left corner) contains the key broker service, attestation
 522 service, and utilities for attesting the hardware root of trust and host measurements. Descriptions
 523 of each component and installation steps are in the [ISecL-DC product guide](#) and [quick start](#)
 524 [guide](#).

525 There are two OpenShift clusters, one comprised of virtual machines (VMs) running on a
 526 VMware cluster, and another comprised of VMs running on kernel-based virtual machines
 527 (KVMs) plus one bare metal host. The first cluster is used as a management cluster, while the
 528 second is a managed cluster.

- 529 • The managed cluster is the cluster in which the trusted workloads will run. There can be
 530 multiple managed clusters governed by the [IBM Cloud Pak for Multicloud Management](#)
 531 (MCM).
- 532 • The management cluster contains the control plane for MCM, as well as DevOps-related
 533 tooling.



534

535

Figure 5: Prototype Implementation Architecture

536 A.2 Hardware Root of Trust: Intel TXT and Trusted Platform Module (TPM)

537 Hardware-based root-of-trust, when coupled with an enabled BIOS, OS, and components,
538 constitutes the foundation for a more secure computing platform. This secure platform ensures
539 BIOS and OS integrity at boot from rootkits and other low-level attacks. It establishes the
540 trustworthiness of the server and host platforms.

541 There are three roots of trust in a trusted platform: root of trust for measurement (RTM), root of
542 trust for reporting (RTR), and root of trust for storage (RTS). They are the foundational elements
543 of a single platform. These are the system elements that must be trusted, because misbehavior in
544 these normally would not be detectable in the higher layers. In an Intel Trusted Execution
545 Technology (TXT) enabled platform, the RTM is the Intel microcode: the Core-RTM (CRTM).
546 An RTM is the first component to send integrity-relevant information (measurements) to the
547 RTS. Trust in this component is the basis for trust in all the other measurements. RTS contains
548 the component identities (measurements) and other sensitive information. A trusted platform
549 module (TPM) provides the RTS and RTR capabilities in a trusted computing platform.

550 Intel TXT is the RTM, and it is a mechanism to enable visibility, trust, and control in the cloud.
551 Intel TXT is a set of enhanced hardware components designed to protect sensitive information
552 from software-based attacks. Intel TXT features include capabilities in the microprocessor,
553 chipset, I/O subsystems, and other platform components. When coupled with an enabled OS and
554 enabled applications, these capabilities safeguard the confidentiality and integrity of data in the
555 face of increasingly hostile environments.

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

- 557 • **Protected execution:** Lets applications run in isolated environments so that no
558 unauthorized software on the platform can observe or tamper with the operational
559 information. Each of these isolated environments executes with the use of dedicated
560 resources managed by the platform.
- 561 • **Sealed storage:** Provides the ability to encrypt and store keys, data, and other sensitive
562 information within the hardware. This can only be decrypted by the same environment
563 that encrypted it.
- 564 • **Attestation:** Enables a system to provide assurance that the protected environment has
565 been correctly invoked and to take a measurement of the software running in the
566 protected space. This is achieved by the attestation process defined in the next subsection.
567 The information exchanged during this process is known as the attestation identity key
568 credential and is used to establish mutual trust between parties.
- 569 • **Protected launch:** Provides the controlled launch and registration of critical system
570 software components in a protected execution environment.

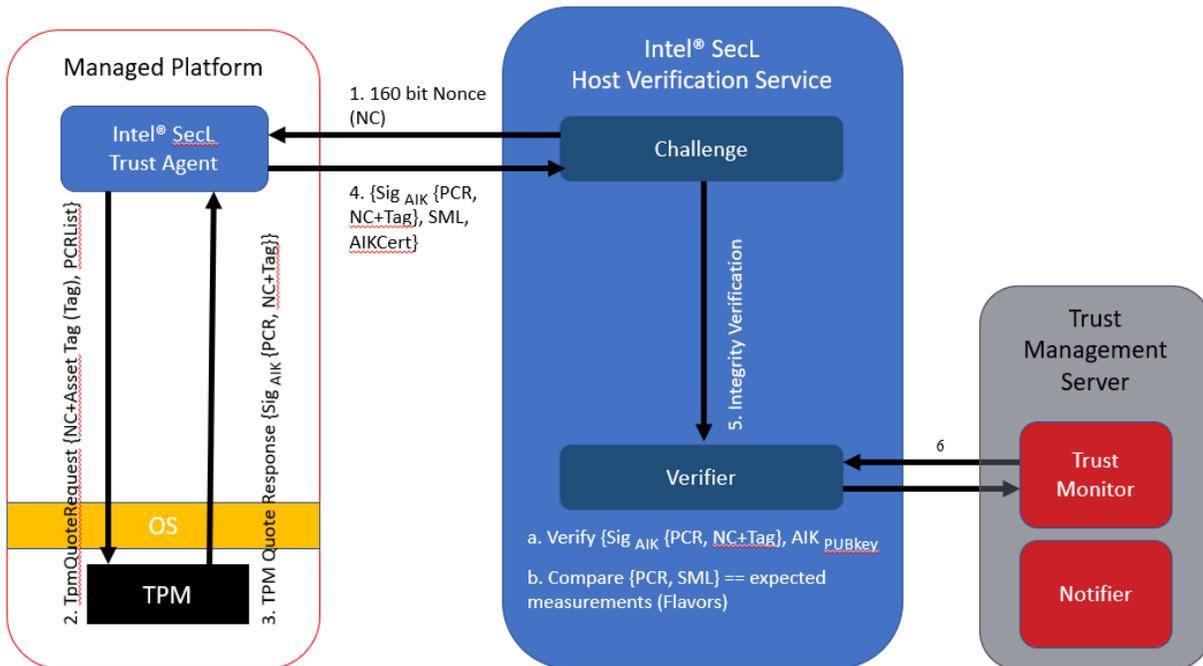
571 Intel Xeon® Platinum Scalable processor series and the previous generation Xeon Processor E3,
572 Xeon Processor E5, and Xeon Processor E7 series processors support Intel TXT.

573 Intel TXT works through the creation of a measured launch environment (MLE) enabling an
574 accurate comparison of all the critical elements of the launch environment against a known good

575 source. Intel TXT creates a cryptographically unique identifier for each approved launch-enabled
 576 component and then provides a hardware-based enforcement mechanism to block the launch of
 577 any code that does not match or, alternately, indicate when an expected trusted launch has not
 578 happened through a process of secure remote attestation. In the latter case, when an attestation
 579 indicates that one or more measured components in the MLE do not match expectations,
 580 orchestration of workloads can be prevented on the suspect platform, even though the platform
 581 itself still launches. This hardware-based solution provides the foundation on which IT
 582 administrators can build trusted platform solutions to protect against aggressive software-based
 583 attacks and to better control their virtualized or cloud environments.

584 **A.3 Attestation: Intel Security Libraries (ISecL)**

585 An attestation authority provides the definitive answers to these questions. Attestation provides
 586 cryptographic proof of compliance, utilizing the root of trust concept to provide actionable
 587 security controls by making the information from various roots of trust visible and usable by
 588 other entities. Figure 6 illustrates the attestation protocol providing the means for conveying
 589 measurements to the challenger. The endpoint attesting device must have a means of measuring
 590 the BIOS firmware, low-level device drivers, and OS and other measured components, and
 591 forwarding those measurements to the attestation authority. The attesting device must do this
 592 while protecting the integrity, authenticity, nonrepudiation, and in some cases, confidentiality of
 593 those measurements.



594
595

Figure 6: Remote Attestation Protocol

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

- 597 1. The challenger, at the request of a requester, creates a non-predictable nonce (NC) and
598 sends it to the attestation agent on the attesting node, along with the selected list of
599 Platform Configuration Registers (PCRs).
- 600 2. The attestation agent sends that request to the TPM as a TPMQuoteRequest with the
601 nonce and the PCR list.
- 602 3. In response to the TPMQuoteRequest, the TPM loads the attestation identity key (AIK)
603 from protected storage in the TPM by using the storage root key (SRK), and performs a
604 TPM Quote command, which is used to sign the selected PCRs and the NC with the
605 private key AIKpriv. Additionally, the attesting agent retrieves the stored measurement
606 log (SML).
- 607 4. In the integrity response step, the attesting agent sends the response consisting of the
608 signed quote, signed NC, and the SML to the challenger. The attesting agent also delivers
609 the AIK credential, which consists of the AIKpub that was signed by a privacy CA.
- 610 5. For the integrity verification step:
 - 611 a. The challenger validates if the AIK credential was signed by a trusted Privacy-
612 CA, thus belonging to a genuine TPM. The challenger also verifies whether
613 AIKpub is still valid by checking the certificate revocation list of the trusted
614 issuing party.
 - 615 b. The challenger verifies the signature of the quote and checks the freshness of the
616 quote.
 - 617 c. Based on the received SML and the PCR values, the challenger processes the
618 SML, compares the individual module hashes that are extended to the PCRs
619 against the “good known or golden values,” and recomputes the received PCR
620 values. If the individual values match the golden values and if the computed
621 values match the signed aggregate, the remote node is asserted to be in a trusted
622 state.
- 623 6. The verifier informs the manager of the trust state of the remote node. The manager
624 records the trust state in its management database and uses it for any individual or
625 aggregated device status requests. If an administrator is subscribed to trust-related events,
626 the manager will also send email notifications when a managed remote node is detected
627 as being untrusted.

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

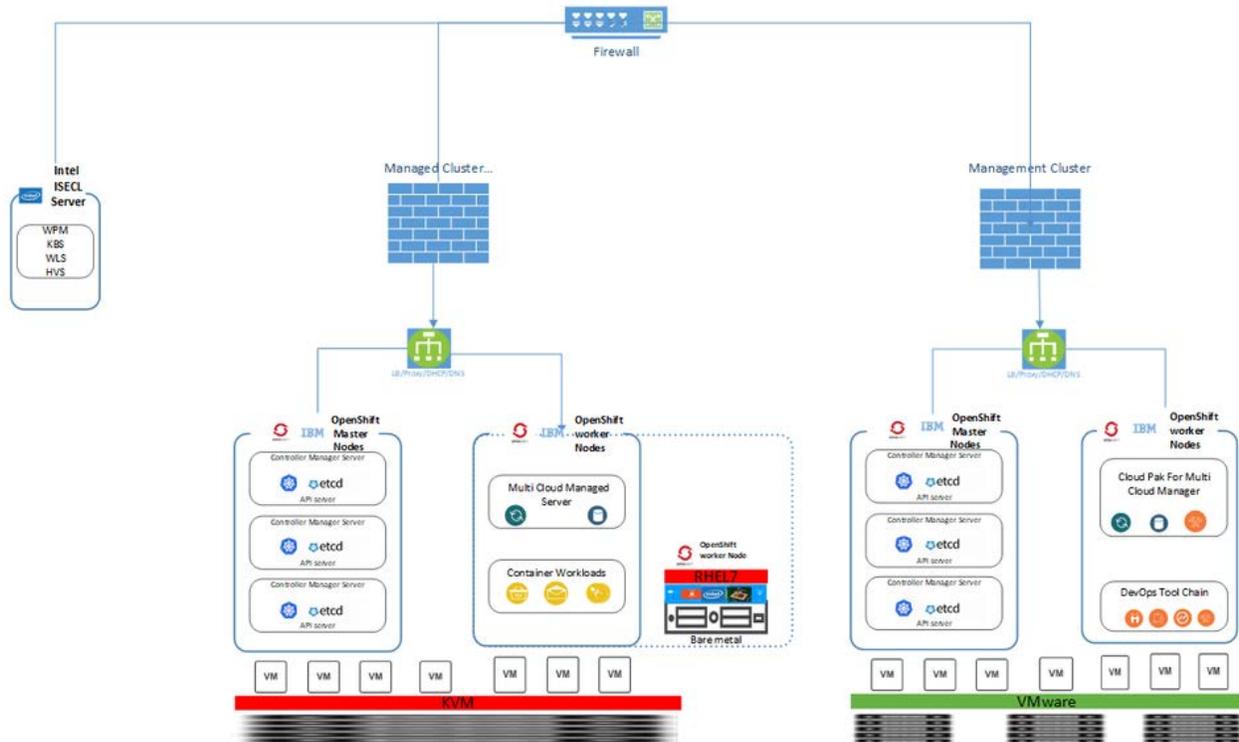
629 Once the ISecL trust agent and host verification service are installed successfully, asset tags can
630 be created and provisioned to each managed server. Section 6.4 in the [ISecL-DC v1.6 product](#)
631 [guide](#) describes the steps of creating and provisioning asset tags.

632 **Appendix B—Workload Orchestration Implementation: OpenShift**

633 This section contains supplementary information describing all the components and steps
634 required to set up the prototype implementation for OpenShift.

635 **B.1 Prototype Architecture**

636 Kubernetes has become a popular source for building large web-scale technologies that enable
637 the enterprise to take advantage of innovations like analytics, artificial intelligence, machine
638 learning, and cloud services. Supporting advanced technologies and building cloud-native
639 applications requires enterprise-grade platforms such as Red Hat OpenShift. This section
640 describes the provisioning and configuration of OpenShift clusters. Figure 7 shows how the
641 nodes from Figure 5 (in Appendix A) in the architecture are logically implemented into two
642 OpenShift clusters and the Remote Attestation Server.



643

644

Figure 7: Prototype Architecture

645 To implement use cases, OpenShift is deployed in two separate clusters: Management Cluster
646 based on VMWare, and Managed Cluster based on KVM infrastructure. Each cluster has three
647 control-plane and three worker nodes, and a VM with load balancer and DNS to simulate the
648 trusted container workload environment. MCM components are deployed on both OpenShift
649 clusters. The Managed Cluster has an MCM Hub component, which aggregates information from
650 multiple clusters using an asynchronous work request model. The hub cluster (Management
651 Cluster) maintains the status of clusters and applications, and provides a set of APIs for the
652 various functions to support as a central controller. The Managed Cluster has an MCM managed-

653 cluster component used to define the cluster, with the MCM Klusterlet and other resources
654 configured to initiate a connection to the hub cluster. The Managed Cluster receives work
655 requests and applies them, then returns the results.

656 **B.2 OpenShift Installation and Configuration**

657 Managing the OpenShift platform in multiple places comes with challenges like complexity,
658 governance, and cost. For example, how do you gain visibility into all the clusters to see where
659 the application's components are running? How do you know which systems are failing? How
660 can you monitor usage across the clouds and clusters? How do you govern the configuration and
661 changes to this environment? IBM Cloud Pak for Multicloud Management/Red Hat Advanced
662 Cluster Management for Kubernetes supports the use cases. It is based on the Kubernetes
663 community direction and includes advanced functions important to running enterprise-grade
664 environments.

665 The Policies repository, which is part of the Manager hub component, resides on the
666 Management Cluster. The repository comes with default compliance and policies templates, but
667 our use cases developed new policies that reflected our environment to integrate with the Intel
668 attestation hub. The repository holds the policy defined for the Managed Cluster, and the policy
669 document is applied by using placement bind.

670 **B.2.1 VMware-Based Management Cluster (Cluster A)**

671 **Hardware Requirement:** For deploying OpenShift Container Platform (OCP) 4.3 on VMware,
672 the minimum recommendation is to provision one ESXi server and one Centos/Red Hat VM on
673 the same virtual local area network (VLAN) in the local datacenter. For this deployment, the
674 setup was an ESXi bare-metal server with 48 CPUs, 256 GB RAM, and 2 TB storage. The
675 Centos/Red Hat VM is only required for a few hours and can be de-provisioned after the install
676 is complete.

677 **Networking:** The IP addresses used in this process and the configuration files came from our
678 NCCoE environment. They are used here for illustration purposes only. Besides setting up your
679 ESXi and vCenter server, you need to have a minimum of 16 IP addresses to assign to the VMs.
680 Each VM node takes one IP address. The recommended minimum of 16 IP addresses is
681 determined by: 1 helper node + 1 boot node + 3 control-plane nodes + 3 worker nodes = 8 nodes.
682 The extra IP addresses are available in case additional worker nodes are required in the future.
683 This installation provisioned the vCenter on the same IP subnet, so a total of 9 IP addresses were
684 used.

685 **VMware OCP VM Requirements:** Table 1 lists the VMs that are instantiated on the VMware
686 server, along with their virtual hardware requirements and the roles they serve in the cluster.

687

Table 1: VMs Instantiated on the VMware-Based Management Cluster

Node Name	vCPU	Mem (GB)	HDD (GB)	Role
Helper Node	4	16	150	LB/DNS/Proxy/DHCP/OCP Installer
Bootstrap-0	4	16	150	Bootstrap OCP
Control-plane-0	4	16	150	Controller OCP
Control-plane-1	4	16	150	Controller OCP
Control-plane-2	4	16	150	Controller OCP
compute-0	4	16	150	Compute OCP
compute-1	4	16	150	Compute OCP
compute-2	4	16	150	Compute OCP

688 **OCP VMware Deployment Playbooks:** To deploy OCP 4.3 on VMware, download the
 689 following Git repository: <https://github.com/fctoibm/ocpvmware4.3> and follow the steps to run
 690 the playbooks. Make sure to change the `vars.yaml` and `host.yaml` files to match the
 691 networking information for your environment.

692 **B.2.2 KVM-Based Managed Cluster (Cluster B)**

693 The second OCP cluster is the managed cluster. It contains an MCM Klusterlet, which ensures
 694 that each managed cluster adheres to the policy in place.

695 **Hardware Requirement:** For this deployment, the lab setup was on a CentOS bare-metal server
 696 with 48 CPUs, 256 GB RAM, and 1 TB storage. KVM will be used to create and manage virtual
 697 machines. The KVM command line tool is `virt-install` and the GUI tool is `virt-manager`.
 698 To use the KVM GUI tool, install Gnome desktop and VNC on the CentOS bare-metal server.
 699 All of the VMs for this managed cluster are deployed on this single KVM host, which has
 700 hostname `wolfpass2` in the table and image in Figure 5 (in Appendix A).

701 **Networking:** The IP addresses used in this process and the configuration files came from our
 702 NCCoE environment. They are used here for illustration purposes only. When you install OCP in
 703 the KVM host environment, you will also need a minimum of 16 portable IP addresses. Each
 704 VM node takes up one IP address. The recommended minimum of 16 portable IP addresses is
 705 determined by: 1 helper node + 1 boot node + 3 control-plane nodes + 3 worker nodes = 8 nodes.
 706 The extra IP addresses are available for additional worker nodes required in the future. You
 707 should plan your IP address space accordingly.

708 **KVM OCP VM Requirements:** Table 2 lists the VMs that are instantiated on the KVM server,
 709 along with their virtual hardware requirements and the roles they serve in the managed cluster.

710

Table 2: VMs Instantiated on the KVM-Based Managed Cluster

Node Name	vCPU	Mem (GB)	HDD (GB)	Role
Helper Node	4	16	150	DNS/Proxy/DHCP/OCP Installer
Bootstrap	4	16	150	Bootstrap OCP
Master0	4	16	150	Controller OCP
Master1	4	16	150	Controller OCP
Master2	4	16	150	Controller OCP
Worker0	4	16	150	Compute OCP
Worker1	4	16	150	Compute OCP

711 Note: The OpenShift cluster requires three worker nodes; however, since this deployment uses
712 an additional physical sever for the third worker node, only two worker node VMs are deployed.

713 **OCP KVM Deployment Playbooks:** To deploy OCP 4.3 on KVM, download the following Git
714 repository: <https://github.com/fctoibm/ocpkvm4.3> and follow the steps to run the playbooks.
715 Make sure to change the `vars.yaml` and `host.yaml` files to match the networking information
716 of your environment.

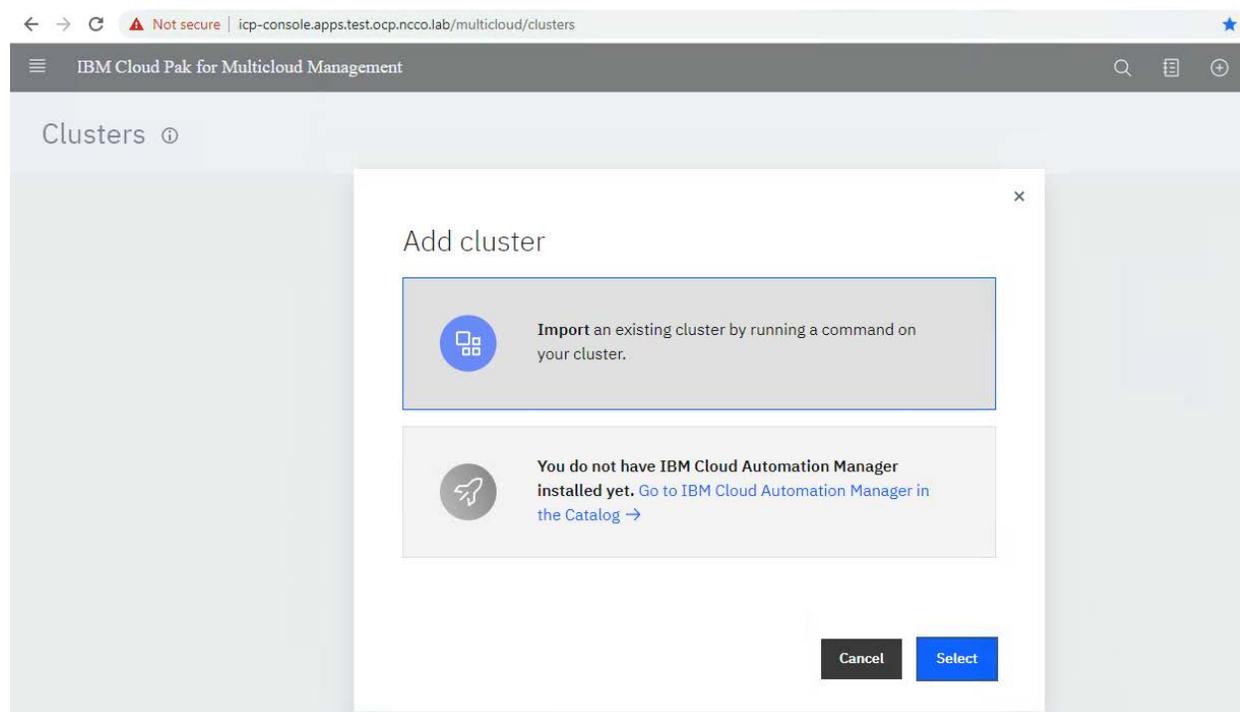
717 The OCP KVM deployment playbook creates all of the worker nodes as virtual machines. In
718 order to create policies bases on hardware roots of trust, a physical server with the Intel TXT and
719 TPM capabilities must be added as an additional worker node to the cluster. This server needs to
720 have the corresponding Red Hat Enterprise Linux (RHEL) OS installed, as well as the Intel Trust
721 Agent as described in Appendix A.2. For the physical server, the OpenShift documentation
722 details [how to add an RHEL compute node to an existing cluster](#).

723 **B.2.3 Installing MCM Pak 1.3 (MCM HUB - VMware)**

724 To install MCM Pak 1.3 on OCP 4.3, the setup assumes OCP 4.3 is already installed and
725 administrator-level access is available to deploy the MCM Pak. The guide assumes OCP 4.3 was
726 installed using the same GitHub repository and the same `vars.yaml` file.

727 **Deploying MCM Pak:** The installation Git repository supports two options, VMware or KVM
728 install, and both will deploy a VM guest if required. The VM guest called PakHelper node will
729 act as a client to install MCM Pak. There is no reason to deploy a VM guest client if you already
730 have a VM guest Centos 7 OS available in the same network. If the Centos 7 VM is already in
731 place, please skip options 1 and 2, but if there is no VM guest available, please execute both
732 options 1 or 2 and 3 from the following Git repository: <https://github.com/fctoibm/mcmpak1.3>
733 and follow the steps to run the playbooks.

734 **Adding KVM OCP as Managed Cluster in MCM:** Once the MCM Pak has been deployed, the
735 KVM OCP cluster can be imported into, and managed by, the IBM MCM. To do so, browse to
736 the web user interface (UI) of MCM and navigate to the Clusters management page. As shown in
737 Figure 8, there is an option to import an existing cluster. To import the existing KVM OCP
738 cluster, perform the steps in [this IBM knowledge center article](#).



739

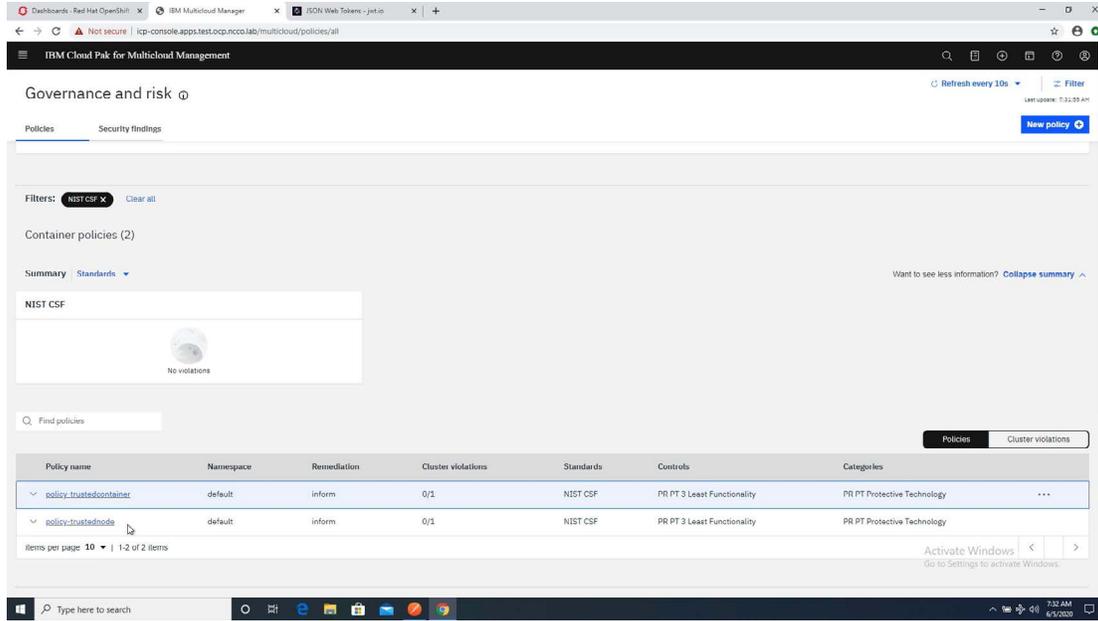
740

Figure 8: MCM Console to Import a Cluster

741 **Creating Policies for Managed Cluster:** Once the KVM OCP cluster has been successfully
 742 imported, specific policies that govern workload orchestration can be created and applied.
 743 Policies are created in the MCM Hub, and these policies are propagated to each managed cluster,
 744 where they are enforced. Two policies were put in place for our prototype implementation:

- 745 1. The “[Trusted Node Policy](#)” ensures that all nodes in the cluster are trusted and attested. In
 746 the inform mode, the policy logs whenever the trust status of a node has been violated. In
 747 the enforce mode, the policy drains and removes the node from the cluster.
- 748 2. The “[Trusted Container Policy](#)” ensures that all workloads run within a namespace are
 749 using a set of images from a particular registry path. The infrastructure is set up so only
 750 encrypted container images are in that path. This makes it so only encrypted images are
 751 run within the namespace.

752 Figure 9 shows the two policies in the web user interface that have been created for the managed
 753 clusters.



754

755

Figure 9: Managed Cluster Policies

756

757

758

In addition to these two policies, there is a Tekton task set up as part of the OpenShift pipeline that does a set of checks and encrypts the image. This secure pipeline does building, vulnerability scanning, and encryption. More details on this pipeline are provided in Appendix C.

759 Appendix C— Workload Encryption Implementation

760 This section contains supplementary information describing all the components and steps
761 required to set up the prototype implementation for container workload encryption.

762 C.1 Prototype Architecture

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

764 C.2 Workload Encryption Configuration

765 Various parts of the container ecosystem allow the enablement of workload encryption via
766 container image encryption. The technology is based on the Open Container Initiative (OCI)
767 container image specification. The components that support the use of this are:

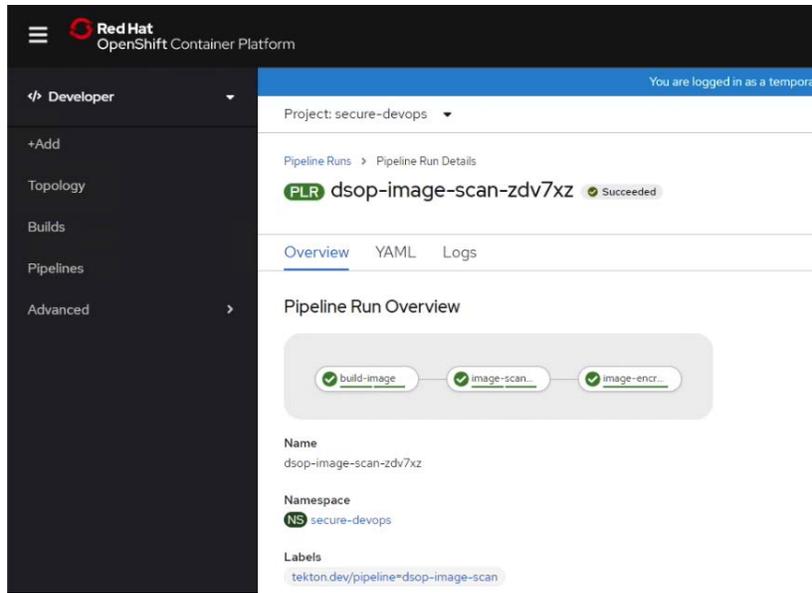
- 768 • **Build:** The Skopeo tool is used to encrypt the container images and push them to the
769 registry.
- 770 • **Runtime:** The Cri-o container runtime came as part of OpenShift and was configured to
771 decrypt the images. It is the default runtime of OpenShift 4.3 worker nodes and supports
772 decrypting OCI container images.
- 773 • **Registry:** The Docker Distribution registry is used to push, pull, and store the encrypted
774 images. The version used was v2.7.1.

775 These are the core components of workload encryption. Several integrations were required with
776 the ISecL Attestation Hub’s APIs to showcase the workload and image encryption with hardware
777 attestation. A custom container encryption metadata scheme was defined to work with the ISecL
778 key broker. The reference implementation code and the document are located at
779 <https://github.com/lumjib/seclkeywrap>.

780 Key integration points are:

- 781 • A custom container encryption metadata scheme was defined to work with the ISecL key
782 broker.
- 783 • The core components CRI-0 and Skopeo were patched to enable the use of the custom
784 ISecL protocol. The patches are available at [https://github.com/lumjib/crilo-
785 o/tree/sample_integration](https://github.com/lumjib/crilo/tree/sample_integration) and https://github.com/lumjib/skopeo/tree/sample_integration,
786 respectively.

787 As part of the DevSecOps cycle, it was integrated with the development flow with the Tekton
788 pipeline to perform builds, security checks, and image encryption. The pipeline workflow can be
789 reached through the OpenShift dashboard by toggling into the Developer role and selecting the
790 “Pipelines” menu, as shown in Figure 10. The definitions of the Tekton objects are available
791 here: <https://gist.github.com/lumjib/22191008f849f240851aec8a1ee0304d>



792

793

Figure 10: Creating Pipeline for Image Decryption

794 **Appendix D—Trusted Service Identity (TSI)**

795 This section contains supplementary information describing all the required components and
796 steps required to set up the prototype implementation for the Trusted Service Identity (TSI).

797 **D.1 TSI Overview**

798 TSI protects sensitive data access by ensuring only attested services with specific location-based
799 restrictions can obtain credentials. This is done through the use of workload identity, composed
800 of the trusted hardware identity data that has been fully attested by Intel TXT, including the data
801 center location and region, and various runtime measurements like the image and cluster name,
802 unique pod IDs, and namespace to identify the application. These measurements are securely
803 signed by a service running on every hosting node, using a chain of trust created during the
804 secure bootstrapping of the environment, then continuously attested and validated.

805 Every container that requires a secret for accessing sensitive data gets assigned a short-lived
806 measured identity, in the form of a JSON Web Token (JWT) token, that is signed with the root of
807 trust by the attested process. This measured identity is a form of an *ephemeral digital biometric*
808 with a short time to live.

809 In this implementation, there is a Kubernetes cluster running on OpenShift, extended with TSI,
810 where each node has an intermediate CA, signed by the root CA during the secure bootstrapping
811 of the cluster.

812 During this bootstrapping, the installation process obtains an attestation report (Security
813 Assertion Markup Language [SAML]) from the Attestation Server for every worker node. This
814 report is checked to verify if all the components are trusted (OS, platform, and software), and the
815 bootstrapping process retrieves worker identity fields from the asset tag of this report. Each
816 worker node also has a JWT Signing Service (JSS) that contains a signing authority that uses the
817 hardware TPM to protect its individual private key.

818 The root CA is securely stored in a vault, which is extended with the TSI Authentication Vault
819 plugin.

820 **D.2 TSI Installation and Configuration**

821 TSI requires an attestation process to accurately define the identity of the worker nodes hosting
822 the application containers. In this implementation, TSI relies on the ISecL server that has been
823 deployed to provide the identity of the worker nodes. Steps detailing the integration of TSI with
824 ISecL can be found here: [https://github.com/IBM/trusted-service-identity/blob/intel-
825 asset/README.md#attestation](https://github.com/IBM/trusted-service-identity/blob/intel-asset/README.md#attestation).

826 This process requires two independent phases:

- 827 • **Asset Registration with Intel Verification Server** – The trusted bootstrapping process
828 responsible for installing the environment must properly set the identity attributes of
829 every worker node. These identity values in the form of asset tags are securely stored in
830 their corresponding TPMs on hosts. As a result, they are included in the SAML

831 attestation report that also includes all the attestation results (OS, platform, software
832 trusted). This process was performed in the steps outlined in Appendix A.

- 833 • **TSI Deployment with Attestation** – The implementation outlined in this document
834 allows for the use of the Intel Attestation Server to obtain the identity of the worker
835 nodes. There are several changes required to configure the TSI installation to support
836 Intel Attestation Server. Additionally, a hardware TPM device is shared between the Intel
837 Trust Agent and the TSI JWT Signing Service, and it requires the use of TPM proxy. For
838 details outlining the suggested configuration changes, visit
839 [https://github.com/IBM/trusted-service-identity/blob/intel-asset/README.md -](https://github.com/IBM/trusted-service-identity/blob/intel-asset/README.md - attestation)
840 [attestation.](https://github.com/IBM/trusted-service-identity/blob/master/examples/vault/README.md#secrets)

841 As a result of these changes, TSI will be installed in the cluster, using an attestation report from
842 the Intel Attestation Service to provide the identities of the workers and to keep the attestation
843 going.

844 Before secrets can be injected into the application container, first they need to be created in the
845 Secret Store (Vault). Follow these steps for injecting secrets to the Vault:
846 [https://github.com/IBM/trusted-service-](https://github.com/IBM/trusted-service-identity/blob/master/examples/vault/README.md#secrets)
847 [identity/blob/master/examples/vault/README.md#secrets](https://github.com/IBM/trusted-service-identity/blob/master/examples/vault/README.md#secrets)

848 Once the application is started, secrets will be injected based on the application identity,
849 including the workload environment and location. As a result, the secret will be delivered to the
850 container runtime memory without ever being stored anywhere in Kubernetes—but from the
851 point of view of the application, no additional changes were needed.

852 Figure 11 shows a sample JWT created by TSI. Notice its three parts: the header, the payload
853 containing the actual claims, and a signature for validation.

Header	<pre>{ "alg": "RS256", "typ": "JWT", "x5c": ["MIIDTCCAjWgA...qCoGa", "MIIDXjC...cMgoO8="] }</pre>
Payload (claims)	<pre> "hd-trusted": "true", "cluster-region": "eu-de", "cluster-name": "EUcluster", "machineid": "266c2075dace453da02500b328c9e325", "pod": "myubuntu-767584864-k9b59", "images": "f36b6d491e0abf1f7130832e9f32d0771de1d7c727a79cc", "images-names": "res-kompass-kompass-docker- local.artifactory.swgdevops.com/myubuntu:latest@sha256:5b224e1 18f1c444d2b88f89c57420a61b1b3c24584c", "exp": 1541689789, "iat": 1541689759, "iss": "wsched@us.ibm.com", "namespace": "appl-ns" }</pre> <div style="border: 1px solid black; padding: 2px; width: fit-content; margin-left: 200px;">with HD RoT attestation</div>
Signature	<pre>RSASHA256(base64UrlEncode(header) + "." + base64UrlEncode(payload),</pre>

854

855

Figure 11: Sample JWT Created by TSI

856 The claims included here represent the measured identity of the application. They contain some
857 static values like `cluster-region` (e.g., Germany, eu-de), `cluster-name`, individual
858 `machineid` which is a unique worker node ID, and several runtime measurements like a
859 namespace, unique pod ID, and the list of images making up the pod. These images include the
860 image signature, so you can validate the image and guarantee that the application is running the
861 code that you want to be running and it was not tampered with. The `hd-trusted` value
862 determines whether all the attested elements are trusted.

863 Values for `cluster-name`, `cluster-region`, and `hd-trusted` are essential for defining the
864 identity of the compute resources, and they are read from Intel's Attestation Server. There is also
865 a token expiration timestamp, typically set to one minute, to make these tokens ephemeral and
866 short-lived, protecting the security from leaking. These are the runtime measurements that
867 represent the identity of the application, signed with the root of trust, and used for evaluation
868 against policies controlling the secrets.

869 The secrets stored in the Vault are protected by policies. Policies are composed of the policy-
870 type and the attributes, the same that are used for building claims, and they represent the path to
871 the secret. If the claims provided in the request matches the policy attribute path, the secret will
872 be released to the application.

873 **Appendix E—Supporting NIST SP 800-53 Security Controls and Publications**

874 The major controls in the NIST SP 800-53 Revision 5, Security and Privacy Controls for
875 Information Systems and Organizations [4] control catalog that affect the container platform
876 security prototype implementation are:

- 877 • AU-2, Event Logging
- 878 • CA-2, Control Assessments
- 879 • CA-7, Continuous Monitoring
- 880 • CM-2, Baseline Configuration
- 881 • CM-3, Configuration Change Control
- 882 • CM-8, System Component Inventory
- 883 • IR-4, Incident Handling
- 884 • SA-9, External System Services
- 885 • SC-1, Policy and Procedures [for System and Communications Protection Family]
- 886 • SC-7, Boundary Protection
- 887 • SC-29, Heterogeneity
- 888 • SC-32, System Partitioning
- 889 • SC-36, Distributed Processing and Storage
- 890 • SI-3, Malicious Code Protection
- 891 • SI-4, System Monitoring
- 892 • SI-6, Security and Privacy Function Verification
- 893 • SI-7, Software, Firmware, and Information Integrity

894 Table 3 lists the security capabilities provided by the prototype:

895 **Table 3: Security Capabilities Provided by the 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
	IC3.3	Policy-Based Workload Decryption
	IC3.4	Policy-Based Workload Access

Capability Category	Capability Number	Capability Name
IC4 – Reporting	IC4.1	Support for Continuous Monitoring
	IC4.2	Support for On-Demand Reports
	IC4.3	Support for Notification of Trust Events

896 Table 4 maps the security capabilities from Table 3 to the NIST SP 800-53 controls in the list at
897 the beginning of this appendix.

898 **Table 4: Mapping of Security Capabilities to NIST SP 800-53 Controls**

NIST SP 800-53 Control	Measurements				Tag Verification	Policy Enforcement				Reporting		
	IC1.1	IC1.2	IC1.3	IC1.4	IC2.1	IC3.1	IC3.2	IC3.3	IC3.4	IC4.1	IC4.2	IC4.3
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
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-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					

899 Appendix F—Cybersecurity Framework Subcategory Mappings

900 This appendix maps the major security features of the trusted geolocation prototype
901 implementation to the following subcategories from the Cybersecurity Framework [5]:

- 902 • ID.GV-1: Organizational information security policy is established
- 903 • ID.GV-3: Legal and regulatory requirements regarding cybersecurity, including privacy
904 and civil liberties obligations, are understood and managed
- 905 • PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and
906 information integrity
- 907 • PR.IP-5: Policy and regulations regarding the physical operating environment for
908 organizational assets are met

909

910 **Appendix G—Acronyms and Other Abbreviations**

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

AIK	Attestation Identity Key
API	Application Programming Interface
BIOS	Basic Input/Output System
CA	Certificate Authority
CPU	Central Processing Unit
CRTM	Core Root of Trust for Measurement
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
FOIA	Freedom of Information Act
GB	Gigabyte
HD	Hard Drive
HDD	Hard Disk Drive
IaaS	Infrastructure as a Service
Intel TXT	Intel Trusted Execution Technology
I/O	Input/Output
IP	Internet Protocol
IR	Interagency or Internal Report
ISecL	Intel Security Libraries
ISecL-DC	Intel Security Libraries for Data Center
IT	Information Technology
ITL	Information Technology Laboratory
JSON	JavaScript Object Notation
JSS	JWT Signing Service
JWT	JSON Web Token
KVM	Kernel-Based Virtual Machine
MCM	Multicloud Management
MLE	Measured Launch Environment
NC	Nonce
NIST	National Institute of Standards and Technology
OCI	Open Container Initiative
OCP	OpenShift Container Platform
OS	Operating System
PCR	Platform Configuration Register
RAM	Random Access Memory
RHEL	Red Hat Enterprise Linux
RTM	Root of Trust for Measurement
RTR	Root of Trust for Reporting
RTS	Root of Trust for Storage
SAML	Security Assertion Markup Language
SML	Stored Measurement Log
SP	Special Publication
SRK	Storage Root Key
TB	Terabyte

TPM	Trusted Platform Module
TSI	Trusted Service Identity
VLAN	Virtual Local Area Network
VM	Virtual Machine

912