

## TCG Public Review

Copyright © TCG 2017

# TCG Guidance for Securing Network Equipment

Version 1.0  
Revision 26b  
July 3, 2017  
Committee Draft

Contact: [admin@trustedcomputinggroup.org](mailto:admin@trustedcomputinggroup.org)

- **Work in Progress:**  
*This document is an intermediate draft for comment only and is subject to change without notice. Readers should not design products based on this document*

## Disclaimers, Notices, and License Terms

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, DOCUMENT OR SAMPLE.

Without limitation, TCG disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this document and to the implementation of this document, and TCG disclaims all liability for cost of procurement of substitute goods or services, lost profits, loss of use, loss of data or any incidental, consequential, direct, indirect, or special damages, whether under contract, tort, warranty or otherwise, arising in any way out of use or reliance upon this document or any information herein.

This document is copyrighted by Trusted Computing Group (TCG), and no license, express or implied, is granted herein other than as follows: You may not copy or reproduce the document or distribute it to others without written permission from TCG, except that you may freely do so for the purposes of (a) examining or implementing TCG documents or (b) developing, testing, or promoting information technology standards and best practices, so long as you distribute the document with these disclaimers, notices, and license terms.

Contact the Trusted Computing Group at [www.trustedcomputinggroup.org](http://www.trustedcomputinggroup.org) for information on document licensing through membership agreements.

Any marks and brands contained herein are the property of their respective owners.

## Acknowledgements

Laffey, Mr. Tom (Co-chair)	Hewlett Packard Enterprise
Eckel, Mr. Michael (Co-chair)	Huawei Technologies Co., Ltd.
Fedorkow, Mr. Guy (Editor)	Juniper Networks, Inc.
Kuntze, Dr. Nicolai	Huawei Technologies Co., Ltd.
Latze, Ms. Carolin	Carolin Latze
Bell, Mr. Bob	Cisco Systems
Pritikin, Mr. Max	Cisco Systems
Sulzen, Mr. Bill	Cisco Systems
Birkholz, Mr. Henk	Fraunhofer Institute for Secure Information Technology (SIT)
Fuchs, Mr. Andreas	Fraunhofer Institute for Secure Information Technology (SIT)
Kotani, Dr. Seigo	Fujitsu Limited
McDonald, Mr. Ira	High North Inc
Rudolph, Dr. Carsten	Huawei Technologies Co., Ltd.
Zhang, Ms. Yi	Huawei Technologies Co., Ltd.
Hanna, Mr. Steve	Infineon Technologies North America Corp.
Rudy, Mr. Gregory	INTEGRITY Security Services, Inc.
Watsen, Mr. Kent	Juniper Networks, Inc.
Shah, Mr. Atul	Microsoft
Thungathurthi, Mr. Bhavani	Microsoft
Proudler, Mr. Graeme	Graeme Proudler
Lorenzin, Ms. Lisa	Pulse Secure, LLC
Schmidt, Mr. Charles	The MITRE Corporation
Fitzgerald-McKay, Ms. Jessica	United States Government

# Table of Contents

1	1.	Scope, Audience and Purpose .....	7
2	1.1	Scope .....	7
3	1.2	Audience and Purpose.....	7
4	2.	Preface.....	8
5	2.1	Network Equipment Reference Model .....	8
6	2.2	Internal Structure of Network Equipment .....	9
7	2.3	Stakeholders .....	10
8	2.4	Key Differences between Network Equipment and PC Applications .....	11
9	2.4.1	Long Life Cycle .....	12
10	2.4.2	Privacy and Networking Equipment .....	12
11	3.	Use Cases .....	14
12	3.1	Device Identity .....	14
13	3.1.1	OEM Device Identity and Counterfeit Protection .....	14
14	3.1.2	Identity for Network Access .....	15
15	3.2	Secure Zero Touch Provisioning.....	15
16	3.3	Securing Secrets .....	16
17	3.4	Protection of Configuration Data .....	17
18	3.5	Remote Device Management.....	17
19	3.6	Software Inventory .....	18
20	3.7	Attestation of Boot Integrity for Network Devices ("Health Check") .....	18
21	3.8	Inventory of Composite Devices .....	18
22	3.9	Integrity-Protected Logs .....	19
23	3.10	Entropy Generation .....	19
24	3.11	Deprovisioning.....	20
25	4.	Building Blocks.....	21
26	4.1	Chain of Trust.....	21
27	4.2	Secure Boot vs Measured Boot .....	22
28	4.3	802.1AR Device Identification .....	23
29	5.	Implementation Guidance .....	25
30	5.1	Device Identity .....	25
31	5.1.1	OEM Device Identity and Counterfeit Protection .....	26
32	5.1.1.1	Configuring Initial Identity Credentials .....	27
33	5.1.1.2	Signing and Encryption Key Types .....	27
34	5.1.1.3	Identity for Attestation .....	28
35	5.1.1.4	Initial Identity with TPM1.2 .....	29

37	5.1.1.5	Key Types in TPM2.0.....	29
38	5.1.1.6	Initial Identity with TPM2.0.....	31
39	5.1.1.7	Loss of IDevID Credentials .....	33
40	5.1.1.8	Compromise of Keys.....	34
41	5.1.2	Local Device Identity .....	34
42	5.1.2.1	How to Install a Local Device ID .....	34
43	5.1.2.1.1	Conveying TPM Certify Information .....	35
44	5.1.2.2	Attestation with Local Identity .....	36
45	5.1.3	Identity for Network Authentication .....	36
46	5.1.3.1	Proving a Link to the TPM.....	36
47	5.1.3.2	Why Not Just Use EK and Platform Cert for Identity? .....	37
48	5.2	Secure Zero Touch Provisioning.....	37
49	5.3	Securing Secrets.....	38
50	5.4	Protection of Configuration Data .....	39
51	5.5	Remote Device Management.....	40
52	5.6	Software Inventory .....	40
53	5.7	Attestation of Integrity for Network Devices (“Health Check”) .....	42
54	5.7.1	Linux Integrity Measurement Architecture (IMA) .....	43
55	5.7.1.1	IMA Appraisal.....	46
56	5.7.2	PCR Definitions.....	47
57	5.8	Composite Networking Devices .....	47
58	5.9	Integrity-Protected Logs .....	49
59	5.9.1	Functional Requirements .....	49
60	5.9.2	Integrity-Protected Log Implementation .....	50
61	5.9.2.1	Timekeeping in TPM 1.2 .....	50
62	5.9.2.2	Timekeeping in TPM 2.0 .....	51
63	5.9.2.3	Assembling the Log .....	51
64	5.10	Entropy Generation .....	52
65	5.11	Deprovisioning.....	53
66	6.	References.....	54
67	7.	Glossary.....	56
68	8.	Appendices .....	57
69	8.1	Use-Cases For Further Study .....	57
70	8.1.1	Secure Identity in Virtual Machines.....	57
71	8.1.2	Distributed Composite Devices.....	57
72	8.2	Standardization Work For Further Study .....	57
73	8.2.1	Known Good Hashes .....	57

74	8.2.2	Composite Systems .....	57
75	8.2.3	Reference Manifests .....	58
76	9.	Revision History .....	59

## 77 **1. Scope, Audience and Purpose**

78 This document, the *TCG Guidance for Securing Network Equipment*, is part of the Trusted  
79 Computing Group's collection of Reference Documents, written by the Network Equipment  
80 Sub Group (NE-SG) to assist architects and designers determine where and how TCG  
81 technology, including the Trusted Platform Module, can best be used to secure network  
82 equipment such as routers, switches, and firewalls.

### 83 **1.1 Scope**

84 The *TCG Guidance for Securing Network Equipment* provides recommendations and detailed  
85 advice on how TCG standards should be used to secure network equipment such as routers,  
86 switches, and firewalls. Physical network functions are considered in this document;  
87 virtualized network functions are not considered in this version.

88 While the TPM's resistance to physical attack can help protect device identity, and can  
89 effectively prevent the leakage of credentials and other secrets, defense against physical  
90 attack is generally beyond the scope of this document.

91 TCG technology users are in the midst of a transition between TPM1.2 and TPM2.0; this  
92 document is constructed to cover both TPM1.2 and TPM2.0 applications, and to highlight  
93 differences when they're important.

94 There is a short glossary of terms in Section 7.

95

### 96 **1.2 Audience and Purpose**

97 Preserving the integrity and security of network equipment is essential to maintaining  
98 customer privacy and network reliability. In the face of increasingly sophisticated attacks on  
99 network equipment, Trusted Computing solutions are desperately needed in this area. Yet  
100 little information is available on how Trusted Computing should be used to secure network  
101 equipment and thus the networks that depend on this equipment. This reference document  
102 provides guidance for using TCG technologies to improve the security of network equipment  
103 and networks..

104

105 

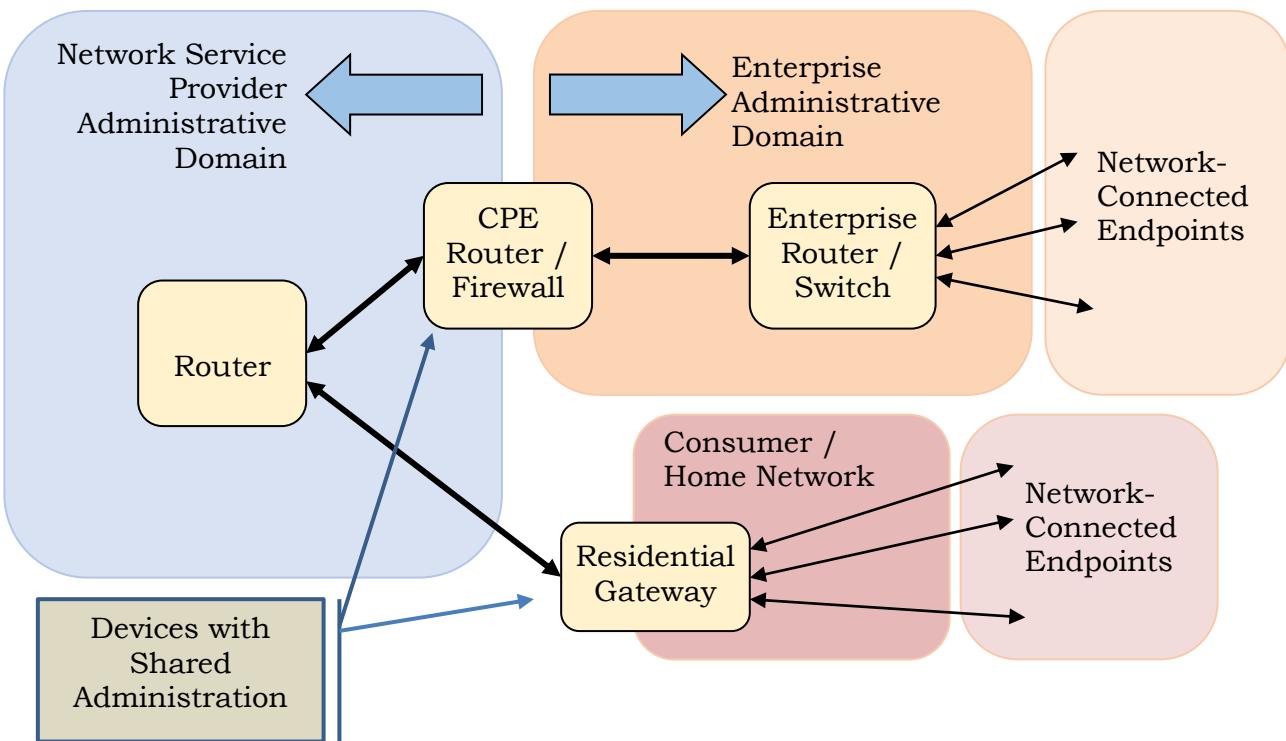
## 2. Preface

106

107 

### 2.1 Network Equipment Reference Model

108 Figure 1 shows a simplified reference model often used for Network Equipment. Customer  
 109 Premise Equipment (CPE) or Residential Gateways are often positioned between  
 110 administrative domains, and may require special attention for management of access and  
 111 identity.

112 

**Figure 1: Simplified Network Reference Model**

113

114 It should be noted that there are often many more administrative sub-domains within each  
 115 side of the domains depicted in Figure 1. Large enterprises in particular often have IT  
 116 departments that act as internal service providers, offering connectivity to departments and  
 117 groups, each of which may have considerable autonomy.

118 In this context, there may be several points in an overall network where different  
 119 administrative domains must share control of a single device; for example, an Enterprise  
 120 Managed Gateway CPE device might have interfaces configured by a Service Provider to  
 121 ensure contracted class of service, while the Enterprise would control interfaces that connect  
 122 to its own equipment.

123 Further, it should be noted that the administrator of a given device normally will not have  
 124 direct physical connectivity to the device, and will require authenticated remote access to  
 125 carry out management functions.

126

127    **2.2 Internal Structure of Network Equipment**

128

129    Networking equipment like routers and switches often comprises several major units:

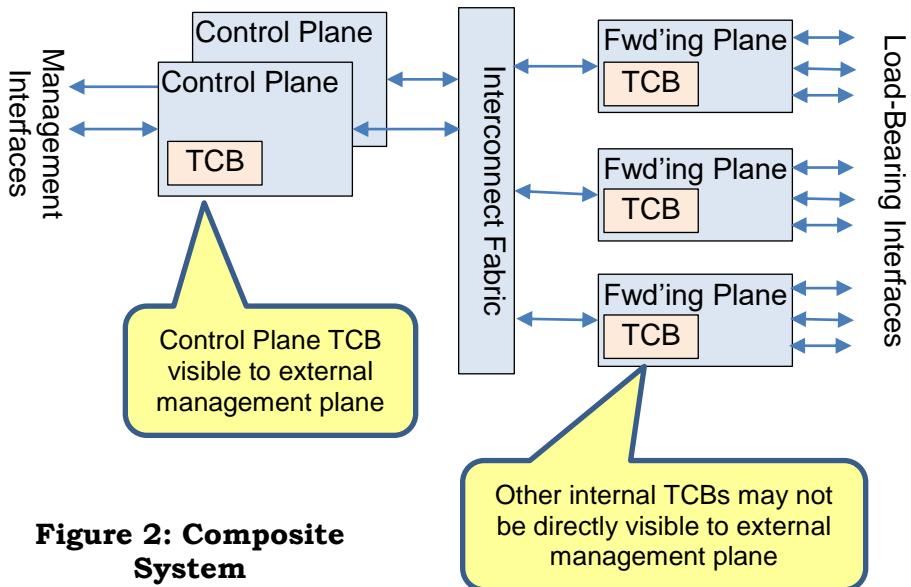
- 130    • The Forwarding Plane inspects each packet and performs the networking function  
131    required (e.g., L3 routing, L2 switching, Firewall security inspection, etc). The  
132    forwarding plane may be implemented in hardware, microcode, or software on a  
133    conventional processor, but in any case, the forwarding plane is typically subservient  
134    to a Control Plane, and would rely on the Control Plane for security.
- 135    • The Control Plane is the part of the device that initializes, monitors, configures and  
136    diagnoses the Forwarding Plane. The Control Plane is typically implemented as  
137    software running on an operating system, hosted on a conventional processor. The  
138    Control Plane typically manages every aspect of Forwarding Plane operation, including  
139    assurance that whatever code runs in the forwarding plane is authenticated.
- 140    • Some devices implement a separate Management Plane that provides an external  
141    interface used to configure and manage the device. A management plane would also  
142    typically execute in a conventional OS and processor environment.

143    Although essentially all networking devices have one or more control-plane processors that  
144    configure and run the device, the term *Network Device* can cover very small to very large  
145    devices:

- 146    • Many small networking devices are packaged as non-modular units with no field-  
147    replaceable units (FRUs). A small switch, wireless access point or customer premise  
148    device might all be non-modular. In this case, the identity (e.g. serial number) of the  
149    chassis encompasses all components of the device.
- 150    • Larger networking devices may have pluggable I/O or feature cards. The base chassis  
151    may have a unique identity, but the pluggable modules usually have their own serial  
152    numbers or identity. I/O or feature cards may also have distributed control plane  
153    processors that should be secured.
- 154    • Many large networking devices have field-replaceable control plane processors, often  
155    with several independent control plane processors for redundancy. In this case, each  
156    control processor may have its own serial number and identity. In some cases, to meet  
157    redundancy goals, the chassis itself may be completely passive, with no active  
158    components to prove identity.

159    Some architectures use a distributed control plane with many independent,  
160    cooperating processors; in that case, the identity and integrity of each processor may  
161    need to be determined separately.

162    Figure 2 shows an abstract view of *composite* networking devices, that is, those comprising  
163    several independent units. While there are many ways to construct a composite system,  
164    ranging from single-chassis devices with plug-in cards through to self-organizing networks of  
165    elements, a key factor in this context is that there are usually a small number of elements  
166    that serve as the external management interface for the system, combined with a larger  
167    number of internal units, which mostly are not directly reachable by the external  
168    management system. Each of the elements in Figure 2 may have its own Trusted Computing  
169    Base (TCB).



**Figure 2: Composite System**

170

171

172 Recommendations in this document relate to security, identity and integrity of the Control  
 173 Plane; if there is Networking Equipment where the forwarding plane can act independently of  
 174 the Control Plane, the forwarding plane may need to be protected as if it were a control plane  
 175 element.

## 176 2.3 Stakeholders

177 There are many stakeholders involved in the use and secure deployment of Network  
 178 Equipment:

- 179 • TPM Manufacturer
  - 180 ○ The TPM manufacturer creates and initializes the TPM silicon. TPM  
 181 manufacturers usually equip their devices with a unique Endorsement Key (EK)  
 182 and EK Certificate.
- 183 • Device Manufacturer
  - 184 ○ Device Manufacturers develop and sell Networking Equipment, often referred to  
 185 also as OEMs (Original Equipment Manufacturers).
  - 186 ○ The Device Manufacturer is usually responsible for supplying the equipment, a  
 187 software operating environment.
  - 188 ○ Device Manufacturers are often called upon to diagnose and remediate complex  
 189 networking problems encountered by Service Provider and Enterprise  
 190 customers.
- 191 • Network Service Provider
  - 192 ○ Service Providers generally provide public network service through defined  
 193 gateways.

- 194           ○ Service Providers may also offer private networking services to enterprises, often  
195            delivered over infrastructure also used for public interconnect.  
196           ○ Security and reliability may be offered as a value added service.  
197       • Enterprise  
198           ○ An Enterprise usually has network equipment for its own use, i.e., not to sell as  
199            a service.  
200           ○ An Enterprise may use specialized monitoring and configuration services  
201            provided by a remote Maintenance Provider as a way to outsource some of its  
202            own IT support.  
203       • Consumer  
204           ○ Network Equipment may be purchased by individual consumers or small  
205            businesses (e.g., WiFi gateway routers).  
206           ○ Consumer devices may need to take care of themselves, with little expectation  
207            for active administration.

208       In this document, the “Administrator” is considered to be the Service Provider, Enterprise  
209       or Consumer that has primary authority over the device, and can authorize or delegate  
210       access for other stakeholders, including Secondary Administrators. The Administrator  
211       must be able to erase private information without making the device unusable for a  
212       subsequent owner. (See Sections 3.11 and 5.11.)

213

## 214   **2.4 Key Differences between Network Equipment and PC Applications**

215       Networking Equipment almost always contains a general-purpose computing environment to  
216       configure and manage the device. But there are distinct differences between Networking  
217       Equipment and the common PC client and server applications:

- 218       • While Network Equipment may be highly modular, it is usually shipped as a closed  
219            embedded system, integrating hardware and software.<sup>1,2</sup>  
220       • The chain of security typically does not stop when the OS boots; what matters is  
221            security of the networking function that’s provided by the unit as a whole.  
222       • Network Equipment typically must boot and operate without manual intervention (e.g.,  
223            no Owner Password can be expected at boot time).  
224       • While Network Equipment has an important role in protecting user privacy, the  
225            equipment itself typically should not have an ability to hide or mask its own identity.  
226            See Section 2.4.2 for more on Privacy.

---

<sup>1</sup> Even though many Device Manufacturers support an SDK environment, and some are moving more aggressively to an open-source model, the manufacturer is still usually responsible for infrastructure security of the device.

<sup>2</sup> For a counter-example, see Open Compute Project Networking,  
<http://www.opencompute.org/projects/networking/>

227

228 **2.4.1 Long Life Cycle**

229 Network Equipment often has a long life cycle, and must stay operational in the network for  
230 many years. Certificates with 10-year expiration are often too short for many Network  
231 Equipment applications, as equipment may stay in operation for upwards of 20 years. Device  
232 Manufacturers should ensure that any immutable certificates have lifetimes long enough to  
233 extend through the expected life of the platform.

234

235 IEEE802.1AR Secure Device Identity [1] requires the following related to the “notAfter” field  
236 in a DevID certificate in Section 7.2.7.2

237 *Devices possessing an IDevID are expected to operate indefinitely into the future and should use the  
238 value 99991231235959Z. Solutions verifying a DevID are expected to accept this value indefinitely. Any  
239 other value in a DevID notAfter field are expected be treated as specified in RFC 5280.*

240

241 Network Equipment manufacturers should consider using long or indefinite lifetimes for any  
242 immutable certificates stored in devices (including the EK Certificate).

243 Along with certificates that don’t expire prematurely, manufacturers may want to consider  
244 mechanisms to accommodate cryptographic agility. Key lengths and algorithms do evolve  
245 over time, and mechanisms to update algorithms and replace dated certificates in fielded gear  
246 may extend useful equipment life.<sup>3,4</sup>

247 Conversely, Administrators may want to configure relatively short lifetimes for locally-  
248 generated certificates such as a Local Device ID (LDevID, see Section 5.1.2), or other vendor-  
249 specific applications, as a mechanism to ensure that devices don’t remain forgotten in their  
250 networks past the time when they should be replaced, or past the time when cryptographic  
251 parameters such as key length may have become outdated.

252

253 **2.4.2 Privacy and Networking Equipment**

254 Networking Equipment such as routers, switches and firewalls has a key role to play in  
255 guarding the privacy of individuals using the network:

- 256 • Packets passing through the device must not be sent to unauthorized destinations.

257 Some examples include:

- 258     ○ Routers often act as Policy Enforcement Points, where individual subscribers  
259         may be checked for authorization to access a network. Subscriber login  
260         information must not be released to unauthorized parties.

---

<sup>3</sup> See National Institute of Standards and Technology, *Report on Post-Quantum Cryptography, NISTIR 8105* [15]

<sup>4</sup> While Algorithm Agility is not part of TPM1.2, TPM2.0 can accommodate multiple cryptographic algorithms.

- 261           ○ Networking Equipment is often called upon to block access to protected  
262           resources, or from unauthorized users.
- 263       • Routing information, such as the identity of a router's peers, must not be leaked to  
264           unauthorized neighbors.
- 265       • If configured, encryption and decryption of traffic must be carried out reliably, while  
266           protecting keys and credentials.
- 267 Functions that protect privacy are implemented as part of each layer of hardware and  
268 software that makes up the networking device.
- 269 In light of these requirements for protecting the privacy of users of the network, the Network  
270 Equipment must identify itself, and its boot configuration and measured device state (for  
271 example, PCR values), to the Equipment's Administrator, so there's no uncertainty as to what  
272 function each device and configuration is configured to carry out<sup>5</sup>. This allows the  
273 administrator to ensure that the network provides individual and peer privacy guarantees.
- 274
- 275

---

<sup>5</sup> For example, serious privacy violations could occur if embedded software was modified to snoop traffic, or if unauthorized configuration commands were added to encrypt traffic with an incorrect key or to replicate traffic to an unencrypted tunnel.

276 

### 3. Use Cases

277 

#### 3.1 Device Identity

278 Providing reliable remotely-verifiable device identity for each piece of network equipment is a  
279 prerequisite for most use cases related to securing network equipment. This section contains  
280 a few use cases that demonstrate the value of device identity for network equipment and the  
281 variety of ways that device identity can be implemented and used. Within this section, the  
282 term “device” is used to mean a piece of network equipment.

283 Before getting into use cases, it’s helpful to distinguish two kinds of device identity:  
284 Manufacturer identity and Owner identity. The Manufacturer identity for a particular device  
285 is established, configured and managed by the Device Manufacturer, although it can also be  
286 used (e.g., verified) by the device owner. The Owner identity for a device is established by the  
287 device Administrator and is generally used only by the Administrator. Manufacturer identity  
288 is generally<sup>6</sup> unique across all products from that manufacturer (e.g., a model number and  
289 serial number) while Owner identity may be unique only within the Administrator’s domain  
290 (e.g., an asset number). These kinds of device identity are defined as Initial Device ID and  
291 Local Device ID, respectively, in IEEE 802.1AR (See Section 4.1).

292 In any case, device identity is generally implemented using a private key kept secret by the  
293 device and a certificate associated with that key, issued to the device by a Certification  
294 Authority (CA).

295 

##### 3.1.1 OEM Device Identity and Counterfeit Protection

296 Both network equipment owners and Device Manufacturers (OEM’s) need to verify the  
297 authenticity of network equipment, determining whether it is “counterfeit” (made by an  
298 unauthorized party or in an unauthorized manner) or “authentic” (made by authorized parties  
299 in an authorized manner).

300 **Business Drivers:** To achieve reliable network operation for end-user customers, Network  
301 Equipment Owners need to be assured that they are using authentic devices from known  
302 manufacturers. At the same time, Device Manufacturers want to reduce lost sales, protect  
303 their brand, prevent returns and warranty repairs of counterfeit equipment. Distributors  
304 want to detect counterfeit devices before they are resold. Auditors want to verify that network  
305 operators are meeting compliance regulations.

306 **How It Works:** During manufacturing, a signed manufacturer identity is instantiated in the  
307 device. Manufacturing processes and key types must be selected to ensure that the identity  
308 cannot be copied to, or instantiated in, another device. Later, the authenticity of the device  
309 can be verified by checking the validity of this identity.

310 This verification can be done by many parties, as detailed above: the network equipment  
311 owner, the manufacturer, the distributor, auditors, network operators, etc. The verification  
312 may be done locally (by connecting a cable directly to the device) or remotely (across a  
313 network).

---

<sup>6</sup> The Manufacturer and Administrator are responsible for managing uniqueness in their respective name spaces, but typically these identifiers will not be ‘cryptographically strong’; the public key associated with the device identity can also be used as a unique identifier where that’s a factor.

314 If a counterfeit device is detected, various actions can be taken: blocking the device from the  
315 network, alerting an operator, etc.

316

317 **3.1.2 Identity for Network Access**

318 **Business Drivers:** Network operators often wish to ensure that only authorized equipment is  
319 connected to their networks. This is true for telecommunications companies but also for cloud  
320 and data center operators, hospitals, enterprises, manufacturing facilities, and other  
321 environments where the network needs to be tightly controlled.

322 **How It Works:** Certificates proving device identity may be used to authorize network access.  
323 When a device connects to the network, its device identity is verified and checked against a  
324 list of authorized devices. If the device is on the list, it's permitted to connect to the network.  
325 Either identity configured by the Device Manufacturer, or identity configured by the  
326 Administrator may be used for this purpose.

327 In the case of CPE equipment, or other situations where administration is shared between  
328 two domains, it may be necessary for a Primary Administrator to grant rights to a Secondary  
329 Administrator to add their own identity.

330 The Identity for Network Access use-case requires the network equipment to authenticate  
331 itself to the upstream network's authentication service, and may require that the upstream  
332 network prove its identity to the particular networking device<sup>7</sup>.

333

334 **3.2 Secure Zero Touch Provisioning**

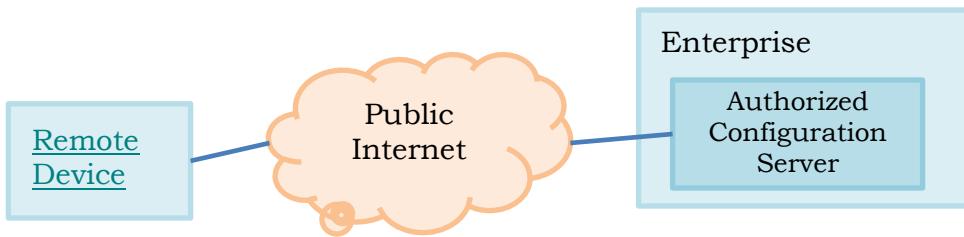
335 **Business Drivers:** There are many cases where a networking device may be shipped with no  
336 unique configuration, but must be configured before it can be used within a network.

337 In these cases, Zero Touch Provisioning (ZTP) may be desirable, requiring a mechanism where  
338 a newly-installed device initiates communication through a public network, but only enough  
339 to obtain the configuration information that would specify policy for operational use. As an  
340 example, downloaded configuration and keys might enable access to a corporate VPN, or  
341 might authorize access to restricted content.

342 Because the result of loading configuration on to the remote device could be access to an  
343 enterprise's private network, careful design is required to ensure that only authorized devices  
344 are able to download a configuration.

---

<sup>7</sup> For example, in the process of establishing an RFC 5246 Transport Layer Security (TLS) session, downstream and upstream devices might exchange DevID certs to prove identity.



345

346 Part of this use-case may require proof that the device is actually in the hands of an  
 347 authorized user before its configured, perhaps through the use of out-of-band authorization  
 348 information.

349 **How It Works:** ZTP can be made more secure with several additions:

- 350 • The networking device should use a Manufacturer-installed credential to authenticate  
 351 itself to the ZTP server.
- 352 • The networking device can check a credential supplied by the ZTP server to ensure  
 353 that it's authorized to serve configuration.
- 354 • The ZTP server can use Remote Attestation (Sections 3.7 and 5.7**Error! Reference  
 355 source not found.**) to verify that the device is running authorized software before  
 356 completing the configuration.

357

### 358 3.3 Securing Secrets

359 **Business Drivers:** Network equipment often contains sensitive information such as traffic  
 360 logs or cryptographic keys (e.g., shared secrets, passwords, VPN keys, SSL keys, and stored  
 361 data encryption keys). Disclosure of these secrets could result in disclosure of confidential  
 362 network traffic and privacy-sensitive information or even enable malicious tampering with  
 363 the network. Network operators (especially Service Providers and Enterprises) must protect  
 364 these secrets against disclosure to keep their networks secure and reliable and also to meet  
 365 regulatory or customer requirements for confidentiality and privacy.

366 **How It Works:** Because network equipment is often located in insecure locations and always  
 367 subject to network attacks and exploitation of software vulnerabilities, protection must be  
 368 provided against physical attacks, network attacks, and software compromise. Many  
 369 techniques can be used to protect keys or shared secrets, depending on the operations that  
 370 must be performed with them, especially throughput requirements.

371 There are two sub-cases that can be considered, depending on scaling expectations:

- 372 • Keys may be stored in the TPM and used in the TPM, providing substantial protection  
 373 against loss, but signing throughput limited by the type of key used, and the TPM's  
 374 low-power crypto engine.
- 375 • Alternately, the TPM may protect a key that's used externally by a high-throughput  
 376 crypto engine. In this case, the key is well protected by the TPM "at rest", but must be  
 377 released by the TPM before it can be used, increasing the risk that it can be snooped.

378 Using a Virtual Private Network gateway device as an example:

- 379     • A remote customer-premise device may log in to a corporate VPN across the internet.  
380       In this case, a small number of authentication operations may suffice, and keys may  
381       be kept and authenticated entirely within the TPM.  
382     • But the centralized VPN gateway may be performance-optimized for high-throughput  
383       authentication from thousands of remote devices; in this case, keys may be held in  
384       encrypted storage, but released for use by platform software.

385   This use-case covers the way these secrets can be managed.

386

### 387   **3.4 Protection of Configuration Data**

388   **Business Drivers:** Network Equipment usually requires configuration, often involving many  
389       parameters stored in a variety of files. The equipment Owner may wish to retain control over  
390       changes to configuration files on the equipment, with the goal of ensuring that unauthorized  
391       configuration changes don't compromise their network.

392   **How It Works:** In this case, the Owner may wish to sign and encrypt configuration files so  
393       each file can only be used on the intended device, and the device will only accept configuration  
394       from the authorized Owner. Signing or Attestation may also be used to ensure configuration  
395       integrity, i.e. to assure the equipment Owner that the configuration has not been modified.

396   In some cases, it may be desirable to prevent rollback to previous versions of configuration  
397       files.

398   Some Network Equipment could also benefit from full-disk encryption to protect data at rest,  
399       to discourage cloning. The document *TCG Storage Opal Integration Guidelines* [17] gives an  
400       overview of the Trusted Computing Group suite of specifications for self-encrypting storage  
401       devices.

402

### 403   **3.5 Remote Device Management**

404   **Business Drivers:** Network equipment Owners with a large number of devices often want to  
405       manage those devices remotely, including the ability to monitor devices and reconfigure them  
406       dynamically. Remote management and reconfiguration is especially important in modern,  
407       flexible computing environments that implement Software-Defined Networking (SDN) or  
408       Network Function Virtualization (NFV).

409   Remote Device Management may also include the need for trustworthy decommissioning or  
410       refurbishment (See Section 3.11).

411   **How It Works:** To perform remote device management securely, the device and management  
412       system must generally establish a secure communications channel with mutual  
413       authentication and confidentiality and integrity protection. This secure channel is used by  
414       the management system to monitor device status and send reconfiguration commands.

415

### 416 3.6 Software Inventory

417 **Business Drivers:** Most Network devices rely on complex embedded software to enable basic  
418 features as well as to enforce security policies. This software is often updated on devices  
419 already in the field, using releases and patches usually supplied by the device manufacturer.

420 Network Administrators are left with the task of keeping track of which devices have been  
421 updated to what revision level, sometimes tracking many independent components on a single  
422 complex device. Automating the process could save money and improve security by making  
423 it harder to overlook out-of-date releases.

424 **How It Works:** Mechanisms can be implemented to allow the Administrator to query devices  
425 to find which revision level of which components are installed on each network device in their  
426 network. The ISO specification *ISO/IEC 19770-2 Information Technology - Software Asset*  
427 *Management - Part 2: Software Identification Tag* [6], (also known as “SWID Tags”), coupled  
428 with the TCG protocol suite *Platform Trust Service* (PTS), can collect verified software image  
429 information from remote devices.

430

### 431 3.7 Attestation of Boot Integrity for Network Devices (“Health Check”)

432 **Business Drivers:** One extension to remote device management enabled by TCG technology  
433 is an ability to monitor the authenticity of software versions and configurations running on  
434 each device. This allows owners and auditors to detect deviation from approved software and  
435 firmware versions and configurations, potentially identifying infected devices.

436 **How It Works:** Networking equipment requires some extension from conventional PC-based  
437 Attestation due to more varied platform architecture, including redundancy, in-service  
438 software updates, and hot-swappable field replaceable components. In addition, the focus  
439 on overall system functionality calls for an extension of the conventional attestation plan to  
440 encompass more objects in PCRs, for example, hashes of application executables, OS  
441 daemons and configuration files.<sup>8</sup>

442 Attestation may be performed by network management systems. Networking Equipment is  
443 often highly interconnected, so it's also possible that Attestation could be performed by  
444 neighboring devices.

445 Attestation depends on a process sometimes called Measured Boot; See Section 4.1 for a  
446 comparison of boot security procedures.

447

### 448 3.8 Inventory of Composite Devices

449 **Business Drivers:** Many network devices are composed of one or more control or  
450 management units plus optional components like line processing units, feature-processing  
451 units and other kinds of field replaceable units (FRUs). The definition of a composite system  
452 and its internal organization, whether it's homogenous or heterogeneous, hierarchical or flat,

---

<sup>8</sup> Following a ‘chain of trust’ model, conventional TCG practice would be to use TPM measurements to prove that a reliable policy engine is running in the device, and then use that engine to verify application resources. However, for embedded devices like Network Equipment, the system designer has quite a bit of latitude in determining which software components are part of the trusted computing base.

453 is up to the system designer, but the behavior of the network device is based upon the  
454 composite behavior of individual components. The security posture of the network device is  
455 therefore only accurately represented by a composite measure that includes the posture of  
456 sub-components. Any change to one of these sub-components or even an undetected  
457 exchange endangers the security of the user's data on the network. A secure inventory allows  
458 determination of the components used and their state.

459 Many network devices allow FRUs to be replaced without triggering a complete system restart  
460 (often called 'hot swap'); for these devices, system-level reboots may be very rare, and the  
461 system's security posture must be re-evaluated every time an individual unit is inserted or  
462 removed from the system.

463 **How It Works:** Mechanisms may be implemented to provide proof of identity and  
464 configuration for each modular component, summarized and reported by the management  
465 process for the device.

466

### 467 3.9 Integrity-Protected Logs

468 **Business Drivers:** Various processes in the day-to-day operation of network equipment are  
469 based on information gathered from the system status of servers, routers and sensors. SACM,  
470 SIEM or even legal interception are based on state information of various components.  
471 Tampering with this information, mostly instantiated as log files, can impact the security  
472 protection (e.g. by suppressing intrusion-detection (IDS) data) or impact the integrity of  
473 information delivered by the legal interception interface.

474 Integrity-protected log files can be used by the management or external entities by providing  
475 information on the authenticity and integrity of the file. An integrity-protected log file can  
476 show on which device it was created, plus evidence that it has not been subject to  
477 unauthorized modification, and other details dependent on the use case.

478 This use case does not address the specifics of a SACM, SIEM or legal interception system  
479 and does not specify any part of such a processing entity, but shows the applicability of  
480 Trusted Computing for the protection of log files and other data used in later processes.

481 **How It Works:** A device's internal log file or system wide log server receives events and  
482 preserves these events for subsequent analysis. The data stored in the log file can be protected  
483 using TPM-based signatures expressing the origin of the data, along with the time of creation  
484 and the state of the device that created the data. All this can later be used in subsequent  
485 analysis to determine the trustworthiness of the data collected.

486 Confidentiality of log information is not covered in this use-case.

487

### 488 3.10 Entropy Generation

489 **Business Drivers:** Businesses and governments have a need to ensure that traffic that passes  
490 across untrusted network connections is reliably protected by strong cryptographic keys.

491 **How It Works:** Many networking protocols such as SSH and IPsec have a need for  
492 cryptographic-quality random numbers to avoid the generation of predictable session keys.

493 In addition, the TCP stack for Network Equipment should use good-quality randomness for  
494 the TCP window starting point as well as in the selection of ephemeral ports. These help to  
495 mitigate SYN and RST attacks against the device.

496 Network Equipment also often requires good-quality random numbers to generate  
497 Cryptographic Security Parameters (CSPs) such as RSA key-pairs.

498 Entropy sources used in Network Equipment may have to comply with standards such as  
499 *NIST Special Publication 800-90 A/B/C*[14] or BSI AIS-31, *A Proposal for Functionality Classes*  
500 *for Random Number Generators* [15] [17].

501 The TPM can act as a source for cryptographic-quality entropy.

502

### 503 3.11 Deprovisioning

504 **Business Drivers:** Networking Devices often contain information that's considered sensitive  
505 by the Administrator, such as customer configurations or routing policies. Once the device  
506 is taken out of service, this information must be reliably destroyed.

507 **How It Works:** Confidential information can include TPM keys themselves, or information  
508 encrypted by TPM keys. The TPM mechanisms for deleting keys can ensure that the  
509 confidential information will become inaccessible.

510

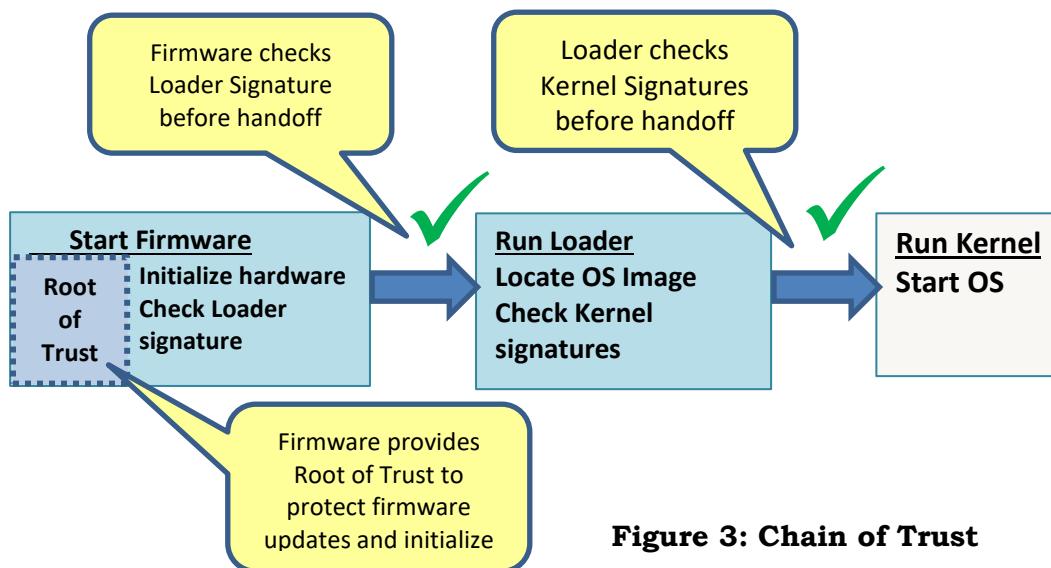
511

## 512 4. Building Blocks

513

### 514 4.1 Chain of Trust

515 Security of Network Equipment Devices is critically dependent on the integrity of the software  
516 running on the device. Software integrity is usually assured with a chain-of-trust model, with  
517 each stage during startup checking the next stage before execution, as shown in Figure 3.  
518 While that model is extensible to as many stages as needed, there's clearly a special case for  
519 the first step in the process, the one before which nothing exists to do any checking. This  
520 stage is referred to as the Root of Trust, and code that runs in the Root of Trust must be  
521 trusted implicitly.



522 **Figure 3: Chain of Trust**

523 Underpinning essentially all of the security mechanisms outlined in this document is a  
524 reliable Root of Trust, often implemented in code called a Basic Input Output System (BIOS),  
525 and usually stored in non-volatile memory so it's available to deliver the very first instructions  
526 to the processor for execution after reset. The root of trust has several functions, among  
527 which are initializing the TPM and doing the first measurement, and ensuring that firmware  
528 is not changed without authorization,

529 Although some systems are able to use an immutable root of trust, which cannot be modified  
530 or hacked, most current devices need a capability to upgrade the software that forms the root  
531 of trust, and must rely on some hardware assistance to lock the flash<sup>9</sup>, as well as code in the  
532 root of trust to protect the device from unauthorized updates or code implants.

533 Techniques for assuring protection of BIOS or other boot firmware can be found in NIST SP  
534 800-147, BIOS Protection Guidelines.

<sup>9</sup> Unlike ordinary software, the root of trust must be protected against attacks from the privileged run-time operating system. This usually results in some kind of hardware feature in the processor to lock flash against OS software.

535     Although TCG technology is not required to secure the root of trust, NIST SP 800-155 does  
536     specify the use of a TPM as part of a comprehensive software integrity strategy, linking to the  
537     Health Check features covered in Section 5.7 of this document.

538

539

540     

## 4.2 Secure Boot vs Measured Boot

541     There are generally two approaches to ensuring that code running on a Networking Device is  
542     authorized, and has not been subject to illicit modification.

- 543
  - **Secure Boot** (also known as Verified Boot) is a process by which each stage of the boot  
544             process checks a cryptographic signature on the next stage before executing it. In a  
545             typical system, there might be a BIOS that does a check of an OS Loader's signature,  
546             which in turn might do a check of the OS kernel before launching it.
  - **Measured Boot** (also known as Attested or Authenticated Boot) is a process by which  
547             each stage of the boot “measures” or computes and stores the hash of components in  
548             the next stage prior to launching it.

550     Secure and Measured Boot can be extended to run-time software through mechanisms such  
551     as the Linux Integrity Measurement Architecture [19]. See Section 5.7.1 of this document for  
552     more on the use of IMA.

553     Secure Boot is a subset of an overall *Secure Computing* architecture, which dictates a  
554     platform's behavior, as opposed to *Trusted Computing*, which enables the Administrator to  
555     deduce a platform's behavior. While implementations of secure boot are relatively common,  
556     complete Secure Computing or Trusted Computing environments are still difficult to design  
557     and implement.

558     Secure Boot and Measured Boot have some commonality and some differences:

- 559
  - While Measured Boot depends critically on the TPM to safely store *measurements*<sup>10</sup>  
560             (See Section 5.7), Secure Boot can be implemented without the use of TPM technology.
  - Both techniques rely on a Root of Trust in the initial BIOS, although Secure boot could  
561             be used to reliably instantiate the Roots of Trust required by Trusted Computing.
  - While measured boot won't stop a corrupted system from starting, it is able to verify  
564             the state of a broader range of components in the boot path (e.g. configuration files  
565             that might impact system security, in addition to executables).
  - Secure Boot and Measured Boot have different risk profiles. Secure boot can convert  
567             an innocent mistake such as a key mismatch into a network outage by refusing to  
568             start, while Measured boot can allow a corrupted network component to continue  
569             running, risking further damage.
  - “Sealing”, or encrypting data such that it can only be decrypted when a platform is in  
571             a known measured state, is a cost-effective way to protect data-at-rest, even with  
572             secure computing.

---

<sup>10</sup> A *measurement* is the cryptographic hash of an object such as an executable file

- 573     • A single system could reasonably execute both Secure Boot and Measured Boot at the  
574        same time. Measured boot and attestation enhance confidence in a computer's  
575        behavior even when it's not practical to cryptographically sign each and every object  
576        accessed as the system starts up.<sup>11</sup>

577 **4.3 802.1AR Device Identification**

578 [This Informative summary of 802.1AR appeared first in TCG Infrastructure WG *TPM Keys*  
579 *for Platform Identity for TPM 1.2*, May 2015 [4]; it's been edited in this section for wording and  
580 for applicability to TPM 2.0]

581 The IEEE 802.1 working group defined the IEEE 802.1AR Standard for Local and  
582 Metropolitan Area Networks Secure Device Identity and it was published on 22 December  
583 2009 [1]. This standard defines a per-device unique identity installed at manufacturing time  
584 and used subsequently in device-to-device authentication exchanges. This standards-based  
585 device identity can also be coupled in multiple ways with user identification.

586 The 802.1AR standard defines a secure device identifier (DevID) as "a cryptographic identity  
587 that is bound to a device and used to assert the device's identity". It further specifies:

- 588     • the DevID is an X.509 credential  
589     • globally unique per-device identifiers and the cryptographic binding of a device to its  
590        identifiers  
591     • the relationship between an initial identity installed during manufacturing and  
592        subsequent locally-significant identities installed by the Administrator  
593     • interfaces and methods for use of DevIDs with existing, and new, provisioning and  
594        authentication protocols

595  
596 The initially installed identity is defined as an IDevID ("I" for initial) and subsequently locally  
597 defined identities are LDevIDs ("L" for local). The IDevID is installed at manufacturing time.  
598 Since the "TPM Keys for Platform Identity" specification allows for an on-premise creation of  
599 the device identity, it will use the term DevID, whether the DevID is an IDevID or an LDevID.  
600 An IDevID will be created at manufacturing time and provides proof that this device has been  
601 manufactured by a certain manufacturer. An LDevID is created on the Administrator's  
602 premises and provides proof that this device is owned by a certain enterprise (or individual).

603 There is another dimension to the IDevID/ LDevID discussion with TPMs prior to and  
604 including version 1.2. In order to create an IDevID, the Device Manufacturer needs to take  
605 ownership of the TPM. However, commands that change ownership make all previously  
606 created keys unusable. Section 5.1.1.4 of this document outlines precautions that must be  
607 taken by the Device Manufacturer to ensure sure that the IDevID cannot easily be made  
608 unusable by the TPM. More on this topic can be found in *TPM Keys for Platform Identity for*  
609 *TPM 1.2* [4], Section 6.

---

<sup>11</sup> E.g., a Device Manufacturer can't sign a customer-mutable configuration file, but the Administrator may know exactly what should be in the file on a fielded machine. Measurement can report a fingerprint of the contents without requiring a signature on each file.

610 The 802.1AR specification DevID module (covering IDevID and LDevID) contains a service  
611 interface, storage holding a DevID secret and credential, secure hashing functions, a random  
612 number generator (RNG) and asymmetric cryptography functions. These functions exist in  
613 the TPM and calls by middleware (such as the TPM Software Stack (TSS) or the TPM API) can  
614 be used to meet interface requirements outlined in 802.1AR.

615 The IEEE Standard 802.1AR-2009 [1] can be used together with TPM-based keys and  
616 certificates. The TPM acts partly as the Secure Device Identifier Module (DevID Module) which  
617 the standard defines as “a logical security component that will secure, store and operate on  
618 one or more DevID Secret(s) [(private key)] and associated DevID credentials”. The document  
619 *TPM Keys for Platform Identity for TPM 1.2* [4] addresses usage of the TPM after provisioning  
620 has occurred and leverages TCG-specified X.509 certificate extensions to prove TPM residency  
621 of the keys, using the Subject Key Attestation Evidence (SKAE) extension [10]. The document  
622 *TPM Keys for Platform DevID for TPM 2* defines how TPM 2.0 key attributes and residency are  
623 proven for use as Device Identifiers.

624

625

## 5. Implementation Guidance

This section outlines approaches for the use of TCG technology to satisfy the Use Cases outlined above. These sections are not normative, but rather suggest ways that existing TCG technology can be put to use.

630

### 5.1 Device Identity

Device Identity requirements can be satisfied by the use of IEEE 802.1AR Device ID certificates (See Section 4.1 above). Appendix B of *IEEE Std 802.1AR-2009* [1] specifies implementation of DevIDs using a TPM, including a MIB and some required behavior (enable/disable for example) that is not specifically addressed in TCG documents.

**Figure 4** shows the relationship of credentials installed in the TPM during its lifetime:

- The EK and its Cert are added by the TPM manufacturer to show that the TPM is authentic.
- Initial Device ID (IDevID) credentials are installed by the device manufacturer to identify the device's manufacturer and serial number.
- The device manufacturer may also install an Initial Attestation key (IAK) and certificate to allow attestation without requiring the administrator set up a CA (See Section 5.1.1.3).
- The device manufacturer can also provide a mechanism to allow an Administrator to configure Local Device ID's (LDevIDs) and their own Attestation keys (LAK) (Section 5.1.2).

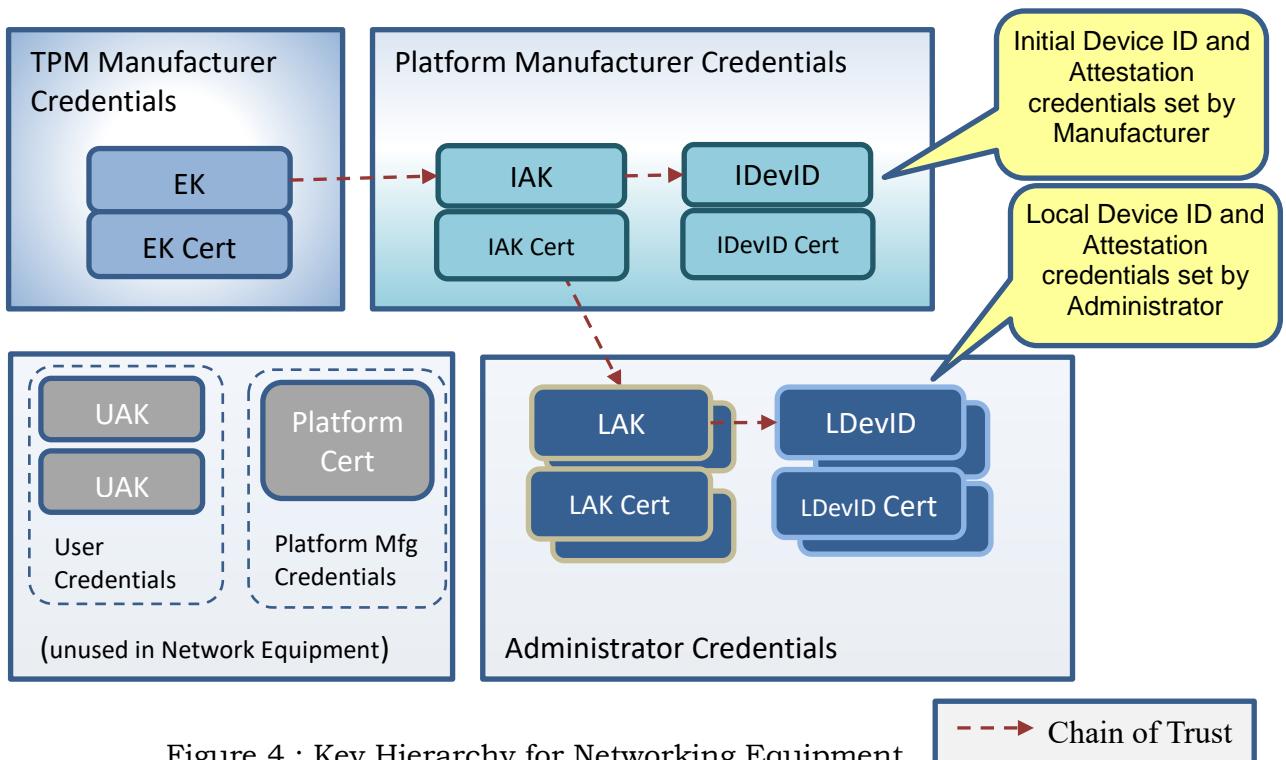


Figure 4 : Key Hierarchy for Networking Equipment

647

648

649

650 There are several TCG documents that are relevant to the use of TPMs to implement 802.1AR  
651 certificates:

- 652 • TCG Infrastructure WG: *TPM Keys for Platform Identity for TPM 1.2*. [2] This  
653 document gives details useful for DevIDs in TPM1.2, including a number of use-cases  
654 for DevID keys, and techniques for provisioning these keys.
- 655 • TCG Infrastructure WG: *TPM Keys for Platform DevID for TPM 2* [5] gives analogous  
656 details for TPM 2.0.
- 657 • TCG Infrastructure WG: *TCG TPM v2.0 Provisioning Guidance* [8]. This document  
658 gives guidance on TPM 2.0 hierarchies, etc, setting out expectations for TPM  
659 manufacturer, platform manufacturer and Administrator in the general case of TPMs  
660 in many kinds of equipment.

661

662 To support Device Identity requirements, the Device Manufacturer must configure the TPM:

- 663 • Either the TPM Manufacturer must ensure that the TPM is configured with an EK  
664 and EK certificate when the chip is manufactured, or the Device Manufacturer must  
665 generate the EK and EK Cert.
- 666 • Although this happens automatically with TPM2, a TPM1.2 must be configured as  
667 “always on”, i.e., Enabled and Activated, when the Device Manufacturer’s product is  
668 built.
- 669 • The Device Manufacturer must ensure that the TPM is configured with an IDevID key  
670 and an IDevID Certificate before the Device leaves the manufacturing facility.

671

672

### 673 5.1.1 OEM Device Identity and Counterfeit Protection

674 Device Identity and Counterfeit protection can be provided by an IEEE 802.1AR Initial Device  
675 Identifier (IDevID) certificate, signed by the Device Manufacturer and installed in the device  
676 while it’s still under the manufacturer’s control. This certificate provides cryptographic  
677 evidence that the specific physical device was manufactured by the specified manufacturer.

678 IEEE 802.1AR certificates include the device serial number and traceability to the  
679 manufacturer’s CA. The combination of manufacturer and Device serial numbers are  
680 expected to be unique.

681 An individual unit of Network Equipment (e.g. a large router) may comprise many field  
682 replaceable modules (FRUs). In this case, the serial number in the certificate should reflect  
683 the serial number of the FRU which houses the TPM. For example, in equipment with  
684 redundant control plane processors, the IDevID should contain the serial number of the  
685 control plane processor module itself, and will change if that processor unit is replaced for  
686 any reason in a fielded installation.

687

688

689

690 **5.1.1.1 Configuring Initial Identity Credentials**

691 To support Device Identity requirements, the Device Manufacturer must configure a  
692 certificate and keys in the TPM. There are differences between TPM1.2 and TPM2.0 in this  
693 process; this section outlines common elements, and the subsequent sections give the family-  
694 specific differences.

695 In all cases, the Device Manufacturer must ensure that the TPM is configured with an EK,  
696 EK Certificate, an IDevID key and certificate, before the network equipment leaves the  
697 manufacturing facility. The manufacturer may also want to configure an Initial Attestation  
698 Key (IAK) and cert (see Section 5.1.1.3).

699 To do this, the Device Manufacturer will create the IDevID key-pair in the TPM, issue a signing  
700 request containing the device serial number and other information to its own CA, and sign  
701 the certificate. Per 802.1AR, the signature should use SHA256. (See Section 5.1.1.4 for how  
702 to do this with TPM1.2). DevID keys must be created with modes that prevent duplication of  
703 the key from one TPM to another, to ensure that the identity of a Device can't be moved from  
704 one instance to another. (See Sections 5.1.1.4 and 5.1.1.6.)

705 *TCG EK Credential Profile For TPM Family 2.0 [9]* offers the following suggestion for EK  
706 credential lifetime, and the same advice should be considered for IDevID, or other immutable  
707 certificates. To quote Section 2.2.3:

708 **2.2.3 EK Credential Lifetime**

709 *An EK Credential contains fields that express the validity period of the credential. The  
710 validity period is at the discretion of the manufacturer. The credential is not expected to  
711 expire during the normal life expectancy of the platform in which it resides. The lifetime can  
712 vary widely between different types of platforms (e.g. while a typical validity period for a  
713 PC Client platform is 5-10 years, non-user device TPMs are expected to operate indefinitely  
714 into the future in which case the value 99991231235959Z should be used as expiration  
715 date). The credential lifetime can also depend on the lifetime of the TPM device and the  
716 algorithm type of the Endorsement Key. The time frame during which the security strength  
717 of the EK is acceptable SHOULD be taken into account by the manufacturer when  
718 determining the credential lifetime (e.g. see SP800-57[10]).*

719

720 It is not believed that there is any Personally Identifiable Information in Network Equipment  
721 PCRs, but that determination must be made by the Device manufacturer. (See Section 2.4.1)

722 **5.1.1.2 Signing and Encryption Key Types**

723 There are two types of keys that could be used for DevIDs in a TPM, as defined in *TPM Keys*  
724 *for Platform Identity for TPM 1.2* [2] Section 2.4, and *TPM Keys for Platform DevID for TPM2.0*.  
725 [5] Section 2.3:

- 726     • Signing keys, which are restricted to signing, and allow no other operations  
727     • General Purpose or Combined Keys (aka Legacy keys) which can be used for signing  
728         as well as encryption and decryption

729 The key type used for DevIDs should be selected based on the application and security needs  
730 for the equipment in question.

731 General-Purpose keys are required in applications that implement different TLS variants that  
 732 may use RSA or Diffie-Helman to convey the shared session secret during setup. Keys  
 733 capable of decryption are also required for features such as Protection of Configuration Data,  
 734 described in Section 5.4 in this document. Although 802.1AR does not impose a requirement,  
 735 Section 7.2.13 of *IEEE 802.1AR-2009 - Secure Device Identity* [1] does suggest *not* imposing  
 736 restrictions on DevID key usage.

737 Applications that are restricted to variants of TLS using Diffie-Hellman or Elliptic Curve Diffie-  
 738 Hellman (DH or ECDH), and are particularly concerned about attacks on keys in the TPM,  
 739 may want to use Signing keys for DevID, and provision separate Encryption DevID key-pairs  
 740 for other functions that need encryption with identity keys.

741 The trade-off is summarized in Table 1:

General Purpose Keys	Separate Signing and Encryption Keys
Fewer Keys to Keep Track Of	Better Security
Use General-Purpose (signing and encryption) keys for DevID	Use separate DevID keys for signing and encryption <ul style="list-style-type: none"> <li>• DevID-Signing for non-RSA variants of TLS</li> <li>• DevID-Encryption for RSA variants of TLS and other applications such as Protection of Configuration Data</li> </ul>
Requires less TPM space	More state to be stored in, or restored to, the TPM

**Table 1: DevID Key Types**

742 NIST Special Publication 800-57 Part 1, Rev 4 Section 5.2 gives additional guidance on use  
 743 of keys.<sup>12</sup>

745

#### 746 **5.1.1.3 Identity for Attestation**

747 TPM mechanisms are arranged so that in privacy-sensitive situations, Remote Attestation can  
 748 be done either without knowing the exact identity of the target system, or by trusting the  
 749 correlation between platform and attestation identity to a special Privacy CA (also known as  
 750 Attestation CA; see Chapter 7 of *IWG Reference Architecture for Interoperability (Part 1)* [20])

751 For Networking Equipment, where clear identity of the devices in question is usually a critical  
 752 requirement, a simpler approach can be used. In this case, the device manufacturer should  
 753 provision an *Initial Attestation Key* (IAK) and certificate that parallels the IDevID, with the  
 754 same device ID information as the IDevID certificate (i.e., the same Subject Name and Subject  
 755 Alt Name, even though the key pairs are different). This allows a quote from the device, signed

---

<sup>12</sup> See <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r4.pdf>

756 by the IAK, to be linked directly to the device that provided it, by examining the corresponding  
757 IAK certificate.<sup>13</sup>

758 Inclusion of an IAK does not preclude a mechanism whereby an Administrator can define  
759 Local Attestation Keys (LAK's) if desired (Section 5.1.2.2).

#### 760 **5.1.1.4 Initial Identity with TPM1.2**

761 TPM1.2 has relatively limited mechanisms for managing Ownership:

- 762 • For TPM1.2, the Device Manufacturer must configure TPM Ownership to make an  
763 SRK, so that an IDevID may be created as a child of the SRK.<sup>14</sup>
- 764 • To ensure that the IDevID credentials are not rendered invalid by deletion of the SRK,  
765 Ownership of the TPM must be locked.

766 In order to prevent the SRK from being deleted, rendering any keys stored under it useless,  
767 TPM Ownership can be locked with the non-volatile TPM\_DisableOwnerClear() ordinal. If the  
768 TPM supports physical presence, the volatile TPM\_DisableForceClear() ordinal may be used  
769 to block ForceClear on each startup.

770 The detailed structure of the 802.1AR certificate for TPM1.2 is given in *TPM Keys for Platform*  
771 *Identity for TPM1.2* [2]

772 The 802.1AR document requires the use of SHA256 for signature generation, although the  
773 TPM1.2 can only perform SHA1 signatures using its built-in functions. However, the TPM 1.2  
774 allows the host to compute the hash, prepend the appropriate ASN.1 encoding, and to submit  
775 this to the TPM for SHA256 RSA signature generation using Distinguished Encoding Rules  
776 (DER).<sup>15</sup>

777 TCG TPM1.2 documents do not define handles that should be used to identify DevIDs, so  
778 these must be assigned by the Device Manufacturer.

779 Identity keys in TPM1.2 should be created as *non-migratable*, or a Certified Migration Key  
780 (CMK), to eliminate the chance that a device identity could be improperly copied from one  
781 device to another.<sup>16</sup>

#### 782 **5.1.1.5 Key Types in TPM2.0**

783 In TPM 1.2, identity keys must be descendants of the Storage Root Key, but in TPM 2.0, there  
784 are a number of options.

---

<sup>13</sup> As long as the IDevID and IAK are signed by equally-reputable CAs (or more simply, the same CA) and contain the same Subject Name and Subject Alt Name, this approach eliminates the Asokan Attack, <https://tools.ietf.org/html/rfc6813>

<sup>14</sup> For TPM1.2, this is an unavoidable overlap of administrative domains; normally, Ownership would be taken by the Administrator, not the Manufacturer, but that's incompatible with a manufacturer-supplied IDevID

<sup>15</sup> See TPM1.2 Part I Design Principles, Section 31.2.2 TPM\_SS\_RSASSAPKCS1v15\_DER

<sup>16</sup> In TPM1.2, it's possible to use a Certified Migration Key (CMK) with suitable Migration Selection Authority (MSA) and Migration Authority (MA) to allow a key to be moved under controlled conditions. The point, of course, is to make sure that an unauthorized copy of a DevID cannot be made.

- Secondary, Internally-Generated keys are like DevID keys in TPM 1.2; the key pair is generated in the TPM, and can be configured so the private key is never released. A Secondary key is always the child of a parent key in the hierarchy.
- Primary keys are unique to TPM 2.0. These keys are also generated in the TPM, but they are generated using a deterministic Key Derivation Function that starts with a generic Template, and uses a primary seed in the TPM to generate the key pair. Primary keys can be deleted and regenerated from the template as long as the seed value is unchanged.
- TPM 2.0 also includes a specialized mechanism where an external entity can generate the key pair and wrap it so it can only be imported to a specific TPM. This can be used to circumvent the non-deterministic time taken to generate an RSA key inside the TPM, or it can be used to meet unique security requirements where an external key generator must be used. Imported key pairs cannot be Primary keys, and they cannot be marked as non-duplicable (i.e., they can't be marked as fixedParent). (Listed in the table below as "Secondary Imported")

Table 2 shows the different properties of the key generation approaches.

Primary Key	Secondary Internally-Generated	Secondary Imported
Must be created from a template; cannot be Imported	Generated internally	Key generated externally, wrapped with the parent key
RSA key generation is non-deterministic, and could take a minute or more of elapsed time		Importing keys is deterministic, as they're wrapped with a symmetric key
DevIDs should be Created as non-Duplicable (FixedTPM, FixedParent)		Imported keys can't be marked as Fixed, and require a Policy to block export or migration.
Key probably needs to be made Persistent to avoid long regeneration each power-on	Key can be saved, wrapped by its parent, then Loaded on power-up or use.	Importing the key produces a version wrapped by the parent, which can be quickly reloaded when needed.
Certs can only be generated with a multi-step Online process using Endorsement Credential and TPM2_ActivateCredential() <sup>17</sup>		Keys and certs can be generated externally and installed off-line across an air-gap if necessary.

<sup>17</sup> Since the key-pair is generated in the TPM, it's not known in advance, implying the need for a multi-step procedure to generate the key and then incorporate the public half in a matching certificate. Conversely, the public key of an imported key pair is known in advance, so a matching certificate can be delivered along with the wrapped key for import.

Key can be re-generated from a generic template in case of TPM2_Clear()	Keys are lost in case of a TPM2_Clear()	Wrapped key can be re-imported after a TPM2_Clear() (Assuming the parent is restored, and the Wrapped key is not lost)
---	---	--

**Table 2: Key Generation Options for TPM 2.0**

For the remainder of this section, we assume the following:

- In most applications, IDevID keys would be Primary keys installed by the device manufacturer. The Clear operation may be used to remove user identity and keys from a device, but the IDevID keys can be regenerated from a generic template after a TPM2\_Clear() operation.
- LDevID keys are Secondary, internally-generated keys, and are intended to be deleted by a TPM2\_Clear() operation.
- Manufacturers can also support Imported DevID keys if they have high-throughput manufacturing processes where the indeterminate wait for RSA key generation would be unacceptable, or for specialized customers that want to manage their own key generation.

In TPM 1.2, special Restricted Attestation keys must be used to sign internally-generated structures such as quotes, to prevent an external actor from spoofing a quote. The TPM 2.0 retains this capability with Restricted Signing Keys, which can sign an internally-generated quote or external data, as long as the external data doesn't contain the fixed pattern that matches an internally-generated structure. (See Trusted Platform Module Library Family 2.0 Part 1 [2], Section 25.1.2)

Thus, a Restricted Signing Key could be used for both Identity and Attestation, with the microscopic possibility that the TPM will refuse to sign a nonce that happens to look like a quote.

For the remainder of this section, we assume that Device IDs (IDevID, LDevID) are separate keys from Attestation keys (IAK, LAK).

### 5.1.1.6 Initial Identity with TPM2.0

To support Device Identity requirements, the Device Manufacturer must configure keys in the TPM as described in the *TPM v2.0 Provisioning Guidance* [8] document.

IDevID keys must be signing keys. If they are also Restricted, then they can serve as attestation keys. However due to complexities of using restricted keys with off-the-shelf crypto software stacks such as OpenSSL<sup>18</sup>, platform vendors may opt to provision IDevID keys as

---

<sup>18</sup> Signing external data with a Restricted key requires a Ticket, and there's some concern that managing tickets would require an involved OpenSSL rewrite. Not impossible, but real work.

832 unrestricted signing keys. Doing so would necessitate provisioning a separate restricted  
833 signing key pair for attestation purposes (identified as an IAK in Figure 3).

834 Except for the case of externally-generated keys noted in Section 5.1.1.5, IDevID keys should  
835 be Primary keys, stored in the Endorsement Hierarchy<sup>19</sup>.

836 The Device Manufacturer must ensure that the TPM is configured with an EK, EK Certificate,  
837 an IDevID Key and IDevID certificate before the network equipment leaves the manufacturing  
838 facility.

- 839     • The Device Manufacturer must provision the Endorsement Hierarchy  
840         authorization.  
841     • The IDevID certificate (if it is stored in the TPM) should be write-locked to prevent  
842         it from being deleted.  
843     • The IDevID Template does not need to be stored in the TPM, as it is generic and  
844         not specific to a unique device.

846 In order to prevent the EPS (Endorsement Primary Seed) from being changed, rendering any  
847 keys stored under it useless, the TPM2\_ChangeEPS() command must be blocked. This  
848 command is optional<sup>20</sup>, and ideally it would *not* be implemented by the TPM vendor. If it is  
849 implemented and enabled, then the Device Manufacturer must implement strict control,  
850 provisioning the Platform Authentication to a random value or un-fulfillable policy upon every  
851 platform reset.

852 TPM2\_Clear() resets the endorsement hierarchy authorization, clears the endorsement  
853 hierarchy (including IDevID keypair if stored here), and removes NV-defined indexes which  
854 are not protected by the TPMA\_NV\_PLATFORMCREATE flag. If the IDevID certificate is stored  
855 in the TPM, it should be placed in the Platform NV hierarchy, by setting its  
856 TPMA\_NV\_PLATFORMCREATE flag, so it's not lost by a TPM2\_Clear().<sup>21</sup>

857 To prevent permanent loss of the IDevID keypair, the Device Manufacturer could either block  
858 execution of this command or recreate the IDevID keypair from a generic template.

859

860 The exact format of an IDevID certificate for TPM2.0 is defined in Section 7 of TCG  
861 specification *TPM Keys for Platform DevID for TPM2.0*. [5]

---

<sup>19</sup> TPM2.0 embodies three independent key hierarchies: a “platform hierarchy” for platform protection, independent of specific users, an “endorsement hierarchy” for exposing information related to identity and a “storage hierarchy” for cryptographic usage related to specific users or applications.

IDevID keys are analogous to the EK, but without the privacy restrictions placed on the EK; as such, they should be stored using the same rules as the EK (i.e., in Endorsement). The Platform hierarchy could be used, but generally objects in the platform hierarchy should not be required after the platform OS has launched.

<sup>20</sup> The ‘optional choice’ is exercised by the TPM Manufacturer; if the TPM Manufacturer implements the ordinal, the Device Manufacturer needs to figure out how to block it.

<sup>21</sup> There is no Endorsement hierarchy for NV RAM, so the only two choices for the IDevID cert is Platform (which does survive TPM2\_Clear) or Storage (which does not). The EK Cert is stored in Platform; the IDevID cert should be too.

862    TCG *TPM v2.0 Provisioning Guidance* [8], Section 7.8 gives defined handles for IDevID  
863    credentials in TPM 2.0:

Description	Reserved Handles
IDevID Key	0x81020000
IDevID Certificate	0x01C90000

864

865    It must not be possible to move identity keys from one TPM to another. In TPM 2.0, the  
866    mechanism to prevent keys from being duplicated depends on the method of creation:

- 867    • If keys are generated within the TPM, FixedParent and FixedTPM must be set (see  
868      *TPM 2.0 Trusted Platform Module Library, Part 2: Structures*, Section 8.3 [3]).
- 869    • Imported keys can't be marked as Fixed, so they must be accompanied by a Policy  
870      that prevents export.

871

### 872    5.1.1.7    Loss of IDevID Credentials

873    OEMs should recognize that loss of IDevID credentials in a fielded device could be hard to  
874    remedy.

875    Although the IDevID key-pair must be protected by the TPM to ensure that the private key  
876    remains secret, it's up to the Device Manufacturer to store the IDevID certificate (and possibly  
877    a wrapped IDevID key) somewhere secure, but not necessarily secret, in the platform. OEM's  
878    could lock the IDevID certificate and wrapped key into NVRAM in a TPM, although there's no  
879    requirement to do so, and NV space is often at a premium.

880    For devices that don't have space in the TPM NVRAM to store Initial identity credentials, there  
881    are a couple of alternatives:

- 882    • If the manufacturer doesn't track device id information in its own database, loss of  
883      these credentials would require return to factory where Initial identity keys could be  
884      regenerated.
- 885    • If the manufacturer does keep track of EK's, it would be possible to re-create an  
886      IDevID in a fielded device using enrollment based on TPM Certify mechanisms, even  
887      if Ownership is lost. (See *TPM Identity Keys for TPM 1.2* [4] or *TPM Platform DevID for*  
888      *TPM 2.0* [5].)
- 889    • For TPM 1.2, as long as Ownership is maintained, the wrapped IDevID key and  
890      matching certificate could be replaced from a manufacturer's database. See Sections  
891      5.1.1.4, 5.1.1.6 on locking Ownership.

892    As noted in Section 2.4.1, manufacturers may want to support upgrade of cryptographic  
893    algorithms in devices with long lives. Cryptographic algorithms cannot be changed in  
894    TPM1.2, but TPM2.0 offers a mechanism to add algorithms, although doing so in fielded  
895    devices might require replacement of any 'permanent' identification.

896 **5.1.1.8 Compromise of Keys**

897 Private keys associated with the EK and IDevID are immutable; once manufactured, these  
898 keys are not easily changed in a fielded device. Furthermore, compromise is unlikely, as they  
899 are protected by the TPM.

900 But manufacturers do need to take care to protect intermediate keys stored in CAs in their  
901 manufacturing facility, since exposure of a signing key could compromise many devices, with  
902 little recourse to re-secure them short of return-to-factory.

903

904 **5.1.2 Local Device Identity**

905 IEEE 802.1AR also includes provision for Local DevID's, customer-specific credentials which  
906 can be installed by the Administrator of the Network Equipment, during or after installation.  
907 LDevID certificates will be signed by the Administrator's CA.

908 While the LDevID can contain anything the Administrator wants, they're probably best  
909 thought of as enterprise "asset tags", providing cryptographic evidence that the Device in  
910 question is one that was legitimately configured for use within the Administrator's enterprise.  
911 Multiple LDevID tags might be applied to a single device, depending on its organizational  
912 role(s).

913 The Device Manufacturer must provide a mechanism to remove Administrator-installed  
914 LDevID keys without deleting IDevID.

915 Modular Network Equipment may contain multiple field replaceable units; as with IDevID  
916 certificates, the basic LDevID can identify the FRU, but not necessarily the overall Device.

917

918

919 **5.1.2.1 How to Install a Local Device ID**

920 An Administrator may want to install LDevID credentials remotely, when the device is located  
921 off-site (as would be the case for Service-Provider's Customer-Premise Equipment, see Figure  
922 1). As such, care must be taken that the LDevID is being installed in the right device.

923 The following steps can be used to install an LDevID into a Device that's already equipped  
924 with an IDevID and IAK. These steps assume Enrolment over Secure Transport (EST; see  
925 IETF RFC 7030).

- 926 • Log into the Device using a TLS session that can't be easily hijacked, authenticated by  
927 the IDevID. Use that TLS session for all of the following steps:

928     ○ Use of the IDevID certificate for TLS authentication ensures that the connection  
929         terminates on the right device by showing the serial number assigned by the  
930         device manufacturer.

931     ○ The Device's identity as expressed by the manufacturer's serial number can be  
932         trusted if:

933         ▪ The certificate can be traced to the Device Manufacturer.

934         ▪ The Device can respond to a challenge, proving it holds the IDevID private  
935         key.

- 936           ○ Check the Manufacturer's serial number against a bill-of-sale to ensure it's the  
937           right device.
- 938           ○ These steps provide proof that there's a secure connection to the exact device  
939           into which the LDevID should be placed.
- 940       • The device should generate the LDevID key pair and a CSR containing the public key  
941           and the asset information selected by the Administrator.
- 942           ○ The CSR may also contain TPM\_CERTIFY\_INFO signed by the IAK key to prove  
943           that the LDevID key is in the same TPM as the IDevID<sup>22, 23</sup>.
- 944       • The CSR should be returned to the Enterprise CA.
- 945           ○ Since an IDevID key must be configured to prevent duplication from one TPM  
946           to another, the enterprise CA can use the TPM\_CERTIFY\_INFO to convince itself  
947           that the IDevID and LDevID are in the same TPM, and that TPM is in the  
948           intended Device.
- 949       • The Enterprise CA can then create the LDevID certificate with OIDs showing TPM  
950           residency, sign it and send it back to the device.

951 This procedure is defined in more detail in *TPM Keys for Platform DevID for TPM2.0*. [5].

952 An analogous procedure can be used for cases where an externally-generated key pair is  
953 desired for LDevID (i.e., LDevID doesn't have to be generated in the TPM for those cases where  
954 an externally-generated key would be more appropriate).<sup>24</sup> This case is also covered in *TPM  
955 Keys for Platform DevID for TPM2.0*. [5], Section 6.2.

956 Chapter 21 of the book *A Practical Guide to TPM 2.0* (W. Arthur & D. Challener) [13] also  
957 outlines techniques that can be used to remotely provisioning DevID credentials.

### 959       **5.1.2.1.1     Conveying TPM Certify Information**

960 While there could be a number of ways to get the TPM\_CERTIFY\_INFO from the TPM to the CA  
961 responsible for signing a DevID, this document suggests uniform practice between TPM 1.2 and  
962 TPM 2.0.

---

<sup>22</sup> This proof is simple if the IAK contains the same device serial number as the IDevID, and they're signed by the same issuing authority. Other mechanisms must be used if the IAK is created independently. See Section 5.1.1.3.

<sup>23</sup> Use of TPM Certify yields a trustworthy mechanism to get the LDevID into the same TPM as the corresponding IDevID, even if the control plane software on the device can't be trusted. If TPM Certify is not used, there's a possibility of variations of the Asokan attack, where the Device might have a valid IDevID, but do something that misuses the LDevID, e.g. not putting it a TPM at all, or passing it off to some other device. TPM Certify must be verified by a CA that puts the TCG OID indicating TPM key residency into the resulting certificate.

Note that TPM Certify info in TPM2.0 (TPMS\_CERTIFY\_INFO) is quite different from TPM Certify info for TPM1.2 (TPM\_CERTIFY\_INFO and TPM\_CERTIFY\_INFO2), but they both satisfy the requirement of proving co-residency.

<sup>24</sup> We note that externally-generated keys don't necessarily have to risk exposure. An HSM (or another TPM) could generate an LDevID key pair, and wrap it with a key known only to the TPM destined to receive the LDevID.

- 963     • For TPM 1.2, TPM\_CERTIFY\_INFO is added to the Certificate Signing Request as an  
964       X.509 extension. This extension is defined in Subject Key Attestation Evidence (SKAE)  
965       [10].  
966     • The TPM2\_Certify structures are more complex, but a new CSR extension is [will be]  
967       defined in *TPM Keys for DevID for TPM 2.0* [5].  
968

### 969     **5.1.2.2     Attestation with Local Identity**

970     Administrators who wish to use LDevIDs for device identity, using their own credentials to  
971       identify devices, may want to provision “local” Attestation Keys (LAK’s) as well, allowing the  
972       Attestation system to reliably identify the device within the Administrator’s context. Use of a  
973       Local AK in addition to the LDevID can prove that Attestation has not been compromised by  
974       an Asokan attack (RFC 6813).

### 976     **5.1.3     Identity for Network Authentication**

977     When presenting identity for network access, the Device may use its IDevID or one of its  
978       configured LDevID certificates, at the discretion of the Administrator. The appropriate  
979       certificate should be chosen by methods controlled by the Device’s operating system software.

980     IETF’s Transport Layer Security (RFC-5246) Section 7.4.4 defines the mechanism by which a  
981       server can negotiate trust anchors, by, for example, specifying a list of the distinguished  
982       names of acceptable certificate authorities.

### 983

### 984     **5.1.3.1     Proving a Link to the TPM**

985     DevIDs may be used to authenticate access to various kinds of network services, in which  
986       case the receiving service must decide how much trust can be placed in the credential.

987     DevID private keys that can be shown to be held by a TPM can presumably be judged more  
988       likely to be credible, and less likely to have been compromised.

989     There are two techniques to demonstrate to a network service that a particular DevID key is  
990       stored in a TPM:

- 991
- 992       • Sign the DevID cert with a CA that will only be used to sign certificates that are  
993           known by the Administrator to originate in a TPM. This may include a specific TCG  
994           OID corresponding to the Certificate Practices Statement (CPS), referenced below in  
995           this section.
  - 996       • Extend the service’s login authentication to parse an additional x.509 extension field  
997           called a Subject Key Attestation Extension (SKAE, TPM1.2 only).

998

999     The use of SKAE to send proof to a CA that a DevID key is stored in a TPM1.2 device is  
1000       described in *TCG Infrastructure Workgroup: Subject Key Attestation Evidence Extension* [10].

1001     The CA can certify that the key is in a TPM using the mechanism described in *TCG*  
1002       *Infrastructure WG TPM Keys for Platform Identity for TPM 1.2* [2]

1003 To quote Section 3.1.1.1:

1004 *In order to assert that the CA verified TPM residency for the key used in a DevID certificate, a  
1005 certificate policy OID unique to the TCG (and for this purpose) MUST be included in the certificate  
1006 policy extension of the DevID certificate. The CA MUST perform verification of the AIK signature  
1007 over the TPM\_CERTIFY\_INFO or TPM\_CERTIFY\_INFO2 prior to including the certificate policy TCG  
1008 OID.*

1009 *The TCG Policy OID is: 2.23.133.6.1.2*

1010

1011 The analogous method to certify that keys are resident in a TPM is described in the TCG  
1012 document *TPM Keys for Platform DevID for TPM 2* [5].

1013 **5.1.3.2 Why Not Just Use EK and Platform Cert for Identity?**

1014 While the Endorsement Key (EK) must be present in every TPM, and must be unique, that  
1015 doesn't mean it's a good candidate for device identity.

- 1016 • The conventional use of the EK credential is to prove the provenance of the TPM  
1017 itself, by linking to a chain of certificates rooted at the TPM manufacturer's CA.  
1018 • But beyond that, the EK can't be used directly as a signing key to prove possession of  
1019 the EK secret key; to protect privacy in privacy-sensitive applications, that proof can  
1020 only be done indirectly through other keys using TPM Certify information.

1021 As a result of these restrictions on EK use, a separate set of credentials traceable to the device  
1022 manufacturer is used to identify the device.

1023

1024 Similarly, some Device Manufacturers may supply a Platform Certificate including a  
1025 manufacturer's serial number (see *TCG Credential Profiles For TPM Family 2.0* [12]). While  
1026 that's similar to a DevID, again, they're not the same:

- 1027 • The Platform certificate is an *Authorization Certificate*, not a Public Key Certificate  
1028 (See RFC 5755).  
1029 • An Authorization Certificate isn't bound to a key pair, so it can't be used for proof of  
1030 identity. The Platform credential does identify the associated EK credential, but has  
1031 no key of its own for signing.

1032 In some cases, the platform credential might be used as a link in the chain to create and  
1033 install an IDevID after the device has left the manufacturer's premises, although that use is  
1034 not described in this document.

1035

1036 **5.2 Secure Zero Touch Provisioning**

1037 There are many cases where a networking device may be shipped with no unique  
1038 configuration, but must be configured before it can be used with a network.

1039 In these cases, Zero Touch Provisioning (ZTP) may be desirable, requiring a mechanism where  
1040 the device communicates through the network as soon as it's activated, to obtain the  
1041 configuration information that would specify policy for operational use. As an example,  
1042 downloaded configuration might enable access to a corporate VPN, or might authorize access  
1043 to restricted content.

1044 Part of this use-case may require proof that the device is actually in the hands of an  
1045 authorized user before it is configured, perhaps through the use of an out-of-band  
1046 authorization code.

1047 The IETF *Zero Touch Provisioning for NETCONF or RESTCONF based Management* and  
1048 *IETF Autonomic Networking Integrated Model and Approach (anima)* bootstrapping work  
1049 provide possible frameworks in which the configuration process can work.<sup>25</sup> See:

1050 <https://datatracker.ietf.org/doc/draft-ietf-netconf-zerotouch/>

1051 <https://datatracker.ietf.org/doc/draft-ietf-anima-bootstrapping-keyinfra/>

1052  
1053 TCG technologies can improve the security of ZTP in two ways:

1054 • A DevID key stored in the TPM (Section 5.1) can provide cryptographic assurance of  
1055 the identity of the device attempting to register for configuration.  
1056 One way to do this is described in IETF RFC 7030, “*Enrollment over Secure Transport*”.

1057 • The Trusted Network Connect IF-M protocol suite can be used to offer assurance to  
1058 the central configuration system that the Device is running authorized software. This  
1059 mechanism is described in IETF RFC 5792, “*PA-TNC: A Posture Attribute (PA) Protocol*  
1060 *Compatible with Trusted Network Connect (TNC)*”. See also Section 5.7 below.

1061  
1062

### 1063 5.3 Securing Secrets

1064 Network Equipment may need to manage many different kinds of secrets, and it may not be  
1065 possible to use a single mechanism to secure all of them.

1066 The same mechanism used to create and store LDevID private keys inside the TPM can be  
1067 used for other uses of asymmetric keys, as long as the rate at which challenges must be  
1068 signed is modest (around five per second for many hardware TPMs<sup>26</sup>). An advantage of  
1069 keeping these keys inside the TPM is that no amount of bus-snooping will reveal the secret.

1070 In cases where higher throughput is needed (e.g. a VPN Gateway that must validate many  
1071 thousands of sessions, at a rate of hundreds of sessions per second), or applications where  
1072 shared-secret symmetric keys are used, it will be necessary to pass secrets to cryptographic  
1073 functions outside of the TPM. In that case, secret keys can be stored in the TPM using a  
1074 mechanism called Sealing, where keys or other secrets are stored in the Device’s file system  
1075 in an encrypted file that can only be decrypted with keys released from the TPM when pre-  
1076 defined criteria are met.

1077 Sealing is a mechanism provided by TPMs to encrypt small data objects (typically symmetric  
1078 keys) using a key that’s kept in the TPM. Sealing and Unsealing offer a couple of capabilities:

---

<sup>25</sup> As of May 2017, these works are still in Draft format.

<sup>26</sup> See *Waltzing the Bear, or: A Trusted Virtual Security Module*, Ronald Toegl, Florian Reimair, and Martin Pirker, pg 155 for some measurements of hardware TPMs.

[https://online.tugraz.at/tug\\_online/voe\\_main2.getvolltext?pCurrPk=67562](https://online.tugraz.at/tug_online/voe_main2.getvolltext?pCurrPk=67562)

- 1079 • The object to be encrypted by the TPM is unstructured; that is to say, the TPM regards  
1080 the object simply as a string of bytes. In TCG documents, this data object is referred  
1081 to as a ‘blob’.
- 1082 • The encrypted object is not stored in the TPM; it’s sent to the TPM along with keying  
1083 and access information, using a “Sealing” operation, and the TPM returns the object  
1084 encrypted, where it can be stored in the host platform’s file system.
- 1085 • The sealed object can be unsealed when returned to the TPM, assuming a number of  
1086 conditions are met:
  - 1087     ○ The relevant asymmetric key must be resident in the TPM.
  - 1088     ○ Authorization to use the parent key must be present.
  - 1089     ○ In addition, the sealed object would usually require that one or more of the PCRs  
1090         have a particular value, indicating that the platform is in a known state.<sup>27</sup>

1091 If the conditions are met, the TPM will return the decrypted blob, which the system designer  
1092 can use as a structure containing symmetric keys for decrypting a larger object. Alternately,  
1093 the decrypted blob could contain keying material for use in high-throughput accelerators or  
1094 other applications.

1095 The system designer has complete freedom in choosing which, if any, PCRs should be  
1096 consulted prior to unsealing the object. Here are two examples:

- 1097 1. Some secrets need to be accessed just once, early in the boot sequence of a platform.  
1098 In that case, the system designer could allocate one PCR specifically for this use, and  
1099 seal the secret against this PCR with a value of zero (the value after reset). After  
1100 retrieving the secret, the system designer could extend that PCR with an arbitrary  
1101 value, preventing the object from being unsealed again by a hacker.
- 1102 2. The system designer could seal the secret against a number of PCRs that represent the  
1103 exact versions and configurations of software booted on the box, preventing the secret  
1104 from being revealed unless the OS is in the exact right configuration. This approach  
1105 needs to be used with caution, as a software update that’s not flawlessly staged could  
1106 render the secret permanently unintelligible.

## 1108 5.4 Protection of Configuration Data

1109 Network Equipment often has many configurable settings that must all be coordinated to  
1110 achieve secure and reliable network connectivity.

1111 Organizations wishing to enforce network policies may want to create configuration files  
1112 centrally, and distribute them to routers in a way that they can’t be changed or moved from  
1113 one router to another.

1114 Transport-level security (e.g. a TLS session authenticated with a DevID) would normally be  
1115 used with network-connected devices to ensure that the right configuration goes to the right

---

<sup>27</sup> Technically, an object can be sealed without requiring reference to any PCR, but that would have the same effect as a simpler encrypt/decrypt operation.

1116 device, but an alternate approach to this that can work across an air-gap is to encrypt and  
1117 sign configuration files so they can only be decrypted and used on the intended router.

1118 This can be accomplished if the organization managing the network devices creates a  
1119 configuration file for each device, and encrypts each one using a symmetric key, which in  
1120 turn can be encrypted using the public portion of an IDevID or LDevID key corresponding to  
1121 the desired device.<sup>28</sup> Authenticity can be assured by signing the package.

1122 Once installed on the device, the configuration file can be decrypted using the TPM\_UnBind()  
1123 (for TPM1.2) or TPM2\_Decrypt() (for TPM2.0) to decrypt and execute the file.

1124 Administrators must build in enough agility in the provisioning process to encrypt  
1125 configuration files with new keys when fielded units must be replaced.

1126

## 1127 **5.5 Remote Device Management**

1128 A management station should be able to perform checks at any time to verify the identity,  
1129 software and configuration of each device. Classic issues like missed updates and deviations  
1130 in configuration parameters can be detected in an automated way, making established  
1131 management systems more secure in the following ways:

- 1132 • Trusted Computing allows monitoring and verifying the state of the devices against  
1133 expected values documented in the management system. Deviations can be reported  
1134 automatically to the Administrator using Remote Attestation and the PTS protocol  
1135 suite (see *TCG Attestation PTS Protocol: Binding to TNC IF-M* [11]), reducing  
1136 operational costs to the customer.
- 1137 • Remote Attestation can also increase confidence in the supply chain, by proving that  
1138 devices have the correct software loaded.
- 1139 • Using a DevID to authenticate the device tightens the ability of the management  
1140 station to check the inventory of devices deployed. Unauthorized or changed devices  
1141 can be detected without manual checks of the physical devices.

1142

## 1143 **5.6 Software Inventory**

1144 The International Standards Organization (ISO) has published document ISO/IEC 19770-  
1145 2:2009 *Information technology -- Software asset management -- Part 2: Software identification*  
1146 *tag* [6] describing the format for Software Identification tags (SWID tags), stored in XML files  
1147 on a device. SWID tags encode version information and expected hashes for each software  
1148 component, and should be updated by the device manufacturer's software update process  
1149 each time a new image or patch is installed.

1150 SWID tags can be retrieved from a device to determine which software was last installed on  
1151 the device. Although techniques beyond the scope of this document must be used to  
1152 determine if the software installed is the version expected, comparison of the hashes in the

---

<sup>28</sup> Some customers may prefer to use separate signing and encryption keys, and not have a single DevID that must be configured for both. See Section 5.1.1.2 for tradeoffs.

1153 SWID tags with hashes retrieved from the TPM would provide assurance that the installed  
1154 software matches a valid release from the manufacturer.<sup>29</sup>

1155 While SWID tags are simply files and can be retrieved with any number of different  
1156 mechanisms, the TCG Trusted Network Communications workgroup has incorporated SWID  
1157 tag retrieval into the TNC protocol suite. Specified in *SWID Message and Attributes for IF-M*,  
1158 the TNC protocol allows a management station to determine whether each network device is  
1159 loaded with the expected release.

1160 The presence of a particular SWID tag indicates which software version is probably installed.  
1161 For purposes of attestation, the SWID tags must be signed by the network device  
1162 manufacturer, or potentially by the Administrator of the device or other organizations,  
1163 providing authority that the hashes claimed by SWID values are the ones that should be  
1164 present on the device.

1165 SWID tags contain hashes of the software modules they describe; with careful alignment by  
1166 the network device manufacturer, these hashes in SWID tags can be used in some cases as  
1167 “known good values” for Attestation, and can be compared to what was recorded in the TPM  
1168 using a measured-boot process.

1169 This allows a degree of automatic checking as part of the Attestation procedure described in  
1170 Section 5.7:

- Use the TNC protocol described in *SWID Message and Attributes for IF-M* to retrieve signed SWID values showing what should be installed on the device.
- Use TNC protocol *TCG Attestation PTS Protocol: Binding to TNC IF-M* [11] to retrieve signed PCR values from the TPM, along with boot logs, to determine what was actually loaded during the boot process.

1176 Signatures, module names and hashes can then be compared by the external Attestation  
1177 Server to detect tampering.

1178

1179 The following documents give details on SWID tags:

- *ISO/IEC 19770-2:2009(en) Information technology — Software asset management — Part 2: Software identification tag* [6]
- *Guidelines for the Creation of Interoperable Software Identification (SWID) Tags*; NIST IR 8060 (<http://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8060.pdf>)
- *TCG Trusted Network Connect: SWID Message and Attributes for IF-M*, Version 1.0, Revision 27, 2 June 2015

1186 [http://www.trustedcomputinggroup.org/files/resource\\_files/F47E6340-1A4B-B294-D0B3F7FF6497F489/SWID\\_Messages\\_For\\_IFM\\_v1r29.pdf](http://www.trustedcomputinggroup.org/files/resource_files/F47E6340-1A4B-B294-D0B3F7FF6497F489/SWID_Messages_For_IFM_v1r29.pdf)  
1187 [http://www.iso.org/iso/catalogue\\_detail.htm?csnumber=53670](http://www.iso.org/iso/catalogue_detail.htm?csnumber=53670)

1189

1190

---

<sup>29</sup> Of course the definition of what's being hashed and the hash algorithms have to align to be useful.

## 1191 5.7 Attestation of Integrity for Network Devices (“Health Check”)

1192 Attestation is a mechanism by which *measurements* stored in a TPM can be retrieved by a  
1193 remote attestation server to determine the configuration of the equipment at boot time as well  
1194 as at runtime. In the case of Network Equipment, remote attestation might be invoked by a  
1195 management station (called the Challenger in TCG documents), causing a networking device  
1196 to report its boot (and runtime) configuration. Alternatively, two peer devices might do Mutual  
1197 Attestation to prove to each other that their configurations are acceptable.

1198 Attestation of both boot time and run time configuration can include a wide range of objects  
1199 that might have a bearing on the security posture of a device.

1200 Boot time objects typically include:

- 1201 • The initial firmware loaded (such as UEFI code)
- 1202 • Options of the initial firmware
- 1203 • Boot loader
- 1204 • OS kernel

1205 When extended to run time, Attestation can include:

- 1206 • Executables launched by the OS
- 1207 • Shared libraries loaded into memory for execution
- 1208 • Policy or configuration files
- 1209 • Scripts
- 1210 • Cryptographic key material and credentials used by the OS, or other relevant sources  
1211 of state information.

1212 As a result of an attestation process, the Challenger will receive a cryptographically signed  
1213 “quote” and log from the Device indicating which versions of which modules were loaded and  
1214 run as the equipment started up.

1215 Attestation requires that the Device implements “Measured Boot” (Section 4.1), in which,  
1216 starting at a Root of Trust for Measurement (RTM), each stage computes a hash of the next  
1217 stage and stores the result in a TPM Platform Configuration Register (PCR) before launching  
1218 it. Each stage of the process can also measure any ancillary files that may be of interest in  
1219 determining the device’s security state, as listed above. Along with extending measurements  
1220 into PCRs, device software is expected to keep a log file which contains the name and version  
1221 of the object along with the hash. The log file does not need to be secure; the management  
1222 station can confirm its contents by computing the hash of the hashes and comparing the  
1223 result to the signed PCR values. While the contents of the log file are not specified, the log  
1224 data structure used by PTS is defined in Section 3.25 of *TCG Attestation, PTS Protocol: Binding*  
1225 to TNC IF-M [11], and could serve as a template.

1226 To obtain attestation information, the Challenger must establish the identity of the Device,  
1227 and then request the signed quote plus log information. This can be done using the TCG  
1228 Platform Trust Services (PTS) protocols, defined in *TCG Attestation PTS Protocol: Binding*  
1229 to TNC IF-M.

1230 Several components are needed for Attestation:

- 1231 • The Device must extend the hash of each boot object into the appropriate PCR, and  
1232 keep the corresponding log entries. If runtime integrity must be attested as well, the

1233        Device must extend the hash of each runtime object into the appropriate PCR, and  
1234        also keep the corresponding log entries.

- 1235     • The Device must have an Attestation Key (either IAK or LAK) configured, to allow it to  
1236        sign the quote.  
1237     • Either a Privacy CA must be used to certify an AK, or the AK should be otherwise  
1238        traceable to the device's identity.<sup>30</sup> (See Sections 5.1.1.3 and 5.1.2.2.)  
1239     • The device must implement the PTS protocol (or equivalent) to deliver attestation  
1240        information.

1241

1242        TCG specifications allow the Challenger to obtain attestation information from the Device,  
1243        but it's the Challenger's task to determine if the hashes are acceptable. Acceptable hashes  
1244        for software modules are called Reference Integrity Metrics (RIMs) in TCG documents<sup>31</sup>, and:

- 1245     • May be obtained directly from the OEM.  
1246     • Could be learned from 'known-good' systems.  
1247     • Could be retrieved from the device under attestation, assuming the reference  
1248        measurements are signed by the entity that's authorized to provide software for the  
1249        product.

1250

1251        Extending Measured Boot into the OS requires additional OS kernel functionality, allowing  
1252        for recording of measurements before software execution or reading<sup>32</sup> of files identified as  
1253        security-critical. Each file identified for analysis is measured and recorded in the TPM, plus  
1254        a measurement log. Unlike earlier stages of the boot process though, the OS measurements  
1255        may continue as long as the system runs, resulting in incremental changes to logs and PCRs  
1256        as different objects are accessed through the system's life.

1257

### 1258        **5.7.1 Linux Integrity Measurement Architecture (IMA)**

1259        In Linux, a part of the security sub-system, called the Integrity Measurement Architecture  
1260        (IMA), carries out the measurement process. IMA extends the principle of Measured Boot into  
1261        the operating system by providing means to measure selected applications started by the OS.  
1262        System designers and administrators can benefit from the following:

---

<sup>30</sup> Devices in which the AK and DevID are not tightly bound are subject to the "Asokan Attack" (RFC 6813) where an infected device tricks a clean device into attesting on its behalf. Mechanisms such as TPM\_Certify and SKAE can be used to provide proof that the AK and DevID are resident in the same TPM

<sup>31</sup> For TPM1.2, see *TCG Infrastructure Working Group Reference Manifest (RM) Schema Specification*, Version 2.0, <https://www.trustedcomputinggroup.org/tcg-reference-manifest-rm-schema-specification/>.

<sup>32</sup> Examples of simple file reads that could have integrity implications might be configuration files or interpreted-language scripts such as shell or Python. These aren't strictly executable, but they may be able to trigger unauthorized execution.

- 1263     • Measurement of any software executable on the system as it is loaded into memory for  
1264       execution, including shared libraries.  
1265     • Measurement of all files read by a particular user, including programs running with  
1266       permissions of that user.  
1267     • A system designer may benefit from that by measuring all files of specific programs,  
1268       such as database systems or web servers, typically running under a dedicated user  
1269       account.

1270

1271   The system designer must decide on the OS and application components to be measured.  
1272   Executable objects would typically be measured, but measurement may also include any  
1273   other kind of object such as those listed in Section 5.7 above. IMA maintains a measurement  
1274   log, referred to as Integrity Measurement Log (IML) in TCG documents<sup>33</sup>.

1275   IMA has many configuration options:

- 1276     • IMA log formats are configurable using *IMA templates*, allowing a system designer to  
1277       adapt the log format to his/her needs.  
1278     • IMA is configurable using a policy, which to a certain extent allows for definition of  
1279       files to be measured. IMA uses that policy to decide whether a file is to be measured  
1280       or not according on each file Open(). The policy can be hard-coded and compiled into  
1281       the Linux kernel or it can be set dynamically on every boot by writing it to a file in the  
1282       Linux security's filesystem (*/sys/kernel/security/ima/policy*). It is advisable to always  
1283       set the policy as early as possible in the boot process, although Linux kernels of version  
1284       4.7 and later allow the policy to be updated after initial boot.  
1285     • IMA also supports policies that are based on Linux Security Modules (LSM) rules. LSM  
1286       provides an interface for security technologies to extend the Linux kernel's capabilities  
1287       with more security-related features. It is used by technologies such as SELinux,  
1288       Simplified Mandatory Access Control Kernel (SMACK), TOMOYO, AppArmor, Yama,  
1289       and IMA, as well as others<sup>34</sup>. IMA is usually used to measure immutable files, so to  
1290       exclude all log files—which are, by definition, mutable—the SELinux attribute “logfile”  
1291       can be used to tag these, allowing the definition of a policy rule to exclude files that  
1292       are tagged with that attribute. Of course, an administrator must create a policy rule to  
1293       set that attribute on all log files. The same applies to scripts and other interpreted  
1294       executables. Using LSM rules an administrator can improve the measurement process  
1295       and define more fine-grained IMA policies by including and excluding files depending  
1296       on their expected usage.

1297   However, system designers need to be aware of some limitations to IMA, and some impact on  
1298   the overall system, including the following:

- 1299     • Files that are accessed by IMA for measuring may cause measurement violations due  
1300       to concurrent exclusive file accesses, called the Time of Measure Time of Use (ToMTToU)

---

<sup>33</sup> also known as Stored Measurement Log (SML) in some documents

<sup>34</sup> See [https://wiki.gentoo.org/wiki/Extended\\_Verification\\_Module](https://wiki.gentoo.org/wiki/Extended_Verification_Module).

1301 problem<sup>35</sup>. Whenever IMA cannot gain read access to a file a measurement violation is  
1302 indicated by a special log entry, containing a file hash of all zeros, but, in turn, the  
1303 PCR gets extended with a hash of all ones. System designers may wish to flag that  
1304 situation when analyzing log files.

- 1305 • Binding the platform state to an IMA-maintained PCR might not make a lot of sense in  
1306 most cases. As program execution order, even during boot, is non-deterministic on  
1307 most modern operating systems, the sequence of measurements in the log is very likely  
1308 not to be the same across boot cycles. This will result in unpredictable PCR values,  
1309 even if the set of executed programs has not changed.
- 1310 • Verification of IMA measurements would require an enhanced verification to a  
1311 management station. Instead of just comparing a single PCR value to a known-good  
1312 one, the management station needs to calculate the accumulated hash over the  
1313 Integrity Measurement Log and compare it to the reported hash from the TPM Quote.  
1314 Furthermore, the management station needs to compare all entries of the  
1315 measurement log to previously defined reference measurements, RIMs. That is, it must  
1316 hold hundreds or even tens of thousands of reference values for a particular system  
1317 configuration. Typically, a search-optimized (indexed) database is used.
- 1318 • As the IMA measurement process continues operation, even during the retrieval of a  
1319 “TPM Quote”, there might be some mismatch between the last recorded entry in the  
1320 PCR and the last entry in the log, when retrieving it. The log might have advanced  
1321 while the “quote” operation was running. Accordingly, the appropriate log entry—the  
1322 last one that was incorporated in the PCR—must be identified to allow the Challenger  
1323 to verify the log. The process for identifying that log entry can look like the following:
  - 1324 1. Retrieve the log for the first time, L1 (with the number of entries  $|L1|$ ).
  - 1325 2. Generate the TPM Quote.
  - 1326 3. Retrieve the log again, L2 (with the number of entries  $|L2|$ ).
  - 1327 4. Check whether the log has advanced ( $|L2| > |L1|$ <sup>36</sup>). If so, then the log entry  
1328 that has been used for the TPM Quote must be identified, which can either  
1329 happen on the device under attestation itself or on the management station,  
1330 which, in turn, needs either both of the log files or just the second one (L2)  
1331 and additionally the number of entries in L1 ( $|L1|$ ). The process would look  
1332 something like what follows:<sup>37</sup>
    - 1333 a. Calculate the accumulated hash from the whole first log (L1).

---

<sup>35</sup> If a file that is opened exclusively for use/writing (at the time of use), the measurement may fail, as the measurement process cannot also open the file for reading in order to measure it (at the time of measurement).

<sup>36</sup>  $|Lx|$  indicates the number of entries in the log.

<sup>37</sup> While the order of entries in the log may be non-deterministic, we do assume that the log file contains entries in *exactly* the same order as they were measured into the TPM.

- 1334                   b. Compare the accumulated hash against the hash from the TPM Quote.  
1335                   If it matches, then return the full log up to the current entry.  
1336                   Otherwise, continue.
- 1337                   c. Update the accumulated hash by incorporating the next log entry from  
1338                   list L2. Then continue with step (b).
- 1339           • In an OS with multiple concurrent processes, the TSS Access Broker should be used  
1340           to manage concurrent access from multiple processes to a single TPM.<sup>38</sup>
- 1341           • While the file hashes are collected, none of the common file metadata is logged by IMA,  
1342           so there's no record of user ID or group ID, either of the file itself or the process  
1343           accessing it, or file attributes.
- 1344           • IMA cannot distinguish between text files, log files, executable scripts, and other files.  
1345           IMA can just detect access to files other than executables. Affected are all kinds of  
1346           interpreted languages, such as Java, Python, Ruby, Lua, shell scripts, etc. IMA's policy  
1347           language is insufficient for that purpose. Depending on the number of files involved,  
1348           IMA's facilities for measuring every file can be used. The Challenger then is responsible  
1349           to identify relevant files, such as scripts. As IMA keeps an internal hash table to avoid  
1350           adding duplicate entries to the log, measuring a lot of files may cause that hash table  
1351           to consume a notable amount of kernel memory and may cause out-of-memory errors.  
1352           LSM rules in an IMA policy, as mentioned above, can be used to restrict measurement  
1353           only to files that have security significance.

1354           IMA is subject to continuous improvement, and may offer additional desirable features  
1355           after this document has been published.

### 1356        **5.7.1.1      IMA Appraisal**

1357           IMA Appraisal is the equivalent of Secure Boot extending into the operating system. In  
1358           contrast to IMA measurements, IMA Appraisal requires a file to match a known-good  
1359           measurement in order to allow access.

1360           Known-good measurements can be managed several ways.

1361           One approach is to simply learn the hash of each file from a clean system. For that purpose,  
1362           a particular boot argument to the Linux kernel must be passed (*ima\_appraise=fix*), instructing  
1363           IMA to record the known-good measurements of accessed files. On any subsequent boot, IMA  
1364           appraisal can be activated to enforce access of a file only if it matches the known-good  
1365           measurement, by passing the corresponding kernel boot argument (*ima\_appraise=enforce*).

1366           By default, IMA allows access to files if the measurements match and will update the known-  
1367           good measurement if the file is modified after having passed the check on the initial file-  
1368           open.<sup>39</sup>

1369

1370           IMA also supports *immutable* files utilizing digital signatures on the known-good  
1371           measurements, generated and signed when the software components are built, prior to

---

<sup>38</sup> For example, see <http://www.trustedcomputinggroup.org/wp-content/uploads/TSS-TAB-and-Resource-Manager-00-91-PublicReview.pdf>

<sup>39</sup> [https://wiki.gentoo.org/wiki/Integrity\\_Measurement\\_Architecture#Registering\\_the\\_file\\_hashes\\_for\\_the\\_system](https://wiki.gentoo.org/wiki/Integrity_Measurement_Architecture#Registering_the_file_hashes_for_the_system)

1372 installation on a particular machine. Private keys stay with the build system; only the signer's  
1373 public key(s) must be available in the IMA keyring (currently *.ima\_root\_ca*), although they  
1374 must be loaded into the keyring early in the boot process, typically in an *initramfs*.  
1375 Digital signatures provide protection against online tampering with the files. Offline  
1376 tampering would still be possible by generating a malicious key-pair, re-signing all the files  
1377 with the malicious private key, and then replacing the public key with the malicious one. It's  
1378 possible to protect against offline changes to extended file attributes in general, including the  
1379 IMA Appraisal known-good measurements, with the Linux kernel Extended Verification  
1380 Module (EVM)<sup>40</sup>. With EVM, IMA can use a cryptographic hash (HMAC) or digital signature,  
1381 with a key loaded at boot time, to verify the extended file attributes on any access. EVM keys  
1382 can be protected by the TPM using Sealing, or in applications where there's an interactive  
1383 user, a password can be used to unlock the EVM key.  
1384 A system designer may utilize IMA Appraisal whenever a system must run only approved  
1385 software that is known prior to device deployment. While systems which use signatures  
1386 generated by the device manufacturer's software build system can readily be updated in the  
1387 field (assuming keys don't change), the software update mechanism on systems that require  
1388 re-measurements of software on the system itself is out of scope for this document and is left  
1389 to be defined by the system designer.

1390

### 1391 **5.7.2 PCR Definitions**

1392 The allocation of various functions to specific PCRs depends a lot on the underlying software  
1393 architecture of the Network Equipment.

1394 There are two definitions available at the time of this writing to serve as starting points for  
1395 Network Equipment:

- UEFI-based *TCG\_PC Client Implementation for BIOS.pdf* Section 3.2.3
- Mobile Platforms: *TPM2.0 Mobile Common Profile* [14] Section 2.5

1398

1399 Guidelines for information to be collected in PCRs can also be found in the following:

- *NIST SP800-155 BIOS Integrity Measurement Guidelines*;  
[http://csrc.nist.gov/publications/drafts/800-155/draft-SP800-155\\_Dec2011.pdf](http://csrc.nist.gov/publications/drafts/800-155/draft-SP800-155_Dec2011.pdf)

1402

1403

### 1404 **5.8 Composite Networking Devices**

1405 In modular Network Equipment, there's usually a control plane that manages field  
1406 replaceable units (FRUs) like line cards, each of which may contain complex control  
1407 processors with TPMs themselves. FRUs, including control plane processors, can usually be  
1408 "hot swapped", that is, inserted and removed while the rest of the system continues to run,

---

<sup>40</sup> See [https://wiki.gentoo.org/wiki/Extended\\_Verification\\_Module](https://wiki.gentoo.org/wiki/Extended_Verification_Module)

1409 yielding a system where components may come and frequently, even though the system as a  
1410 whole rarely goes down or starts up all at once. This results in several challenges:

- 1411 • Each time an FRU is inserted or removed, a control plane process typically notes the  
1412 change in configuration, initializes the new unit and may alert other components in  
1413 the system of the new configuration. These new components should not be admitted  
1414 to the system unless they can show that they are in a known-good state.
- 1415 • The conventional definition of attestation may need to be extended with the goal of  
1416 reporting not just the attested state of one control plane processor, but including also  
1417 the state of the subsidiary FRUs, each with their own processor state.

1418

1419 TCG offers a number of mechanisms that are useful to the design of such systems:

- 1420 • **Integrity-Preserved logs** could be used to keep a tamper-evident record of modules  
1421 that were added or subtracted during the life of the system.
- 1422 • **Identity** - The system designer must set the rules and procedures for determining  
1423 membership in the composite device, but Device Identity (e.g., IDevID or LDevID)  
1424 could provide a solid basis for determining if specific devices are eligible to join a  
1425 composite or not.
- 1426 • **Internal Attestation** among elements – System designers may want elements of a  
1427 system to prove not only their identity, but also that they're running authorized  
1428 software.<sup>41</sup>
- 1429     ○ In some cases, control processes may already know what versions of software  
1430       they'll accept on subsidiary units (e.g. a router control plane may only work  
1431       with line cards with certain software revisions, and if the software isn't up to  
1432       rev, the control plane may force an upgrade.)
- 1433     ○ In other cases, units may need to provide evidence that software versions  
1434       being run are authentic (e.g. signed SWID tags). This might be applicable to  
1435       peer nodes in a distributed system, where there's no obvious leader to dictate  
1436       expected software versions.
- 1437     ○ Except for the most aggressively distributed systems, these mechanisms  
1438       would function inside the composite "device", and would be designed to suit  
1439       the manufacturer's software environment, without the need for interoperability  
1440       specifications. As a result, extensions to standards are not likely needed for  
1441       this application.
- 1442 • **External Attestation** – Composite systems will usually elect a leader, either through  
1443       hardware or software means, to communicate with the outside world and represent  
1444       that state of the system (i.e. provide a management plane and management  
1445       interface). One of the functions that might be desired is Remote Attestation, where  
1446       an external management device can identify the system, its components and the  
1447       authenticity of software running on the system, including all the software running on  
1448       subsidiary units which may not be visible to an external management agent.

1449

1450

---

<sup>41</sup> For example, a Device shouldn't use a crypto-accelerator FRU unless it can prove that it's not been tampered with.

## 1451 5.9 Integrity-Protected Logs

1452 Use cases like Security Information and Event Management (SIEM), IETF *Security*  
1453 *Automation and Continuous Monitoring* (SACM) or even Lawful Interception rely on log files  
1454 created by devices deployed in the field outside the direct, physical administration of the  
1455 operator or owner. Supporting these use cases requires integrity-protected or non-  
1456 repudiable log files created by the device. Protecting the integrity requires evidence of the  
1457 authenticity, order, and time of the log file entries.

1458 An integrity-protected log can be utilized in cases like Lawful Interception proving the correct  
1459 operation of a particular device or link at the time of data interception, for later forensic or  
1460 court-ordered evaluation.

1461

### 1462 5.9.1 Functional Requirements

1463 Integrity-Protected logged data must be protected against undetected manipulation such as  
1464 deletion, insertion of modification of log events, and can additionally provide evidence of the  
1465 state of the system at the time of the logged event.

1466 Each entry in the log file provides the following information:

- 1467 • Type of event
- 1468 • Date and time
- 1469 • Order relative to other events

1470 The TPM-issued signature covers the following information, showing it has not been modified  
1471 after being signed:

- 1472 • Identity of the log device, based on the DevID
- 1473 • A non-volatile monotonic sequence number showing the order of events and proving  
1474 that the current event immediately follows the previous event, so that causal  
1475 relationships can be determined, and deletion of events can be detected
- 1476 • (optional) time of signature
- 1477 • (optional) state of device during signature operation

1478 Network Devices may reboot unexpectedly, or they may operate without a reboot for years; as  
1479 such, logging mechanisms must typically meet two additional goals.

- 1480 • Log files must persist through a reboot, and retain integrity through the reboot.
- 1481 • But at the same time, devices may produce a lot of logging information, so  
1482 mechanisms that “trim” the logs, discarding old entries, are often provided.

1483 The need to discard old log information implies that there can be no expectation of a  
1484 permanent log that goes all the way back to the first use of the Device; it’s the Administrator’s  
1485 responsibility to collect and archive log information periodically if a long-term record is  
1486 needed. The intent of the mechanism described here is to allow the Administrator to  
1487 demonstrate that a set of log files provides a complete record.

1488

1489 **5.9.2 Integrity-Protected Log Implementation**

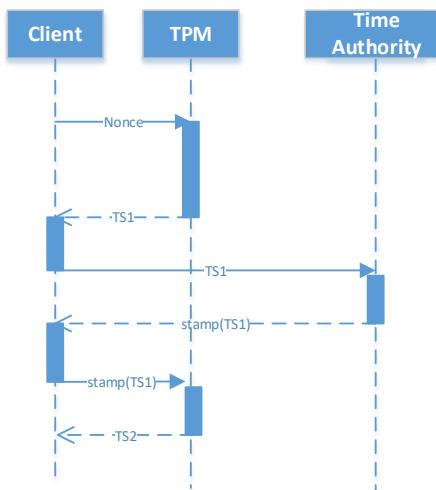
1490 Implementation of Integrity-Protected Logs uses a number of TPM mechanisms:

- 1491     • The TPM time-keeping mechanism can be used to generate tamper-resistant time  
 1492        stamps on log entries.
- 1493     • The Monotonic Counter mechanism can be used to sequence log entries and to detect  
 1494        missing or added entries, even across boot cycles.
- 1495     • The TPM signing mechanism can be used to sign log entries so that any changes in  
 1496        an individual entry can be detected.

1497

1498 **5.9.2.1 Timekeeping in TPM 1.2**1499 While the TPM does not keep a real-time clock, it does provide a mechanism that keeps track  
 1500 of elapsed time since the last initialization. This relative time measure is called a Tickstamp  
 1501 (*TPM 1.2 Part 1 Design Principles* Chapter 20).1502 The tick counter increments at a specified rate, but cannot be “set”, and does not relate to  
 1503 any external real time, so there is usually a requirement to establish the correlation between  
 1504 the monotonic tick-counter of the TPM and a trustworthy external real time clock.1505 The correlation process works by measuring the difference between an authorized external  
 1506 time source and the local Tickstamp timer. A Tickstamp is taken prior to the request to an  
 1507 external source, and then again after the response has been processed, yielding a bounded  
 1508 uncertainty in time correlation. The internal tick counter of the TPM is subject to drift, so  
 1509 correlation may have to be repeated periodically to maintain adequate sync. (See *TPM 1.2*  
 1510 *Part 1 Design Principles* Section 20.3 for a detailed description.)

1511 Figure 5 shows the process:



1512

1513 **Figure 5 : Time Correlation**

1514

1515 The TPM supports the generation of signed tick stamps that include a signature over user-  
 1516 supplied data associated with the current tick count value, the tick increment rate and tick

1517 session nonce. Every time the TPM is reset a new tick session nonce is generated. (See *TPM*  
1518 *1.2 Part 1 Design Principles* Section 20.1 for a detailed description.)

1519

1520 **5.9.2.2 Timekeeping in TPM 2.0**

1521 TPM 2.0 offers additional timekeeping facilities, described in *TPM 2.0 Part 1 Architecture*  
1522 Chapter 36.7. *Time* contains the time in milliseconds since the last startup of the TPM.  
1523 *resetCount* is a counter that increments every time the TPM is successfully reset. Its purpose  
1524 is, among other things, to indicate possible discontinuity in *Clock*. *restartCount* is used to  
1525 provide an indication of discontinuities due to (1) TPM Resume, (2) TPM Restart, or (3)  
1526 *\_TPM\_Hash\_Start*; all allowing TPM *Time* to fall behind real time (*TPM 2.0 Part 1 Architecture*  
1527 Chapter 36.5).

1528 *Clock* is a volatile value that increments in memory each millisecond. The non-volatile value  
1529 *NV Clock* is updated periodically from that value. On every reboot the value of *Clock* is set to  
1530 the value of *NV Clock*. However, in case of unexpected power loss the value of *Clock* may be  
1531 reported twice as the final *Clock* value might not have been written to *NV Clock*. The TPM flag  
1532 *safe* of the *TPMS\_CLOCK\_INFO* structure is used to indicate such a situation of unexpected  
1533 power loss. Mitigations are described in *TPM 2.0 Part 1 Architecture* Chapter 36.2. In contrast,  
1534 an orderly shutdown of the TPM is possible using the *TPM2\_Shutdown()* command, which  
1535 tells the TPM to preserve an appropriate state.

1536 Neither the TPM 1.2 nor the TPM 2.0 has a real-time clock that is able to advance time in the  
1537 power-off state. The TPM1.2 clock retains no state through a power cycle, but in TPM 2.0,  
1538 *Clock* is non-volatile, and can be set forward (never backward, though, except by installing a  
1539 new owner) by external software managing the TPM, allowing for correlation to an external  
1540 real-time clock. As the TPM may be driven by an imprecise frequency source, the clock is  
1541 subject to drifting. To compensate, the TPM 2.0 allows external software to adjust the rate of  
1542 advancement by +/- 15 %.

1543 Attestation data in TPM 2.0 includes the *resetCount* value, although it is obfuscated to  
1544 preserve privacy (*TPM 2.0 Part 1 Architecture* Chapter 36.7). Nevertheless, an attesting  
1545 management station can detect a reboot, and hence a possible discontinuity in time, by  
1546 checking whether the *resetCount* value has changed. TPM 2.0 also allows for usage of timing  
1547 data in policies, e.g. to limit the usage of keys for a certain amount of time.

1548

1549 **5.9.2.3 Assembling the Log**

1550 While this document does not specify the exact format of Integrity-Protected log files, the logs  
1551 can be assembled using TPM Monotonic Counters for sequencing, TPM signing for integrity,  
1552 and Tickstamps for temporality.

1553 To detect addition or deletion of events in the log, TPM Monotonic Counters (*TPM1.2 Design*  
1554 *Principles* Chapter 17, *TPM Rev 2.0 Part 1 – Architecture* Section 37.2.4) can be used. These  
1555 counters can be incremented by a TPM user, but can never be decremented or cleared. This  
1556 mechanism allows log entries to be created where each entry has a sequence identifier that's  
1557 exactly one larger than the previous entry, allowing deletions or additions to be easily  
1558 detected. Old log entries may be deleted, but in the remaining logs, the monotonic count

1559 should increment uniformly from the oldest entry through each intermediate entry to the  
1560 newest.

1561 For high-value, low-rate<sup>42</sup> log information, a simple mechanism can be used:

- 1562 • Each new log entry can be assembled with a Tickstamp or Timestamp to show when  
1563 it happened, and the next value from a non-volatile Monotonic Counter.
- 1564 • The log entry can be signed with an IDevID or LDevID key, and appended to the log  
1565 file.
- 1566 • Reboots of the system and restarts of the TPM can be detected using the tick session  
1567 nonce of a Tickstamp (TPM 1.2 and TPM 2.0) or the values *resetCount* and  
1568 *restartCount* (TPM 2.0 only).

1569 Any subsequent tampering will invalidate signatures, or show a gap in the sequence numbers.

1570 It should be noted that the TPM also incorporates a mechanism which can be used to audit  
1571 the command codes which it has been called upon to execute; this mechanism works by  
1572 keeping an internal digest of commands, which can be compared to an external log file of  
1573 what commands were executed.

1574

1575 High-rate logs would require a more complex block-signing mechanism, something beyond  
1576 the scope of this document.

1577

## 1578 **5.10 Entropy Generation**

1579 Many network protocols rely on a dependable source of random numbers for correct  
1580 operation. For security related protocols, random numbers may be used by a cryptographic  
1581 element. For other protocols, this might be a TCP sequence identifier, or a source port  
1582 number, or some other nonce used to identify a particular session or traffic flow.

1583 The TCG TPM definition includes a requirement for a Random Number Generator, which can  
1584 be accessed after the TPM self-test sequence is complete, simply by invoking the  
1585 TPM\_GetRandom() or TPM2\_GetRandom() ordinals.

1586 TCG specifications allow TPM vendors considerable flexibility in implementation choices. A  
1587 common design approach, not mandated by TCG, would be a mechanism where a physical  
1588 noise source of some sort (a True Random Number Generator, or TRNG) can be used to seed  
1589 a Deterministic Random Bit Generator (DRBG), ensuing predictable statistical properties.

1590 For many Network Equipment vendors, compliance to standards such as NIST Special  
1591 Publication 800-90A/B/C [14] or BSI AIS-31, *A Proposal for Functionality Classes for Random*  
1592 *Number Generators* [15], is a requirement; in this case, the OEM designer should consult the  
1593 TPM vendor for compliance certifications.

1594 There are two aspects of the standards that OEM designers should consider:

- 1595 a) The standards require continuous testing of the components of the RNG, in addition  
1596 to a self-test at startup. These tests must be done inside the TPM, as they require

---

<sup>42</sup> Average rate of a few log entries per hour; TPMs can't sign quickly, and the Monotonic Counter in TPM1.2 will wrap and/or wear out if it's incremented more than once every few minutes.

1597 access to internal state. If there's a failure of the self-test at any point, the TPM will  
1598 report a failure in response to further Requests.

- 1599 b) The standards require that entropy from a physical source (TRNG) should be added  
1600 periodically to the DBRG, so that repeated reads will not exhaust the entropy supply.

1601 A compliant TPM is required to monitor and replenish the entropy level as needed, and  
1602 may return fewer bytes of random numbers than requested should requests outpace  
1603 generation of new entropy.

1604 Rates at which new entropy is generated within the TPM is not specified by TCG or the  
1605 standards, so it's up to the OEM designer and the TPM vendor to ensure that the TPM  
1606 in consideration generates entropy at an adequate rate for the application.

1607

1608 The TPM has a mechanism with which external entropy can be added to the TPM's DRBG  
1609 output (TPM\_StirRandom() and TPM2\_StirRandom() ). As long as the TPM vendor certifies  
1610 that the device generates entropy on its own at an adequate rate, the OEM designer is not  
1611 likely to need this ordinal.<sup>43</sup>

1612 It's always good practice for the OEM's operating system to collect and merge entropy from  
1613 as many sources as are available. For example, the Linux OS rng-tools can be configured to  
1614 merge entropy from a number of sources such as arrival time of network traffic, or disk  
1615 seek times, with the entropy obtained from the TPM to supply /dev/random for OS and  
1616 application needs.

1617

## 1618 **5.11 Deprovisioning**

1619 Networking equipment often contains sensitive information that an Administrator would  
1620 want to protect from disclosure, and must be destroyed when the device is deprovisioned.

1622 The Device Manufacturer should provide a mechanism to enable deprovisioning.

1623 Upon deprovisioning, the device:

- 1624 • Should delete any LDevID, LAK or other customer-generated keys from the TPM.
- 1625 • Should *not* erase IDevID or IAK keys, since these indicate the manufacturer's device  
1626 identity, not the Administrator's.

1627 Any installation-specific configuration files should also be deleted.

1628 Deletion of files from modern flash-based memory systems is not easy; an approach to  
1629 ensure the deletion is reliable would be:

- 1630 • Create a local storage key as a child of the SRK.
- 1631 • Store all configuration files in a partition using disk-encryption technology.

1632 Upon deprovisioning, the TPM can be instructed to delete the local storage key, rendering  
1633 the partition unintelligible.

1634

---

<sup>43</sup> Assuming a platform has additional sources of entropy, StirRandom can also be used to defend against a hidden failure of the TPM's internal entropy source.

## 1635 6. References

- 1636
- 1637 [1] IEEE Standard for Local and Metropolitan Area Networks  
1638 Borza, Mike (ed.) and Max Pritikin (ed.), "802.1AR-2009 – IEEE Standard for Local  
1639 and Metropolitan Area Networks - Secure Device Identity", IEEE Computer Society,  
1640 New York, New York, December 10, 2009,  
[1641 http://www.ieee802.org/1/pages/802.1ar.html](http://www.ieee802.org/1/pages/802.1ar.html)
- 1642 [2] Trusted Computing Group, *Trusted Platform Module Library, Part 1: Architecture,*  
1643 Family "2.0", Level 00 Revision 01.16
- 1644 [3] Trusted Computing Group, *Trusted Platform Module Library, Part 2: Structures,*  
1645 Family "2.0", Level 00 Revision 01.16
- 1646 [4] Trusted Computing Group, *TPM Keys for Platform Identity for TPM 1.2*, Version 1.0,  
1647 Revision 2, 12 May 2015
- 1648 [5] Trusted Computing Group, *TPM Keys for Platform DevID for TPM 2*, Version ??,  
1649 Revision ??, [Not published as of July, 2017]  
1650 [http://members.trustedcomputinggroup.org/apps/org/workgroup/infra\_wg/downlo  
1651 ad.php/31089/TPM%20Keys%20for%20Platform%20DevID%20for%20TPM%202.0%  
1652 20--%20Current%20Draft]
- 1653 [6] The International Organization for Standardization/International Electrotechnical  
1654 Commission (ISO/IEC), *Information Technology - Software Asset Management - Part 2:*  
1655 *Software Identification Tag*, ISO/IEC 19770-2, November 2009  
1656 <https://www.iso.org/obp/ui/#iso:std:iso-iec:19770:-2:ed-1:v1:en>
- 1657 [7] Registry of Reserved TPM 2.0 Handles and Localities  
1658 Trusted Computing Group, "Registry of Reserved TPM 2.0 Handles and Localities,"  
1659 Trusted Computing Group, Beaverton, OR, October 11, 2013.
- 1660 [8] Trusted Computing Group, TCG TPM v2.0 Provisioning Guidance, Version 1.0,  
1661 Revision 1.0, Mar 15, 2017
- 1662 [9] EK Credential Profile 2.0  
1663 Trusted Computing Group, *TCG EK Credential Profile For TPM Family 2.0, revision 14*,  
1664 Trusted Computing Group, Beaverton, OR, 2014.
- 1665 [10] Trusted Computing Group, *TCG Infrastructure Workgroup Subject Key Attestation*  
1666 *Evidence Extension*, Specification Version 1.0, Revision 7, 16 June 2005
- 1667 [11] Trusted Computing Group, *TCG Attestation PTS Protocol: Binding to TNC IF-M*,  
1668 Specification Version 1.0, Revision 28, August 24, 2011
- 1669 [12] Trusted Computing Group, *TCG Credential Profiles for TPM Family 2.0; Level 2*,  
1670 [https://members.trustedcomputinggroup.org/apps/org/workgroup/infra\\_wg/downl](https://members.trustedcomputinggroup.org/apps/org/workgroup/infra_wg/downl)  
1671 oad.php/30794/Credential\_Profile\_Platform\_Att\_V1.0-000-17-03-28-Draft.docx [not  
1672 published yet]
- 1673 [13] *A Practical Guide to TPM 2.0*, W. Arthur & D. Challener, Apress; 1st Edition, January  
1674 28, 2015
- 1675 [14] National Institute of Standards and Technology, *Recommendation for Random*  
1676 *Number Generation Using Deterministic Random Bit Generators*, Revision 1, June  
1677 2015, <http://dx.doi.org/10.6028/NIST.SP.800-90Ar1>

- 1678 [15] National Institute of Standards and Technology, *Report on Post-Quantum*  
1679 *Cryptography, NISTIR 8105*  
1680 <http://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8105.pdf>
- 1681 [16] Bundesamt für Sicherheit in der Informationstechnik (BSI) AIS-31, *A Proposal for*  
1682 *Functionality Classes for Random Number Generators*, Version 2.0, 18 September  
1683 2011,  
1684 [https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Zertifizierung/Interpretationen/AIS\\_31\\_Functionality\\_classes\\_for\\_random\\_number\\_generators\\_e.pdf](https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Zertifizierung/Interpretationen/AIS_31_Functionality_classes_for_random_number_generators_e.pdf)
- 1685 [17] ISO/IEC 18031:2011 Information technology -- *Security techniques -- Random bit*  
1686 *generation*, published on: 2011-11,  
1688 [http://www.iso.org/iso/catalogue\\_detail.htm?csnumber=54945](http://www.iso.org/iso/catalogue_detail.htm?csnumber=54945)
- 1689 [18] Trusted Computing Group, *TCG Storage Opal Integration Guidelines*, Version 1.00,  
1690 Revision 1.00, March 16, 2016  
1691 [https://trustedcomputinggroup.org/wp-content/uploads/TCG\\_Storage\\_ReferenceDocument\\_Opal\\_Integration\\_Guidelines\\_v1.00\\_r1.00.pdf](https://trustedcomputinggroup.org/wp-content/uploads/TCG_Storage_ReferenceDocument_Opal_Integration_Guidelines_v1.00_r1.00.pdf)
- 1694 [19] Linux *Integrity Measurement Architecture*, <http://linux-ima.sourceforge.net>,  
1695 <https://wiki.gentoo.org/wiki/Project:Integrity>
- 1696 [20] Trusted Computing Group, *IWG Reference Architecture for Interoperability (Part 1)*,  
1697 Specification Version 1.0,  
1698 [http://www.trustedcomputinggroup.org/resources/infrastructure\\_work\\_group\\_reference\\_architecture\\_for\\_interoperability\\_specification\\_part\\_1\\_version\\_10](http://www.trustedcomputinggroup.org/resources/infrastructure_work_group_reference_architecture_for_interoperability_specification_part_1_version_10), June 2005.
- 1700 [21] TCG Specification, *TPM 2.0 Mobile Common Profile, Family “2.0” Level 00 Revision 31*,  
1701 21 December, 2015  
1702
- 1703

## 1704 7. Glossary

1705

1706 There are a few terms that are specific to Network Equipment:

1707 • Device – a unique piece of network equipment

1708 • CPE – Customer Premise Equipment

1709 • FRU – Field Replaceable Unit

1710 • From IEEE 802.1AR:

1711     ○ **Initial Secure Device Identifier (IDevID):** The Secure Device Identifier  
1712       installed on the device by the manufacturer1713     ○ **Locally Significant Secure Device Identifiers (LDevIDs):** A Secure Device  
1714       Identifier credential that is unique in the local administrative domain in which  
1715       the device is used1716     ○ **Secure Device Identifier (DevID):** A device identifier that is cryptographically  
1717       bound to the device and is composed of the Secure Device Identifier Secret and  
1718       the Secure Device Identifier Credential

1719 TCG-specific acronyms and abbreviations can be found in the TCG Glossary

1720     [https://www.trustedcomputinggroup.org/wp-content/uploads/TCG\\_Glossary\\_Board-  
1721       Approved\\_12.13.2012.pdf](https://www.trustedcomputinggroup.org/wp-content/uploads/TCG_Glossary_Board-Approved_12.13.2012.pdf)

## 1722 8. Appendices

### 1723 8.1 Use-Cases For Further Study

1724

#### 1725 8.1.1 Secure Identity in Virtual Machines

1726 Manufacturers are starting to incorporate Virtual Machine technology into Network  
1727 Equipment, either as an element of an integrated hardware/software product or as a  
1728 software-only product running entirely within a virtual machine environment, as in Network  
1729 Function Virtualization.

1730 In either case, there's often a need to identify the hardware resource underlying the virtual  
1731 machine, to ensure that the function is actually executing in the expected environment, and  
1732 to determine if the software is running in a virtual machine or directly on a physical machine.

1733

#### 1734 8.1.2 Distributed Composite Devices

1735 In today's infrastructure, composite devices make use of wired or wireless networks as  
1736 interconnects to merge management planes or sub-sets of control plane functions. Therefore,  
1737 the physical components of a composite device can be geographically distributed, but act as  
1738 a single entity on the management plane. In addition, network interfaces can be associated  
1739 with more than one composite device; each managed via independent management planes  
1740 interfaces.

1741 There is a need to take into account the type of composition used to merge management  
1742 planes or control plane functions of distinguishable physical units. Every distributed  
1743 composition relies on a specific form of interconnect between physically separated units,  
1744 which can be used as a basis for corresponding solutions.

1745

### 1746 8.2 Standardization Work For Further Study

#### 1747 8.2.1 Known Good Hashes

1748 There are currently no existing protocols that are being used to convey known-good hashes  
1749 from the Network Device Manufacturer to Attestation systems that might use these values.

1750 One suggested approach could be for the manufacturer to provide an "oracle", which could  
1751 return a signed quote and a log with an entry for every possible object that could be measured,  
1752 using the PTS protocol; the Attesting system then would use this list as a data base to look  
1753 up expected values.

#### 1754 8.2.2 Composite Systems

1755 External Attestation of a composite system needs new protocol work. The current attestation  
1756 mechanism assumes attestation is being carried out on an integrated unit with one TPM. In  
1757 a composite system, attestation might require multiple steps:

- 1758     • Attest the control plane element that was elected (either by hardware or software) to  
1759       represent the device to the outside world. This can be done using conventional PTS.  
1760     • Ask the control plane element for a list of devices that it currently has accepted into  
1761       its configuration. This might be a list of line cards, or a list of stacked switches, or a  
1762       bunch of servers. The list might be flat or hierarchical, as decided by the system  
1763       designer. All that the management station needs is a list of the elements and how to  
1764       identify them.  
1765     • The management system might then ask for attestation results from some or all of  
1766       the units reported. Some units may include a separate TPM, which can be leveraged  
1767       to enrich the attestation of the composite device or even allow for direct remote  
1768       attestation of a sub-component.

1769  
1770     In most cases, the elements that make up the system (e.g. line cards) will not be directly  
1771       visible externally, i.e., they won't have routable IP addresses to their own control plane  
1772       processors, resulting in a need for a mechanism to send a request for Attestation that can  
1773       be forwarded by the receiving management plane processor to the appropriate internal FRU.  
1774     For PTS, this would probably be done by adding an "element id" field to the request/response.  
1775     For an SNMP-based mechanism, it might be handled by introducing a layer of hierarchy into  
1776       the MIB.

### 1777 1778     8.2.3 Reference Manifests

1779  
1780     The document *TCG Infrastructure Working Group Reference Manifest (RM) Schema*  
1781       *Specification*, Version 2.0, <https://www.trustedcomputinggroup.org/tcg-reference-manifest-rm-schema-specification/> defines the XML schema with which integrity information is communicated  
1782       between entities for TPM1.2.

1783  
1784     There currently is not an equivalent document for TPM2.0

1786

## 9. Revision History

Name	Description	Date	Revision
S. Hanna & G. Fedorkow	First draft	2015-05-01	v0.1
G. Fedorkow		2015-05-07	v0.2
G. Fedorkow	Integrate TCG template and email outlines	2015-05-12	V0.2, R3
S. Hanna & G. Fedorkow	Elaborate and reorganize use-cases	2015-05-28	V1.0 R4
N. Kuntze	Update on the Use Cases	2015-05-29	V0.2, R5
G. Fedorkow	Merge Nicolai's updates	2015-06-01	V1.0 R6
N. Kuntze	Minor Comments and additional Use Case	2015-06-01	V1.0 R7
G. Fedorkow	Restructure Implementation section to reflect use-cases	2015-06-07	V1.0 R8
G. Fedorkow	Merged comments from Nicolai, Carsten, Steve H and others		
G. Fedorkow	Many updates at Face-to-Face meeting	2015-06-18	V1.0 R9b
G. Fedorkow	New material in Device ID section, proposed text on Privacy, background on 802.1AR and comments from T Laffey and S Hanna	2015-07-10	V1.0 R10b
Max, Henk, Carsten	New material from Max, Henk and Carsten		V1.0 R11
G Fedorkow	Merged other comments and changes		V1.0 R11b
G Fedorkow	Revised RNG Implementation section, updated DevID after Aug 12 phone conf	2015-08-24	V1.0 R12a
G Fedorkow	Outlined approaches to many implementation sections.  Merged in text on license management implementation from Steve Hanna	2015-09-07	V1.0 R12c
G Fedorkow	Revision to License Manager implementation section 6.6 to allow operation on air-gapped networks  Revision to Composite Inventory section 6.9 to respond to comments  Updates responding to T Laffey's comments on v11b	2015-09-014	V1.0 R13a
G Fedorkow	Merged in updates from Nicolai  - Remote Management  - Integrity Logs	2015-10-03	V1.0R13c
G Fedorkow	Merged TPM2 material from Bill Sulzen, and addressed many comments from T Laffey (Thanks Tom!)	2015-10-15	V1.0R14a

G Fedorkow	Restructured Section 6.1.1 to eliminate redundancy	2015-10-18	V1.0R14c
G Fedorkow	Updates after Montreal F2F Added Software Inventory sections	2015-12-8	V1.0R15a
G Fedorkow	Highlighted open issues prior to F2F		V1.0R16
G Fedorkow	Comment resolution following San Francisco F2F Eliminated sections on VM (it's new work, not guidance) Revised Identity section to address many comments Simplified Composite Systems and Integrity Logs based on F2F discussion	2016-03-16	V1.0R17a
G Fedorkow	Cleared change bars Moved some of the 'future' work to Appendix Started to clean the bibliography		V1.0R18a
G Fedorkow	Added "How to Install an LDevID" Added review comments from Bill Sulzen	2016-05-13	V1.0R19b
G Fedorkow	Editorial comments resolved and "accepted"; doc change-tracker shows comments that might need additional review  Added a short sub-sub-sub-section on EST approaches for conveying TPM_Certify; this section may need normative work somewhere else (probably IETF)	2016-05-25	V1.0R19c
H. Birkholz	Added details on current practice for Composite Devices, sub-component composition and a corresponding subsection 8.1.3. to "Use-Cases For Further Study"	2016-06-01	V1.0R19d
G Fedorkow	Added comments from D Challener  Merged updates from Henk on Composite Systems	2016-06-01	V1.0R19e
G Fedorkow	Merged corrections and comments from Graeme Proudler	2016-06-05	V1.0R19f
G Fedorkow	Merged corrections from W Sulzin Started to reduce TPM1.2 bias	2016-06-05	V1.0R19g
G Fedorkow	Merged corrections from S Hanna		V1.0R19h
G Fedorkow	Merged corrections from Nicolai & Michael Eckel  Updates from Vienna F2F	2016-06-19	V1.0R19i
G Fedorkow	Replaced IMA section	2016-10-19	V1.0R21a
G Fedorkow	Editorial updates from Seoul F2F call (hmm, that would be Face-to-Phone...)	2016-10-24	V1.0R21b
G Fedorkow	Change bars removed from R21b	2016-10-24	V1.0R22a
G Fedorkow	Closed more reviewer comments  Moved New Work items to appendix	2016-11-03	V1.0R22b

G Fedorkow	Reviewer comments from Graeme, TC, Huawei and Cisco	2016-12-28	V1.0R23a
G Fedorkow	Reviewer comments from Kent Watsen, JNPR	2017-01-21	V1.0R23b
G Fedorkow	Updates after February 2017 F2F with change bars	2017-03-22	V1.0R23d
G Fedorkow	Clean copy with change-bars removed, ready for the next round of revision	2017-03-22	V1.0R24a
G Fedorkow	Removed License Management section; needs more restrictions on use. Many editorial corrections from Graeme (thanks!)	2017-05-12	V1.0R25a
G Fedorkow	Editorial corrections from Rob Spiger (thanks!)	2017-05-21	V1.0R25b
G Fedorkow	Microscopic Editorial corrections	2017-05-23	V1.0R25c
G Fedorkow	No functional changes; cleared change bars, removed comment bubbles that have already been addressed	2017-05-24	V1.0R26a
G Fedorkow	Editorial corrections from review comments – commas, periods and semicolons, some clarification of wording	2017-06-05	V1.0R26b
G Fedorkow	Cleared change bars for review	2017-07-03	V1.0R27a

1787

1788