

The attached DRAFT document (provided here for historical purposes), released on May 30, 2017, has been superseded by the following publication:

Publication Number:     **NIST Special Publication (SP) 800-193**

Title:                     ***Platform Firmware Resiliency Guidelines***

Publication Date:        **May 2018**

- Final Publication: <https://doi.org/10.6028/NIST.SP.800-193> (which links to <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-193.pdf>).
- Related Information on CSRC:  
Final: <https://csrc.nist.gov/publications/detail/sp/800-193/final>

1 **Draft NIST Special Publication 800-193**

---

2 **Platform Firmware Resiliency**  
3 **Guidelines**

---

4  
5  
6 Andrew Regenscheid  
7  
8  
9  
10  
11  
12  
13

---

14 **C O M P U T E R S E C U R I T Y**

---

17  
18  
  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40

# Draft NIST Special Publication 800-193

# Platform Firmware Resiliency Guidelines

Andrew Regenscheid  
*Computer Security Division  
Information Technology Laboratory*

May 2017



41  
42  
43  
44  
45  
46  
47  
48

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

National Institute of Standards and Technology  
*Kent Rochford, Acting NIST Director and Under Secretary of Commerce for Standards and Technology*

49

## Authority

50 This publication has been developed by NIST in accordance with its statutory responsibilities under the  
51 Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 *et seq.*, Public Law  
52 (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including  
53 minimum requirements for federal information systems, but such standards and guidelines shall not apply  
54 to national security systems without the express approval of appropriate federal officials exercising policy  
55 authority over such systems. This guideline is consistent with the requirements of the Office of Management  
56 and Budget (OMB) Circular A-130.

57 Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and  
58 binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these  
59 guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce,  
60 Director of the OMB, or any other federal official. This publication may be used by nongovernmental  
61 organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would,  
62 however, be appreciated by NIST.

63 National Institute of Standards and Technology Special Publication 800-193  
64 Natl. Inst. Stand. Technol. Spec. Publ. 800-193, 46 pages (May 2017)  
65 CODEN: NSPUE2

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

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

76 NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at  
77 <http://csrc.nist.gov/publications>.  
78

79

80 **Public comment period: May 30, 2017 through July 14, 2017**

81 National Institute of Standards and Technology  
82 Attn: Computer Security Division, Information Technology Laboratory  
83 100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930  
84 Email: [sp800-193comments@nist.gov](mailto:sp800-193comments@nist.gov)

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

86

## Reports on Computer Systems Technology

87 The Information Technology Laboratory (ITL) at the National Institute of Standards and  
88 Technology (NIST) promotes the U.S. economy and public welfare by providing technical  
89 leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test  
90 methods, reference data, proof of concept implementations, and technical analyses to advance the  
91 development and productive use of information technology. ITL's responsibilities include the  
92 development of management, administrative, technical, and physical standards and guidelines for  
93 the cost-effective security and privacy of other than national security-related information in federal  
94 information systems. The Special Publication 800-series reports on ITL's research, guidelines, and  
95 outreach efforts in information system security, and its collaborative activities with industry,  
96 government, and academic organizations.

97

### Abstract

98 This document provides technical guidelines and recommendations supporting resiliency of  
99 platform firmware and data against potentially destructive attacks. The platform is a collection  
100 of fundamental hardware and firmware components needed to boot and operate a system. A  
101 successful attack on platform firmware could render a system inoperable, perhaps permanently  
102 or requiring reprogramming by the original manufacturer, resulting in significant disruptions to  
103 users. The technical guidelines in this document promote resiliency in the platform by  
104 describing security mechanisms for protecting the platform against unauthorized changes,  
105 detecting unauthorized changes that occur, and recovery from attacks rapidly and securely.  
106 Implementers, including Original Equipment Manufacturers (OEMs) and component/device  
107 suppliers, can use these guidelines to build stronger security mechanisms into platforms. System  
108 administrators, security professionals, and users can use this document to guide procurement  
109 strategies and priorities for future systems.

110

111

### Keywords

112 BIOS; Code signing; Firmware; Option ROM; Platform Firmware

113

114

115

## **Acknowledgements**

116 The author, Andrew Regenscheid of the National Institute of Standards and Technology (NIST),  
117 wishes to thank his colleagues who reviewed drafts of this document and contributed to its  
118 technical content. In particular, NIST appreciates the contributions from experts from industry  
119 and government who helped guide this work. These experts included Gary Simpson from AMD;  
120 Chirag Schroff from Cisco; Mukund Khatri from Dell; CJ Coppersmith, Gary Campbell, Shiva  
121 Dasari, and Tom Laffey from Hewlett Packard Enterprise; Jim Mann from HP Inc.; Charles  
122 Palmer from IBM; Bob Hale, David Riss, and Vincent Zimmer from Intel Corporation; and Paul  
123 England and Rob Spiger from Microsoft.

124 NIST would also like to acknowledge and thank Kevin Bingham, Cara Steib, and Mike Boyle,  
125 from the National Security Agency, who also provided substantial contributions to this  
126 document, as well as Jeffrey Burke from Noblis NSP.

127

128

## **Audience**

129 The intended audience for this document includes system and platform device vendors of  
130 computer systems, including manufacturers of clients, servers and networking devices. The  
131 technical guidelines assume readers have expertise in the platform architectures and are targeted  
132 primarily at developers and engineers responsible for implementing firmware-level security  
133 technologies in systems and devices.

134

135

## **Trademark Information**

136 All product names are registered trademarks or trademarks of their respective companies

137

138 **Executive Summary**

139 Modern client and server computing system architectures can be thought of in layers. The top  
140 layers are *software*, composed of the operating system and applications. While these provide  
141 most of the functional capabilities employed by users, they rely on functions and services  
142 provided by the underlying layers, which this document collectively refers to as the *platform*.  
143 The platform includes the hardware and firmware components necessary to initialize  
144 components, boot the system, and provide runtime services implemented by hardware  
145 components.

146 Platform firmware, and its associated configuration data, is critical to the trustworthiness of a  
147 computing system. Much of this firmware is highly privileged in the system architectures, and  
148 because this firmware is necessary for the system to operate, repairing this firmware can be  
149 challenging. A successful attack on platform firmware could render a system inoperable,  
150 perhaps permanently or requiring reprogramming by the original manufacturer, resulting in  
151 significant disruptions to users. Other sophisticated malicious attacks could attempt to inject  
152 persistent malware in this firmware, modifying critical low-level services to disrupt operations,  
153 exfiltrate data, or otherwise impact the security posture of a computer system.

154 Earlier NIST publications have addressed the threat of attacks on one particular type of platform  
155 firmware: boot firmware, commonly known as the Basic Input/Output System (BIOS).  
156 However, the platform consists of many other devices with firmware and configuration data.  
157 These devices, including storage and network controllers, graphics processing units, and service  
158 processors, are also highly-privileged and needed for systems to behave securely and reliably.

159 This document provides technical guidelines intended to support resiliency of platforms against  
160 potentially destructive attacks. These guidelines are based on the following three principles:

- 161 • **Protection:** Mechanisms for ensuring that Platform Firmware code and critical data  
162 remain in a state of integrity and are protected from corruption, such as the process for  
163 ensuring the authenticity and integrity of firmware updates.
- 164 • **Detection:** Mechanisms for detecting when Platform Firmware code and critical data  
165 have been corrupted.
- 166 • **Recovery:** Mechanisms for restoring Platform Firmware code and critical data to a state  
167 of integrity in the event that any such firmware code or critical data are detected to have  
168 been corrupted, or when forced to recover through an authorized mechanism. Recovery  
169 is limited to the ability to recover firmware code and critical data.

170 These guidelines are intended to address platforms in personal computer (PC) client, servers, and  
171 network devices, but should be broadly applicable to other classes of systems. Implementers,  
172 including Original Equipment Manufacturers (OEMs) and component/device suppliers, can use  
173 these guidelines to build stronger security mechanisms into platforms. System administrators,  
174 security professionals, and users can use this document to guide procurement strategies and  
175 priorities for future systems.

176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208

**Table of Contents**

**Executive Summary ..... iv**

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

    1.1 Purpose .....1

    1.2 Audience.....1

    1.3 Applicability and Scope.....1

    1.4 Document Structure.....2

**2 Platform Architecture ..... 3**

    2.1 Platform Devices .....4

    2.2 Code and Data in Platform Devices .....7

        2.2.1 Code ..... 7

        2.2.2 Data..... 7

**3 Principles and Key Concepts ..... 10**

    3.1 Principles Supporting Platform Resiliency.....10

    3.2 Resiliency Properties .....10

    3.3 Roots of Trust and Chains of Trust .....11

    3.4 Device Relationships .....12

    3.5 Firmware Update Mechanisms.....14

        3.5.1 Authenticated Update Mechanism..... 14

        3.5.2 Authorized Update Mechanism ..... 15

        3.5.3 Secure Local Update..... 16

    3.6 Other Considerations for Platform Resiliency.....16

        3.6.1 Management..... 16

        3.6.2 Authorization Mechanisms ..... 17

        3.6.3 Network-Assisted vs. Local Recovery..... 18

        3.6.4 Automated vs. Manual Recovery..... 18

        3.6.5 Event Logging..... 19

**4 Firmware Security Guidelines for Platform Devices..... 20**

    4.1 Roots of Trust and Chains of Trust .....20

        4.1.1 Roots of Trust (RoT) and Chains of Trust (CoT) ..... 20

        4.1.2 Root of Trust for Update (RTU) and Chain of Trust for Update (CTU) ..... 21

        4.1.3 Root of Trust for Detection (RTD) and Chain of Trust for Detection (CTD) 21



209 4.1.4 Root of Trust for Recovery (RTRec) and Chain of Trust for Recovery  
 210 (CTRec)..... 22  
 211 4.2 Protection.....22  
 212 4.2.1 Protection and Update of Mutable Code..... 22  
 213 4.2.2 Protection of Immutable Code ..... 24  
 214 4.2.3 Runtime Protection of Critical Platform Firmware ..... 24  
 215 4.2.4 Protection of Critical Data ..... 25  
 216 4.3 Detection.....25  
 217 4.3.1 Detection of Corrupted Code ..... 25  
 218 4.3.2 Detection of Corrupted Critical Data..... 27  
 219 4.4 Recovery.....27  
 220 4.4.1 Recovery of Mutable Code ..... 27  
 221 4.4.2 Recovery of Critical Data ..... 28

222  
 223

**List of Appendices**

224 Appendix A— Acronyms ..... 31  
 225 Appendix B— Glossary ..... 32  
 226 Appendix C— References ..... 38

227

228

**List of Figures**

229 Figure 1 – High-Level System Architecture..... 3  
 230 Figure 2: Roots of Trust..... 12  
 231 Figure 3: Trust Chains ..... 13  
 232 Figure 4: Example Update Scenario ..... 14

233

## 234 **1 Introduction**

### 235 **1.1 Purpose**

236 Modern computing and information technology systems are built upon a variety of hardware  
237 components that provide the fundamental capabilities required by the system to operate. Many  
238 of these hardware components have firmware and configuration data that drive their behavior,  
239 and which must remain in a good state in order for the system to function properly. One  
240 example of such firmware is commonly referred to as the Basic Input/Output System (BIOS),  
241 which is used facilitate the hardware initialization process and transition control to the operating  
242 system. Depending on the system, there may be tens or hundreds of microcontrollers with other  
243 kinds programmable firmware which support the overall system architecture. That collection of  
244 hardware and firmware components is typically called the *platform*.

245 The devices which made up the platform are crucial to integrity and availability of the systems  
246 built upon these platforms. Without these devices, systems may fail to operate correctly, or may  
247 not operate at all. Targeted attacks at certain devices within the platform could significant  
248 impact the security posture of these systems, possibly allowing a low-level, persistent malware  
249 presence. Destructive attacks which aim to damage or remove platform firmware have the  
250 potential to render systems permanently damaged, incurring substantial costs to the affected  
251 parties.

252 The purpose of this document is to provide security guidelines to make platforms more resilient  
253 to such destructive attacks. These guidelines address mechanisms to protect firmware and  
254 configuration data from attacks, as well as mechanisms to detect and recover from successful  
255 attacks.

### 256 **1.2 Audience**

257 The intended audience for this document includes system and platform device vendors of  
258 computer systems, including manufacturers of clients, servers and networking devices. The  
259 technical guidelines assume readers have expertise in the platform architectures and are targeted  
260 primarily at developers and engineers responsible for implementing firmware-level security  
261 technologies in systems and devices.

262 The material may also be of use when developing enterprise-wide procurement strategies and  
263 deployment. The material in this document is technically oriented, and it is assumed that readers  
264 have at least a basic understanding of computer security principles and computer architectures.  
265 The document provides background information to help such readers understand the topics that  
266 are discussed.

### 267 **1.3 Applicability and Scope**

268 The goal of this document is to provide principles and guidelines that can support platform  
269 resiliency. These principles and guidelines directly apply to the individual devices that make up  
270 a platform (see Section 2.1 for a list of examples). Specifically, they describe security  
271 mechanisms aimed at protecting each device from unauthorized changes (whether inadvertent or

272 due to malicious intent to its firmware or critical data and restoring the platform to a state of  
273 integrity.

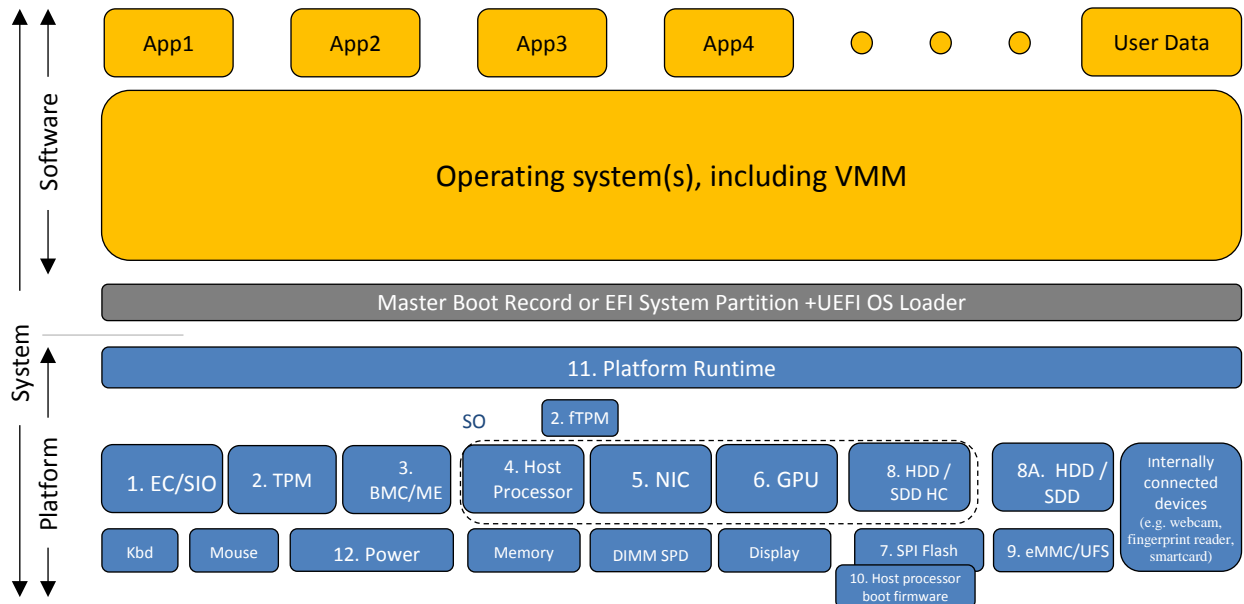
#### 274 **1.4 Document Structure**

275 The remainder of this document is organized into the following major sections:

- 276 • Section 2 provides informative material describing platform components and  
277 architectures.
- 278 • Section 3 describes the security principles that form the basis for the guidelines in this  
279 document, and describes key concepts for applying these principles to platform  
280 resiliency.
- 281 • Section 4 contains technical security guidelines for protection of firmware code and  
282 critical data, detection of authorized changes, and recovery to a state of integrity.
- 283 • Appendix A provides an acronym and abbreviation list for the document.
- 284 • Appendix B presents a glossary of selected terms from the document.
- 285 • Appendix C contains a list of references for the document.

286 **2 Platform Architecture**

287 Ensuring a platform’s firmware code and critical data are always in a state of integrity is critical  
 288 to ensure that a computing system can be operated free from malware. Modern client and server  
 289 computing systems can be considered to be separated into two high-level logical constructs,  
 290 *platform* and *software*. For the purposes of this document, we will describe the combination of  
 291 these two logical constructs as a system. Note that Figure 1 is merely illustrative and is not  
 292 intended to represent all possible devices in a platform, nor is it intended to represent an  
 293 exemplary architecture for any particular device. At a high-level, items in blue-shaded boxes are  
 294 devices to be considered part of a platform.



295 **Figure 1 – High-Level System Architecture**

296 Broadly speaking, the *platform* is comprised of hardware and firmware necessary to boot the  
 297 system to a point at which software, or an operating system, can be loaded<sup>1</sup>; *software* is  
 298 comprised of elements required to load the operating system and all applications and data  
 299 subsequently handled by the operating system. Note that some firmware continues to execute  
 300 once software has started. Existing industry best practice, as well as NIST publications such as  
 301 NIST SP 800-147 and NIST SP 800-147B, already address the issue of protecting the integrity of  
 302 a platform’s host processor boot firmware (traditionally called BIOS, and more recently UEFI  
 303 [3])<sup>2</sup> and its update mechanisms, but protection is only one of three key elements of cyber  
 304 resiliency (the other two being detection and recovery). Additionally, the resilience of other

<sup>1</sup> The OS loader is not included as part of the platform because it is loosely coupled to the platform. While the various implementations of Secure Boot makes that coupling tighter, ultimately the OEM of the platform is typically not responsible for the OS or OS loader in the same way that the OEM is responsible for the hardware and firmware on a platform. An end user can completely replace the OS and OS loader independent of the hardware and firmware, keeping the platform entirely intact.

<sup>2</sup> For purposes of this document, host processor boot firmware will be used generically to refer to either legacy Basic Input/Output System (BIOS) or Unified Extensible Firmware Interface (UEFI).

305 critical firmware on the platform has not yet been addressed to the level the host processor boot  
306 firmware has been. While it is beyond the scope of this document to identify and define every  
307 category and architecture of computing device which contains firmware, this document has  
308 applicability to any computing device (platform) which contains firmware, inclusive of PC's,  
309 servers, networking devices, smartphones, tablets, etc.

## 310 **2.1 Platform Devices**

311 As noted above, the platform is a collection of devices that provide the functional capabilities  
312 and services needed by the operating system and applications. While resiliency of the platform  
313 as a whole is the ultimate objective, it is important to recognize that the platform is composed of  
314 many different devices, often developed and manufactured by different vendors. For that reason,  
315 the technical guidelines in this document are described in terms of guidelines for individual  
316 platform devices.

317 For the purposes of describing a resilient platform, this section provides a list of devices which  
318 are often critical to the normal and secure operation of a platform. These devices typically  
319 contain mutable firmware, and are covered by the intended scope of the security guidelines in  
320 this document.

321 However, this should not be considered an exhaustive list of all devices in every platform of  
322 interest and platform vendors will need to carefully consider other devices which should be  
323 considered in scope for their particular platform.

324 In the case of a traditional x86-based platform (desktop, notebook, server, network switch), these  
325 devices are identified in Figure 1, and are defined below. Note that the numbers here reference  
326 the devices identified by their like numbers as indicated in the device box. The ordering is not  
327 meant to imply any priority or sequencing.

### 328 **1. Embedded Controller (EC) / Super I/O (SIO)**

329 An EC is typically associated with mobile platforms (notebooks, convertibles, tablets),  
330 while a SIO is typically associated with desk-based platforms (desktops, desk-based  
331 workstations, All-in-Ones, Thin Clients). This is not universally true, but is generally  
332 true enough to establish the type of client system in which one might find an EC or SIO.  
333 An EC or SIO typically controls functions in the platform such as the keyboard, LEDs,  
334 fans, battery monitoring/charging, thermal monitoring, etc. Additionally, it is typically  
335 the first system board device in the platform to execute code, even holding the host  
336 processor in reset until the EC/SIO is ready for the host processor to fetch its first line of  
337 host processor firmware code.

### 338 **2. Trusted Platform Module (TPM)**

339 A TPM [4] is a security coprocessor capable of securely storing and using cryptographic  
340 keys and measurements of the state of the platform. These capabilities can be used,  
341 among other things, to secure data stored on the system, provide a strong device identity,  
342 and to attest the state of the system. While not all platforms include or make use of a  
343 TPM, on any system in which a TPM is included and used, its firmware must be  
344 protected given its criticality in helping ensure the trustworthiness of the platform.  
345

346 TPM's also contain non-volatile memory storage which may contain critical data and, if  
347 so, must be protected. TPM's can be either discrete hardware devices, or may be realized  
348 in firmware executed on a platform host controller or other microcontroller (the latter are  
349 sometimes referred to as firmware TPM's, or fTPM).

350

### 351 **3. Baseboard Management Controller (BMC) / Management Engine (ME)**

352 A BMC is associated with server platforms while an ME is typically associated with  
353 client platforms. In both cases, a core aspect of their functionality is to serve as an out-  
354 of-band management device enabling platform administrators to manage a platform  
355 without requiring an operating system to be running. While not always strictly necessary  
356 to the basic computing function of a server or client platform, most modern server and  
357 client platforms include a BMC/ME, making it critical that their firmware cannot affect  
358 the state of integrity of the host processor's security domain.

359

### 360 **4. Host Processor [aka Central Processing Unit (CPU), aka Application Processing 361 Unit (APU)]**

362 The host processor is the primary processing unit in a typical platform, traditionally  
363 called a CPU, and now also sometimes referred to as an APU or an SoC. This is the  
364 processing unit on which the primary operating system (and/or hypervisor), as well as  
365 user applications run. This is the processor that is responsible for loading and executing  
366 the host processor firmware.

367

### 368 **5. Network Interface Controller (NIC)**

369 Whether discrete or integrated as part of an SoC, most modern client and server platforms  
370 have at least one NIC (wired or wireless), and could have multiple, including multiple  
371 types (wired, Wi-Fi, Wireless WAN). While having a NIC is not strictly required to boot  
372 a platform, in today's connected world it's important to have some form of connectivity  
373 at some point after a system has booted. More importantly, a compromised NIC  
374 firmware image could serve as a launch pad for other exploits in the system, be used to  
375 exfiltrate data, serve as a man-in-the-middle, etc. In addition to firmware run by a  
376 microcontroller, a NIC may include expansion ROM firmware which is loaded during  
377 boot and executed by the host processor. It is critical that a NIC's expansion ROM  
378 firmware is also protected. The expansion ROM firmware may be stored with the host  
379 processor boot firmware (in the case of an integrated NIC), or may be stored separately  
380 with the NIC itself as in the case of an add-in card. A NIC often also contains critical  
381 data, for example a Media Access Control (MAC) address may be stored in mutable  
382 memory. An attack on this critical data could result in a denial of service (DoS) both of  
383 this platform as well as of another system with a matching MAC address.

384

### 385 **6. Graphics Processing Unit (GPU)**

386 A GPU is a device that serves as the primary 'output' human interface device (HID) in  
387 client platforms. In addition, GPUs could serve as a launch pad for other exploits in the  
388 system. In addition to firmware run by a microcontroller, a GPU may include an  
389 expansion ROM firmware which is loaded during boot and executed by the host  
390 processor. The expansion ROM firmware may be stored with the host processor boot

391 firmware (in the case of an integrated GPU), or may be stored separately with the GPU  
392 itself as in the case of an add-in card.

393

#### 394 **7. Serial Peripheral Interface (SPI) Flash**

395 Most modern platforms include some amount of SPI flash to store firmware, typically for  
396 host processor boot firmware, though it could be used for other purposes.

397

#### 398 **8. A) Host Controller (HC) for mass storage devices**

399 For most modern platforms, some form of local mass storage in the form of either a HDD  
400 or SSD is required in order to boot an operating system and hold a user's applications and  
401 data. In order for the data to get stored on the mass storage device, a Host Controller  
402 (HC) is used to move the data from the platform's main memory to the physical storage  
403 medium over some storage bus (e.g. SATA, SCSI, PCIe). This HC has its own  
404 microcontroller and associated firmware. The HC could either be integrated into an SoC,  
405 or could be a separate device or on an add-in card.

406

#### 407 **B) Hard Disk Drive (HDD) / Solid State Drive (SSD)**

408 An HDD or SSD represents current state of the art in a traditional platform for storage of  
409 large quantities of data. These devices are coupled with the Host Controller. Within the  
410 HDD or SSD, a microcontroller and associated firmware are used to perform the actual  
411 storage operation of data sent from the platform's main memory to the mass storage  
412 device. Compromising a HDD's or SSD's firmware can also be used as a launch pad for  
413 other exploits in the system, or could be used to compromise a user and/or platform data.

#### 414 **9. embedded MultiMedia Card (eMMC) / Universal Flash Storage (UFS)**

415 eMMC and UFS are emerging as the standard mass storage devices for mobile systems.  
416 Each of them may include their own expansion ROM firmware and/or microcontroller  
417 with associated firmware.

418

#### 419 **10. Host Processor Boot Firmware [aka Basic Input/Output System (BIOS), aka Unified 420 Extensible Firmware Interface (UEFI)]**

421 In most modern platforms, host processor boot firmware is contained in a SPI flash  
422 device.

423

#### 424 **11. Platform Runtime Firmware**

425 In addition to boot firmware, there is platform runtime code. This is code which remains  
426 resident (typically in DRAM) and executable after the platform has booted. This is most  
427 typical for microcontrollers where firmware is required to execute to perform some  
428 function while the system is fully operational. An example of host processor firmware  
429 which is considered as runtime code would be System Management Mode (SMM) code.

430

#### 431 **12. Power Supply**

432 Some power supplies have their own microcontroller and associated firmware.

433

#### 434 **13. Glue Logic (CPLD's, FPGA's) – *not pictured***

435 Modern embedded systems use, programmable logic components to provide glue logic

436 functionality. There are two types of programmable logic components, Field  
437 Programmable Gate Arrays known as FPGAs and complex Programmable Logic Devices  
438 known as CPLDs. FPGAs are typically loaded with bitstream programs from attached  
439 flash devices on power up. CPLDs on the other hand are programmed with a bitstream  
440 once and then they retain the function until programmed again in the field. Typically this  
441 functionality is needed for basic operations of the system and if corrupted could result in  
442 bricking of a platform.

443

#### 444 **14. Fans – *not pictured***

445 Some fans have their own microcontroller and associated firmware.

446

## 447 **2.2 Code and Data in Platform Devices**

448 The devices describe above will typically contain some set of firmware and data on nonvolatile  
449 storage, either resident on the device itself or on a shared storage device (e.g., the SFI flash).

450 This section describes firmware code and data, and briefly discusses the scope of the document  
451 related to these components.

452

### 453 **2.2.1 Code**

454 Firmware code is the set of instructions used by the any device’s processing unit to perform the  
455 operations required by the device. Historically, firmware in platform devices is rarely modified,  
456 although system or component vendors may develop firmware updates which patch  
457 vulnerabilities, fix bugs, or add new functionality.

458 Because firmware is large part drives the behavior of a device, it is important that it remain in a  
459 trustworthy state on the platform. Attacks on the firmware code could render a device inoperable  
460 or inject malicious functionality into the device. Firmware should only be loaded from  
461 authorized source, typically either the manufacturer of the system or of the platform device.

462 The guidelines in this document describe mechanisms to protect firmware code by verifying  
463 updates using digital signatures. They also describe mechanisms to detect unauthorized changes  
464 to firmware, and secure methods of recovery.

### 465 **2.2.2 Data**

466 Data are pieces of information that Platform Firmware code uses to carry out its operation, as  
467 instructed by the code. Data can be further categorized as critical and non-critical. Critical data  
468 includes configuration settings and policies that are needed to be a valid state for the device to  
469 maintain its security posture. Non-critical data includes all other data.

#### 470 **2.2.2.1 Critical Data**

471 Critical data may be used for various purposes, including:

- 472 • **Configuration settings:** Data which tells the code how to configure operational aspects  
473 of the device

474 *Example: Enabling or disabling a peripheral when doing so would violate the*



475 *enterprise's security policies.*  
 476 *Example: The table of non-functional sectors in a hard drive.*

- 477 • **Policies:** Data which tells the code what path to take or how to respond
- 478 *Example: The system's boot order describes the valid devices to attempt to boot from as*
- 479 *well as the order.*
- 480 *Example: UEFI Secure Boot, a set of security configurations which control which third*
- 481 *party code the BIOS will hand control to.*

482 Critical data is difficult to precisely define because data that may be critical for one device may  
 483 not be critical for another. However, common characteristics of critical data include:

- 484 • It must be in a valid state for the proper booting and run-time operation of the device;
- 485 • It persists across power cycles (e.g. stored in non-volatile memory)
- 486 • It modifies the behavior or function of the device
- 487 • It must be in a valid state to support protection, detection and/or recovery of platform
- 488 firmware and associated data.

489  
 490 Some critical data is hard-coded in code, and updated only by means of a firmware image  
 491 update. For the purposes of this document, hard-coded data is considered part of the code, and  
 492 protected according to the firmware code protection, detection and recovery guidelines.

493 Platform devices often have other data that is configurable during its normal operation by  
 494 Platform Administrators, hardware, firmware or software. Because corruption of critical data  
 495 can interfere with the normal or secure operation of a system, it is important to protect critical  
 496 data from corruption, and to be able to recover when problems are detected. However, strong  
 497 protection of some forms of critical data can be architecturally difficult, due to expectations that  
 498 some entities, such as operating systems and device drivers, have access to change these settings.

499 Some configuration data can only be changed through defined interfaces controlled by platform-  
 500 level code. For example, UEFI runtime variables fall into this category. This basic level of  
 501 protection guards against attackers from directly modifying configuration data, and allows  
 502 Platform Firmware to validate input before committing changes to storage. However, malicious  
 503 entities may be able to use these defined interfaces to make well-formed, but incorrect,  
 504 configuration changes.

505 To guard against such tampering, changes to some particularly sensitive configuration data may  
 506 require authorization before changes can be applied using the defined interfaces described above.  
 507 In some cases, platform devices, such as host processor boot firmware or service processor  
 508 firmware, may be capable of authenticating Platform Administrators prior to allowing them to  
 509 make changes. Stronger authentication techniques may allow Platform Firmware to  
 510 cryptographically verify the source and integrity of changes.

511 Some critical data is managed by the firmware with no programmatic exposure through external  
 512 interfaces (e.g., wear leveling data) and if lost or damaged can result in permanent loss of service  
 513 of the device. This type of state data needs to be protected at the highest level and is not  
 514 writeable from the rest of the platform.

515 In practice, all data consumed by platform firmware may be security-sensitive, including some  
516 data that does not directly impact the correct and secure operation of the platform. Errors or  
517 malicious attacks in any data consumed by platform firmware could expose and exploit  
518 vulnerabilities in that code. As such, particular care needs to be given to any non-trusted input or  
519 data consumed by the platform.

#### 520 **2.2.2.2 Non-Critical Data**

521 Non-critical data may be used for various purposes, including:

- 522 • **Informational / UI:** Data which is merely informational or used as part of a user  
523 interface for the end user  
524 *Example: An asset tag name of “Property of NIST” is displayed during boot*
  
- 525 • **State:** State settings which do not affect the integrity of the platform  
526 *Examples: The state of the Num Lock key upon system boot; whether the BIOS performs*  
527 *a fast boot or standard boot*

528 Non-critical data is not deemed critical to the secure booting or operation of a platform and is  
529 therefore not addressed by the security guidelines in this document.

## 530 **3 Principles and Key Concepts**

531 This section provides a brief description of the driving principles for platform resiliency which  
532 provide the foundation for the guidelines in this document. It also discusses major architectural  
533 concepts and considerations used throughout the document.

### 534 **3.1 Principles Supporting Platform Resiliency**

535 The security guidelines in this document are based on the following three principles:

- 536 • **Protection:**  
537 Mechanisms for ensuring that Platform Firmware code and critical data remain in a state  
538 of integrity and are protected from corruption, such as the process for ensuring the  
539 authenticity and integrity of firmware updates.  
540
- 541 • **Detection:**  
542 Mechanisms for detecting when Platform Firmware code and critical data have been  
543 corrupted.  
544
- 545 • **Recovery:**  
546 Mechanisms for restoring Platform Firmware code and critical data to a state of integrity  
547 in the event that any such firmware code or critical data are detected to have been corrupted,  
548 or when forced to recover through an authorized mechanism. Recovery is limited to the  
549 ability to recover firmware code and critical data.

550 The technical guidelines found in Section 4 are organized around these principles. The first  
551 principle, protection, is similar in scope and purpose to the guidelines found in NIST SP 800-147,  
552 *BIOS Protection Guidelines*. The basic principle of protection is expanded in this document to  
553 apply to a broader set of firmware and configuration data within the platform.

554 While protection mechanisms are intended to prevent destructive or malicious attacks against  
555 platform firmware and critical data, these mechanisms may be imperfect or impractical to  
556 implement on all categories of devices. In those cases, detection and recovery mechanisms are  
557 intended to discover and remediate attacks to regain normal and secure operation on the device.

### 558 **3.2 Resiliency Properties**

559 The technical guidelines in this document are written in terms of guidelines for individual  
560 platform devices in order to make them broadly applicable to a variety of devices, platforms and  
561 systems. Despite the narrow focus on devices, the intent of this document is to establish  
562 guidelines supporting overall resiliency in systems against destructive attacks by ensuring that  
563 the underlying platform is resilient.

564 Platforms may not be able to fully provide the protection, detection and recovery capabilities for  
565 all platform devices. A loss of functionality in even one device may be sufficient to render the  
566 complete system permanently inoperable if that particular device plays a crucial role in booting

567 or operating the platform. For a platform as a whole to claim resiliency to destructive attacks,  
568 the set of platform devices necessary to minimally restore operation of the system, and sufficient  
569 to restore complete functionality, should themselves be resilient. We call this set of devices  
570 *critical platform devices*. The particular resiliency properties may vary from platform-to-  
571 platform.

### 572 ***Protected***

573 For a platform to be considered *Protected*, all critical platform devices should meet the protection  
574 guidelines found in Sections 4.1 and 4.2, but may not fully offer capabilities to recover the device's  
575 firmware and/or critical data. In the event that a device is able to detect a corrupted state but cannot  
576 recover, then the device may at least be able to signal a user or IT administrator about this event  
577 and/or stop the device from operating in order to prevent the platform from booting to allow the  
578 administrator to attempt to manually recover the device.

### 579 ***Recoverable***

580 For a platform to be considered *Recoverable*, all critical platform devices should provide the means  
581 to detect corruption as described in Sections 4.1 and 4.3, and provide the means to recover from  
582 this corruption in compliance with the guidelines in Sections 4.1 and 4.4.

### 583 ***Resilient***

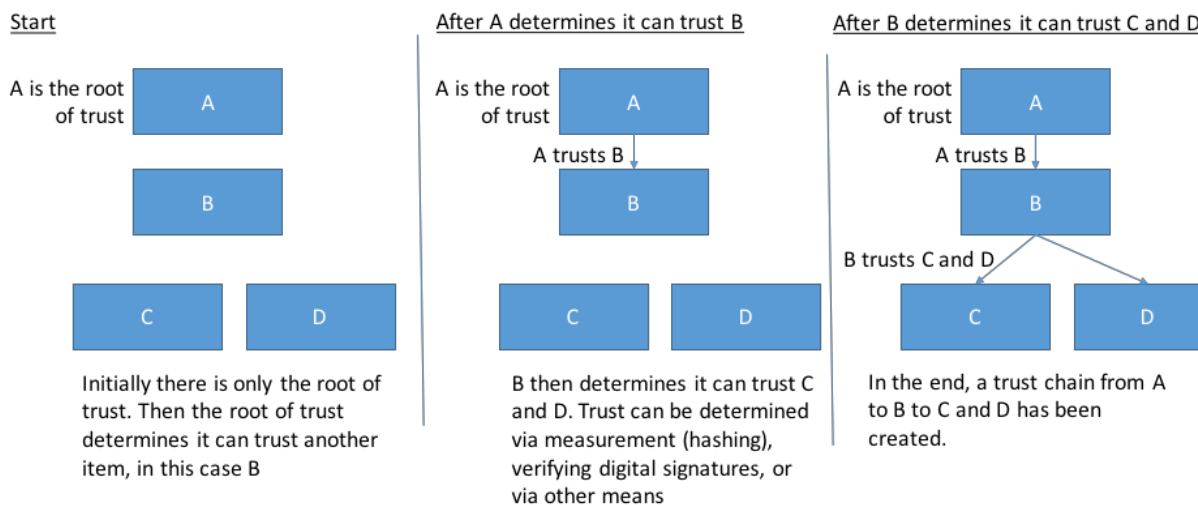
584 In this mode, all critical platform devices should meet all of the guidelines in Section 4. This mode  
585 is the ideal situation and the desired end-state for all devices in a platform.

## 586 **3.3 Roots of Trust and Chains of Trust**

587 The security mechanisms described in this document are founded in Roots of Trust (RoT). A  
588 Root of Trust is an element that forms the basis of providing one or more security-specific  
589 functions, such as measurement, storage, reporting, recovery, verification, update, etc. A RoT is  
590 trusted to always behave in the expected manner because its proper functioning is essential to  
591 providing its security-specific functions and because its misbehavior cannot be detected. A RoT  
592 is typically just the first element in a Chain of Trust (CoT) and can serve as an anchor in a chain  
593 of trust to deliver more complex functionality. The responsibilities and capabilities of a RoT  
594 may be implemented entirely within the RoT or may be performed by a delegate or agent  
595 spawned by its RoT via a chain of trust anchored in the RoT. For example, a RoT for recovery  
596 (RTRec), when triggered, will initiate a recovery process by launching another element that  
597 determines an appropriate recovery sequence and launches a chain of successive elements that  
598 perform the recovery actions.

599 Generally, successive elements are cooperative in maintaining the chain of trust started by the  
600 RoT. Components in a chain of trust are privileged to perform security critical functions like  
601 performing device updates that are not available to less trusted software. Roots of Trust and  
602 Chains of Trust may have mechanisms to relinquish these privileges once the security function is  
603 complete, or if it is determined that the security function is not required. A chain of trust may  
604 also relinquish privileges before passing control to a non-cooperative element.

605



**Figure 2: Roots of Trust**

606 Because Roots of Trust are essential to providing critical security functions, they need to be  
 607 secure by design. Major considerations for determining confidence in Roots of Trust are an  
 608 analysis of the attack surface of a Root of Trust and an evaluation of the mitigations used to  
 609 protect that attack surface. The responsibility of ensuring the trustworthiness of a Root of Trust is  
 610 on the vendor which provides the Root of Trust. Vendors typically protect Roots of Trust by  
 611 either making them immutable, or by ensuring that the integrity and authenticity of any changes  
 612 to a RoT are verified prior to performing such updates. Often, Roots of Trust run in isolated  
 613 environments, at greater privilege level than anything which could modify it, and/or complete  
 614 their function before anything can modify it to ensure that other devices cannot compromise their  
 615 behavior of them during operation.

616 Section 4.1 of this document provides specific guidelines on the capabilities and properties of the  
 617 Roots of Trust that support platform resiliency.

618 Platforms are often composed of numerous devices, often with isolation boundaries between  
 619 devices and different manufacturers. A platform may need multiple independent Roots of Trust  
 620 and Chains of Trust to provide comprehensive coverage for resiliency. For example, a hard disk  
 621 controller may have a separate microcontroller and firmware than the host platform. Both the  
 622 hard disk controller and the host platform may need their own independent chain of trust for  
 623 recovery if their individual critical data become corrupted.

### 624 3.4 Device Relationships

625 Due to lack of capability or functionality, some platform devices may not have their own root(s)  
 626 of trust to perform an update, detection, or recovery. We refer to devices needing assistance as  
 627 symbiont devices and those lending assistance as host devices. A dependency may be established  
 628 whereby a host device and symbiont device jointly fulfill the guidelines for protection, detection

Examples of different trust chains in a platform




Figure 3: Trust Chains

629 and/or recovery that the symbiont device cannot fulfill independently. Such dependencies might  
 630 leverage a secure communication channel or other techniques. To be effective at lending  
 631 assistance, the host device needs to meet the guidelines itself for the mechanisms it helps convey  
 632 to the symbiont device. Together, the host and symbiont device provide a CoT that implements  
 633 the security guidelines for protection, detection and recovery.

634 There may be relationships between devices where the trust is implicit- that is, where trust is  
 635 provided by the architecture of the system. A device may receive indication of unambiguous  
 636 physical presence from a device where an implicit trust relationship already exists. The fact that  
 637 the other device sent the message means that the device can trust the request.

638 There may also be other relationships which do not imply nor require any level of trust. Consider  
 639 a device responsible for receiving updates. That device may then propagate those updates to  
 640 other devices. Since each device (or the symbiont device along with its host device) is  
 641 responsible for verifying its own updates, there is no requirement for trust between the device  
 642 distributing updates and the devices those updates are provided to.

### 643 3.5 Firmware Update Mechanisms

644 A central tenet to the firmware protection guidelines is ensuring that only authentic and  
 645 authorized firmware update images may be applied to platform devices. An update image is  
 646 authentic if the source (e.g., the device or system manufacturer) and integrity can be successfully  
 647 verified. Technical processes to verify images before using applying updates are called  
 648 *authenticated update mechanisms*.

649 Authorization, however, is the permission to perform an update. While authentication is  
 650 typically rooted in the device or system manufacturer, authorization to perform updates is  
 651 typically rooted in the device or system owner.

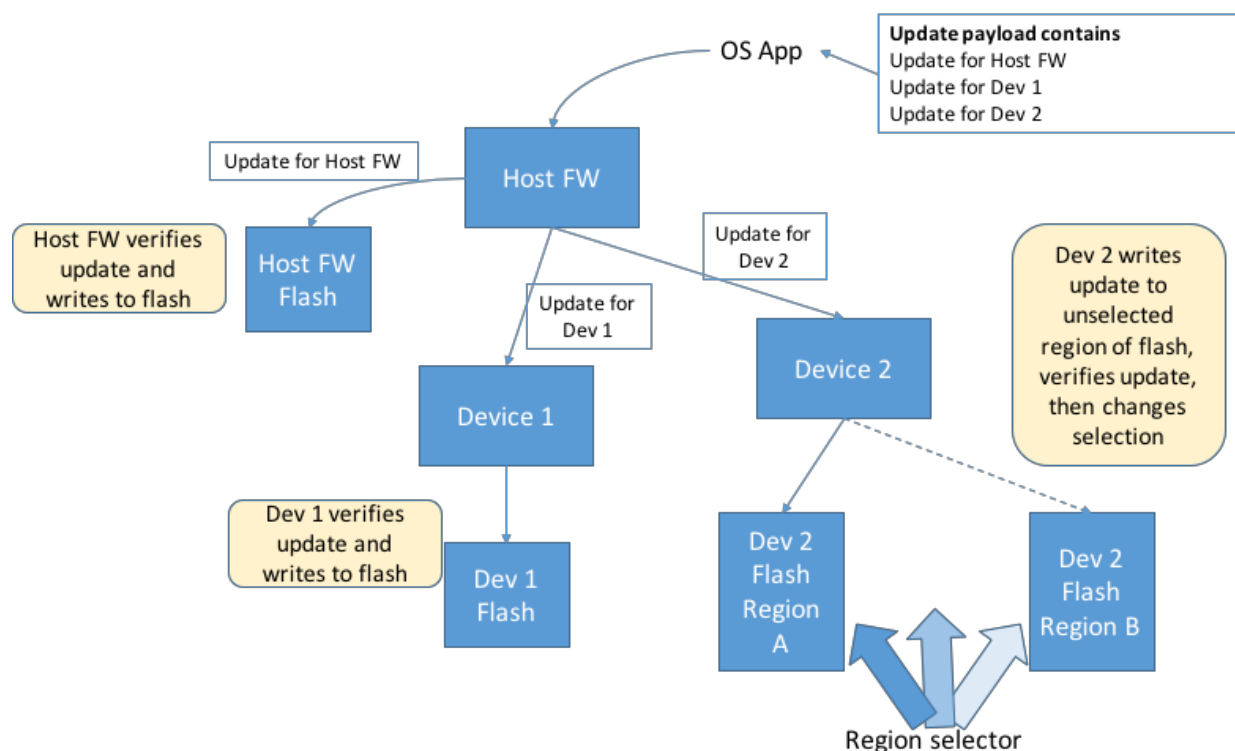


Figure 4: Example Update Scenario

652

#### 653 3.5.1 Authenticated Update Mechanism

654 An authenticated update mechanism employs digital signatures to ensure the authenticity of the  
 655 firmware update image. An update of the firmware image using an authenticated update  
 656 mechanism relies on a Root of Trust for Update (RTU) that contains a signature verification  
 657 algorithm and a key store that includes the public key needed to verify the signature on the  
 658 firmware update image. The key store and the signature verification algorithm are stored in a  
 659 protected fashion on the computer system and are modifiable only through use of an  
 660 authenticated update mechanism or a secure local update mechanism.

661 The key store in the RTU includes a public key used to verify the signature [7] on a firmware  
662 update image or includes a hash [6] of the public key if a copy of the public key is provided with  
663 the firmware update image. In the latter case, the update mechanism hashes the public key  
664 provided with the firmware update image and ensures that it matches a hash which appears in the  
665 key store before using the provided public key to verify the signature on the firmware update  
666 image.

667 It is possible that the private key corresponding to the public key in the key store may become  
668 “compromised”, for example, by the private key being stolen and published. An attacker that  
669 gains access to this key could sign invalid firmware that might damage platform devices or inject  
670 malware into the platform. Proper use of signatures thus necessitates provisions to recover from  
671 a key compromise. A variety of techniques may be used to recover from these cases. Examples  
672 range from the complex, including key hierarchies, to simpler, including updating the key store  
673 when recovering (or updating) the rest of an image.

### 674 **3.5.2 Authorized Update Mechanism**

675 A system and its supporting management software and firmware may provide several authorized  
676 mechanisms for legitimately updating a firmware image. These include:

- 677 1) **User-Initiated Updates:** Vendors typically supply end users with utilities capable of  
678 updating a firmware image. This could be from external media to perform these updates,  
679 or via utilities that can update the firmware image from the user’s normal operating  
680 system. Depending on the security mechanisms implemented on the system, these  
681 utilities might directly update the firmware image or they may schedule an update for the  
682 next system reboot. The updated code will encounter critical data written by a different  
683 revision of the code. The updated code should ensure that the platform continues to  
684 function by remaining compatible with the critical data, by updating the critical data to be  
685 compatible with the updated code, or, at least by resetting the critical data values to their  
686 defaults.
- 687 2) **Managed Updates:** A given computer system may have hardware and software-based  
688 agents that allow a system administrator to remotely update the firmware image without  
689 direct involvement from the user.
- 690 3) **Rollback:** Implementations that authenticate updates before applying them may also  
691 check version numbers during the update process. In these cases, the firmware image  
692 may have a special update process for rolling back the installed firmware to an earlier  
693 version. For instance, the rollback process might require the physical presence of the  
694 user. This mechanism guards against attackers flashing old firmware with known  
695 vulnerabilities.
- 696 4) **Manual Recovery:** To recover from corrupt or malfunctioning firmware, computer  
697 systems may provide mechanisms to allow a user with physical presence during the boot  
698 process to replace a firmware image with a known good version and configuration.



699 5) **Automatic Recovery:** Some computer systems are able to detect when a firmware image  
700 has been corrupted and recover from a backup firmware image stored in a separate  
701 storage location from the corrupted image (e.g., a second flash memory chip, a hidden  
702 partition on a hard drive).

### 703 **3.5.3 Secure Local Update**

704 While this document recommends firmware updates be done through an authenticated update  
705 mechanism as described in Section 3.5.1, some devices may also support a secure local update  
706 mechanism. These mechanisms instead authorize firmware updates through a process  
707 demonstrating unambiguous physical presence. A secure local update mechanism can be used,  
708 for example, to recover a corrupted firmware image that cannot be updated using an  
709 authenticated update or an automated recovery mechanism. The secure local update mechanism  
710 could also be used by a physically-present administrator to update to an earlier firmware image  
711 on a device that does not allow rollback.

712 To protect against remote attacks exploiting secure local update mechanisms, it is important that  
713 these mechanisms verify that a user has physically authorized the update. Remote mechanisms,  
714 such as interacting with a device or system via a remote console, do not satisfy this requirement  
715 for physical presence. Similarly, mechanisms that can be spoofed by malware running on the  
716 system or device do not satisfy this requirement. See Section 3.6.2 for more details.

717 However, note that devices that implement the secure local update mechanism are potentially  
718 vulnerable to attacks by rogue administrators or other attacks with physical access to the device  
719 or system. Therefore, additional physical, environmental and technical security measures are  
720 essential to protecting these devices, but they are beyond the scope of this document.

## 721 **3.6 Other Considerations for Platform Resiliency**

722 This document does not address certain other considerations that a purchaser, user, or IT  
723 administrator may take into account as it pertains to platform cyber resiliency. A non-exhaustive  
724 list and discussion of these other considerations follows.

### 725 **3.6.1 Management**

726 Vendors should carefully consider their target customers when designing resilient platforms to  
727 ensure proper management and control of policies and configuration settings can be administered  
728 in the way which best serves their needs. Management of policies and configuration settings can  
729 be performed either locally or remotely. Depending on platform type, customers may expect to  
730 be able to fully administer a platform securely from a remote location. Some customers may  
731 expect to require a physically present user to approve a change in policy. Customers may expect  
732 to be able to remotely extract any log data, or they may wish to prevent the exfiltration of log  
733 data except through authorized local mechanisms.

### 734 3.6.2 Authorization Mechanisms

735 Some recovery and administrative actions can make significant changes to either the Platform  
736 Firmware or software. For example, firmware settings may control the boot order and a  
737 software recovery agent may restore a backup erasing recently created data. Modifying these  
738 settings could require Platform-level Authorization to demonstrate that the entity requesting a  
739 change is authorized to do so. For some environments, like large organizations or data centers, a  
740 professional Platform Administrator may authorize actions remotely using credentials  
741 provisioned to manage the platform. In other environments, like for consumers or smaller  
742 enterprises, there may not be a remote Platform Administrator. Some systems, however, may  
743 have Platform Administrator credentials that can be used locally. Alternatively, some systems  
744 may allow users to assert platform-level authorization by ensuring that a physically-present user  
745 has issued a command or requested a change. On these systems, the platform must  
746 unambiguously verify that a physically present user has authorized the action. If done correctly,  
747 malware cannot impersonate an authorization check that involves confirmation from a physically  
748 present user. We use the term Unambiguous Physical Presence to indicate a local user that  
749 cannot be impersonated by malware.

750 Unambiguous Physical Presence allows for assertion of Platform-level Authorization (or a  
751 portion of Platform-level Authorization) by demonstrating that a person is physically interacting  
752 with a device or platform. By ensuring that recovery actions or critical data changes are  
753 authorized by a physically present person, Unambiguous Physical Presence provides a  
754 management path that is intended to be protected from influence by malware.

755 Creating platforms and devices that properly and reliably verify confirmation of a physically  
756 present person is complex. Dedicated physical buttons or hardware jumpers could provide a  
757 relatively direct and explicit method by which to demonstrate physical presence. Platform design  
758 or deployment considerations may prevent a person from having direct physical mechanisms to  
759 interact with each device supporting a function that relies on Unambiguous Physical Presence. In  
760 these cases, there will need to be a trusted path between the mechanisms used to verify UPP, and  
761 the device which will perform an action on behalf of the Administrator. To satisfy the non-  
762 bypassability guidelines found later in this document, this trusted path, which could include I/O  
763 devices (e.g., human interface device, graphics card, etc.) and internal buses, needs to be  
764 protected from manipulation by malware.

765 There are a variety of techniques that could provide a trusted path between a physical mechanism  
766 that verifies Unambiguous Physical Presence and a platform or device. One example could be to  
767 accept or confirm commands from a physically-present person only when the platform can be  
768 trusted to be in a good state, before malware could disturb these processes, such as early in the  
769 boot process. In other cases, system architectures may provide trusted paths between a Service  
770 Processor (e.g., EC, BMC) and other platform devices.

771 Devices which rely on UPP in place of Platform Administrator credentials to authorize  
772 administrative actions may be vulnerable to attacks by individuals with physical access to the  
773 device. As such, its usage may not be suitable in applications or environments that lack strong  
774 physical security.

### 775 3.6.3 Network-Assisted vs. Local Recovery

776 In most cases, the ability to recover locally from corruption will be the most expedient, provide  
777 the highest level of customer satisfaction, and may be necessary if there is no network  
778 connectivity. But it is recognized that this is not always possible, particularly given storage  
779 limitations that many devices will have. In instances where local recovery is not possible,  
780 network-assisted recovery can be implemented if done in a secure and trustworthy manner,  
781 which may include the use of encryption, digital signatures, secure transport methods, etc. While  
782 either local or network-assisted recovery are acceptable implementation mechanisms, the ability  
783 for a device to support both provides for an even higher level of resiliency and is therefore  
784 recommended.

### 785 3.6.4 Automated vs. Manual Recovery

786 Recovery can proceed in one of three ways:

- 787 1. **Fully automated** -- no user interaction required to initiate recovery or during recovery  
788 process
- 789 2. **Partially automated** -- recovery initiated automatically but requires user interaction at  
790 some point during the recovery process
- 791 3. **Manual** -- user interaction required to initiate recovery

792 Fully automated recovery mechanisms may be preferred by some users, as this can allow for  
793 faster recovery at scale in the event of a mass attack.

794 Fully automated recovery may not be supported by all systems or desired by all users. For  
795 example, systems may require administrative credentials or authorization to continue with the  
796 recovery process.

797 Manual recovery may be preferred by some users so that a Platform Administrator is informed  
798 that something is wrong, and then wait for that administrator to decide what steps will be taken  
799 next. This can also be useful in the event that a Platform Administrator wishes to capture  
800 information in order to help with forensic analysis.

801 Administrator-defined policies typically define the behavior and privilege requirements for  
802 manual recovery. Such policies may also affect automatic recovery. For example, an  
803 administrator-defined policy may constrain the versions of firmware that may be installed during  
804 the recovery. Those who set recovery policies must do so with care. The firmware version set in  
805 the policy might well be the version that was successfully attacked necessitating the recovery.  
806 Simply rewriting the vulnerable version may lead to an attack/recovery cycle.

807 Policy itself may be a target of attack so the design of the recovery implementation must account  
808 for the case that policy is not available. Also, since recovery can be a multi-step process, a policy  
809 requirement that will be met by the end of the recovery process might not be met during  
810 intermediate steps.

811 Recovery schemes which write over firmware images might, in the process, destroy evidence  
812 that would be useful in the analysis of the attack. Recovery schemes should, where practical,  
813 provide methods to retain or record attacked images and other information in cases where  
814 recovering the firmware may lose that information.

### 815 **3.6.5 Event Logging**

816 Logging firmware and recovery-related can often be useful for multiple purposes, including but  
817 not limited to:

- 818 • Forensic analysis which allows for a Platform Administrator of a system to capture  
819 information which might have led to an attack on a platform or actual platform  
820 compromise. This can be useful in determining if the platform might contain an  
821 unknown security vulnerability, or understanding if there might be a widespread attack of  
822 similar nature.
- 823 • Providing an audit trail to know when an event has occurred and if an update or recovery  
824 was authorized, who authorized it and when.

825 Platform and device manufacturers need to determine what level of event logging might be  
826 required for their systems, taking into account the intended users environments for those  
827 systems. Events which are logged should be recorded in a manner which provides assurances of  
828 their integrity and allows for the secure recovery and transmission of logged events. Care must  
829 be taken to insure even log access is controlled. Unauthorized personnel can use event log data  
830 analysis to broaden the attack surface.

## 831 **4 Firmware Security Guidelines for Platform Devices**

832 This section details the technical security guidelines for devices in a platform for each of the  
833 three elements of resiliency: protection, detection, and recovery. Devices may implement the  
834 requirements in one or more of these sections based on the firmware resiliency properties, as  
835 defined in Section 3.2, they aim to support. Section 4.1 provides foundational security  
836 guidelines for the Roots of Trust that support those properties. Section 4.2 provides security  
837 guidelines for the protecting firmware code and critical data. Guidelines for mechanisms to  
838 detect authorized changes to firmware and data are described in Section 4.3. Finally, Section 4.4  
839 specifies security guidelines for firmware and data recovery mechanisms.

840 While the guidelines are written in terms of applying to individual devices, a device may  
841 implement these guidelines with assistance from another device. Security functionality may be  
842 done by the device itself (self-contained) or it may rely upon a security architecture whereby  
843 another platform device provides some or all those security functions for that device. The  
844 reliance on another device to provide necessary security functionality creates a highly-trusted  
845 device relationship, as described in Section 3.4. These guidelines refer to the device relying on  
846 security functionality from another device as a *symbiont*, and the device providing that  
847 functionality for the symbiont as a *host device*. In these cases, the symbiont and host together  
848 form the Root of Trust or Chain of Trust responsible for implementing the security functions.  
849 The host device must additionally meet all of the requirements for a self-contained device.

850 The use of **shall**, **should**, and **may** are used as defined in RFC 2119 [5].

### 851 **4.1 Roots of Trust and Chains of Trust**

852 This section provides foundational guidelines on the Roots of Trust and Chains of Trust that  
853 support the subsequent guidelines for Protection, Detection, and Recovery.

#### 854 **4.1.1 Roots of Trust (RoT) and Chains of Trust (CoT)**

- 855 1) The security mechanisms **shall** be founded in Roots of Trust (RoT).
- 856 2) If Chains of Trust (CoT) are used, a RoT **SHALL** serve as the anchor for the CoT.
- 857 3) All RoTs and CoTs **shall** either be immutable or protected using mechanisms which  
858 ensure all RoTs and CoTs remains in a state of integrity.
- 859 4) All elements of the CTU, CTD and CTRec in non-volatile storage **shall** be implemented  
860 in platform firmware.

861 *Note: This guideline that RoTs and CoTs be implemented as part of platform*  
862 *firmware applies only to elements that implement the platform resiliency functions*  
863 *described in this paper. Platform vendors are encouraged to maintain a chain of*  
864 *trust from boot firmware through the Operating System to provide resiliency*  
865 *against various forms of attacks.*

- 866 5) Software running under, or as part of, the operating system on the host processor **shall**  
867 **not** be capable of interfering with the functions of the RoTs or CoTs.

- 868 6) Information transferred from the software on the host processor to the platform firmware  
869 **shall** be treated as untrusted.
- 870 7) CoTs **may** be temporarily extended to include elements that are not from non-volatile  
871 storage. Before use, those elements **shall** be cryptographically verified by an earlier  
872 element of the CoT.
- 873 8) RoTs and CoTs that cross device boundaries, or that provide services to a symbiont  
874 device, **shall** use a secure communication channel between components.

#### 875 4.1.2 Root of Trust for Update (RTU) and Chain of Trust for Update (CTU)

- 876 1) Each platform device with mutable firmware **shall** rely on either a Root of Trust for  
877 Update (RTU), or a Chain of Trust for Update (CTU) which is anchored by an RTU, to  
878 authenticate firmware updates.
- 879 2) If the RTU or CTU is mutable, then the RTU or CTU elements **shall** be updated using an  
880 authenticated update mechanism, absent physical intervention through a secure local  
881 update. During such an update, the RTU or CTU **shall** always be operational or  
882 recoverable upon a subsequent reboot even in the event of an unexpected, catastrophic  
883 event (e.g., power loss in the middle of a flash write operation).
- 884 3) The RTU or CTU **shall** include a key store and an approved digital signature algorithm  
885 implementation from FIPS 186-4 [7] to verify the digital signature of firmware update  
886 images.
- 887 4) If the keystore is updateable, then the keystore **shall** be updated using an authenticated  
888 update mechanism, absent unambiguous physical presence through a secure local update.
- 889 *Note: Updatable keystores provide a means to recover from compromise of the*  
890 *signing key, but may the device's keystore more vulnerable to tampering.*  
891 *Implementers that use a non-updatable key store are encouraged to design their*  
892 *devices to be resilient to disclosure of the firmware signing keys.*
- 893 5) An authenticated update mechanism anchored in the RTU **shall** be the exclusive means  
894 for updating device firmware, absent unambiguous physical presence through a secure  
895 local update.

#### 896 4.1.3 Root of Trust for Detection (RTD) and Chain of Trust for Detection (CTD)

- 897 1) Each platform device which implements a detection capability **shall** rely on either a Root  
898 of Trust for Detection (RTD), or a Chain of Trust for Detection (CTD) which is anchored  
899 by an RTD, for its detection.
- 900 2) The RTD or CTD **shall** include or have access to information necessary to detect  
901 corruption of firmware code and critical data.
- 902 3) Detection mechanisms anchored in the RTD **shall** provide the detection capabilities  
903 specified in Section 4.3.

904 *Note: This document provides minimum requirements for detection capabilities*  
 905 *rooted in low-level hardware and firmware to provide resiliency against*  
 906 *destructive attacks. However, nothing in this document should be construed as*  
 907 *disallowing other detection capabilities that are outside this trust chain.*

#### 908 **4.1.4 Root of Trust for Recovery (RTRec) and Chain of Trust for Recovery (CTRec)**

909 1) Each platform device which implements a recovery capability SHALL rely on either a  
 910 Root of Trust for Recovery (RTRec), or a Chain of Trust for Recovery (CTRec) which is  
 911 anchored by an RTRec, for its recovery.

912 2) The RTRec or the CTRec SHALL perform the recovery.

913 *Note: RTR was not chosen as the acronym for Root of Trust for Recovery because*  
 914 *RTR is typically used to denote Root of Trust for Reporting. As such, throughout*  
 915 *this document we will disambiguate RTR by using RTRec to denote Root of Trust*  
 916 *for Recovery.*

## 917 **4.2 Protection**

918 While previous efforts have addressed protection of BIOS (e.g., NIST SP 800-147 [1], NIST SP  
 919 800-147B [2]), there remains other security-critical firmware in the platform that has not been  
 920 addressed. These include option ROMs, management controllers, service processors, firmware  
 921 on disk/flash drives, network controllers, and graphics processing unit. Protection must also  
 922 extend to critical data associated with the firmware being protected, as some of this data could be  
 923 a vector of attack which can compromise the integrity of the platform.

924 All platform devices which provide protection of firmware code and critical data must meet the  
 925 requirements which follow.

### 926 **4.2.1 Protection and Update of Mutable Code**

927 This section specifies guidelines for firmware protection based on the principles of authenticated  
 928 firmware updates, integrity protection, and non-bypassability of security mechanisms.  
 929 Authenticated update mechanisms use digital signatures verify the integrity and authenticity of  
 930 firmware update images. Firmware integrity protections prevent unintended or malicious  
 931 modification of firmware outside the authenticated BIOS update process. The final principle,  
 932 non-bypassability, ensures that there are no means for an attacker to bypass the protective  
 933 mechanisms.

#### 934 **4.2.1.1 Authenticated Update Mechanism**

935 One or more authenticated update mechanisms anchored in the RTU **shall** be the exclusive  
 936 means for updating device firmware, absent unambiguous physical presence through a secure  
 937 local update, as defined in Section 3.5.3. Authenticated update mechanisms **shall** meet the  
 938 following authentication guidelines:

939 1) Firmware update images shall be signed using an approved digital signature algorithm as  
 940 specified in FIPS 186-4 [7], Digital Signature Standard, with security strength of at least

- 941 112 bits in compliance with SP 800-57, Recommendation for Key Management – Part 1:  
942 General [8].
- 943 2) Each firmware update image shall be signed by an authorized entity – usually the device  
944 manufacturer, the platform manufacturer or a trusted third party - in conformance with SP  
945 800-89, Recommendation for Obtaining Assurances for Digital Signature Applications  
946 [9].
- 947 3) The digital signature of a new or recovery firmware update image shall be verified by an  
948 RTU, a CTU prior to the non-volatile storage completion of the update process. For  
949 example, this might be accomplished by verifying the contents of the update in RAM and  
950 then performing an update to the active flash. In another example, it could also be  
951 accomplished by loading the update into a region of flash, verifying it, and then selecting  
952 that region of flash as the active region.

#### 953 4.2.1.2 Integrity Protection

954 To prevent unintended or malicious modification of the firmware, nonvolatile storage regions  
955 containing device firmware need to be protected from such modifications outside of an  
956 authorized update mechanism.

- 957 1) The flash regions that contain device firmware **shall** be protected so that it is modifiable  
958 only through an authenticated update mechanism or by a secure local update mechanism  
959 that ensures the authenticity and integrity of the firmware update image by requiring that  
960 an authorized user physically touch the system itself to conduct the update.

961 *Note: To ensure integrity protections cannot be bypassed, integrity protections*  
962 *must either always be enabled, or must be engaged prior to execution of code*  
963 *outside of the CTU. Hardware integrity mechanisms may provide higher*  
964 *assurance than software or firmware-based mechanisms. These integrity*  
965 *protection mechanisms must ensure that firmware can only be modified as part of*  
966 *authenticated update, or a secure local update.*

#### 967 4.2.1.3 Non-Bypassability

968 The principle of non-bypassability is that it should not be possible for an attacker to modify  
969 device firmware outside of the authenticated update mechanism or, if supported, a secure local  
970 update. Any mechanisms capable of bypassing the authenticated update mechanism could create  
971 a vulnerability allowing malicious software to modify device firmware with a malicious or  
972 invalid image.

- 973 1) The protection mechanisms **shall** ensure that authenticated update mechanisms are not  
974 bypassed.
- 975 2) The authenticated update mechanism **shall** be capable of preventing unauthorized updates  
976 of the device firmware to an earlier authentic version that has a security weakness or  
977 would enable updates to a version with a known security weakness.

978 *Note: Updates to earlier firmware versions, sometimes called “rollback,” may*  
979 *provide a means to recover from a firmware update that is not functioning*



980                    *correctly. However, unauthorized rollback could allow an attacker to restore a*  
 981                    *vulnerable firmware image, which in turn could allow the attacker to damage the*  
 982                    *device or inject malware. As such, devices that support rollback should include*  
 983                    *appropriate security controls to ensure it cannot be exploited by an unauthorized*  
 984                    *entity in an attack.*

#### 985    **4.2.2 Protection of Immutable Code**

986    Code could be stored in field non-upgradable memory, such as Read Only Memory (ROM). While  
 987    the protections for this type of storage are strong, the trade-off is the inability to update the code  
 988    to fix bugs and patch vulnerabilities. Manufacturers of systems and devices should carefully weigh  
 989    the advantages and disadvantages of using nonvolatile storage that is not field upgradable.

990            1) If used, the write protection of field non-upgradable memory **shall not** be modifiable.

#### 991    **4.2.3 Runtime Protection of Critical Platform Firmware**

992    To satisfy the principle of non-bypassability described in Section 4.2.1.3, it is important that  
 993    software or bus-mastering hardware under the control of Software not be capable of interfering  
 994    with the intended function of *Critical Platform Firmware*. Critical Platform Firmware is the  
 995    collection of all Platform Firmware that either (a) performs the functions of protection, detection,  
 996    recovery and update of any Platform Firmware, (b) maintains the security of Critical Data or (c)  
 997    implements interfaces for Critical Data that are non-bypassable.

998    Devices that claim conformance with the Protection requirements, and rely on critical platform  
 999    firmware to protect the firmware image and/or critical data at OS runtime, must meet the  
 1000    guidelines in this subsection. The goal of these guidelines is to establish an environment for  
 1001    critical platform firmware to execute in which is isolated (protected) from software. Such  
 1002    isolation (protection) may be provided either logically (e.g., use of System Management Mode in  
 1003    x86-based platforms, or TrustZone in ARM-based platforms), or physically (e.g. in RAM  
 1004    attached to a non-host processor which is physically or logically isolated from the host  
 1005    processor).

1006    This subsection does not necessarily apply to firmware that is classified as non-critical (for  
 1007    instance, the majority of the BIOS on a PC-style platform is typically non-critical).

1008            1) If Critical Platform Firmware code in non-volatile memory is copied into RAM to be  
 1009            executed (for performance, or for other reasons) then the firmware program in RAM  
 1010            **shall** be protected from modification by Software or complete its function before  
 1011            Software starts.

1012            2) If Critical Platform Firmware uses RAM for temporary data storage, then this memory  
 1013            **shall** be protected from Software running on the Platform or until the data's use is  
 1014            complete.

1015            3) Software **shall not** be able to interfere with the functioning of Critical Platform  
 1016            Firmware. For example, by denying execution, modifying the processor mode, or  
 1017            polluting caches.

1018 *Note: This guideline does not preclude the use of RAM that is writable by*  
1019 *Software specifically for communication with Firmware or device hardware,*  
1020 *including using memory as a staging area for updates. The requirement is*  
1021 *intended to prevent the unauthorized modification of executing code or private*  
1022 *state used by the Critical Platform Firmware.*

#### 1023 4.2.4 Protection of Critical Data

1024 Unauthorized changes to critical data stored and used by devices could also seriously impact the  
1025 security posture of a device. Such changes could modify or disable important security-relevant  
1026 functions provided by the platform, or prevent the device from functioning at all. While critical  
1027 data may need to be modifiable by operating systems and other components, the guidelines in  
1028 this section aim to provide a controlled interface for these changes and guard against changes  
1029 that would put the device in an invalid state.

- 1030 1) Critical data **shall** be modifiable only through the device itself and/or defined interfaces  
1031 provided by device firmware. Examples of defined interfaces include proprietary or  
1032 public application programming interfaces (API's) used by the device's firmware, or  
1033 standards-based interfaces. Symbiont devices may rely on their host devices to meet this  
1034 requirement.
- 1035 2) Critical data updates **shall** be validated either by the device or a symbiont's host device  
1036 prior to committing changes to critical data to ensure that the new data is well-formed.  
1037 Examples of validation can include range or bounds checking, format checking, etc.
- 1038 3) Critical data updates **shall** be authorized by a Platform Administrator.
- 1039 4) Critical data updates **may** employ mechanisms to authenticate the critical data before it is  
1040 used.
- 1041 5) The device **shall** protect its factory defaults at least as well as it protects its code. The  
1042 factory defaults **shall** be able to be updated in the same manner as the code.

#### 1043 4.3 Detection

1044 The detection guidelines in this section describe mechanisms which can detect authorized  
1045 changes to device firmware and critical data before it is executed or consumed by the device.  
1046 When unauthorized changes are detected, a device could initiate a recovery process, as described  
1047 in Section 4.4. Detection mechanisms are particularly important for devices that lack strong  
1048 protections on their firmware or critical data. However, these mechanisms can also provide a  
1049 means to detect failures in firmware or code protection for devices that attempt to implement the  
1050 guidelines in Section 4.2.

1051 All devices which provide detection of corruption of their firmware code and critical data must  
1052 meet the guidelines which follow.

##### 1053 4.3.1 Detection of Corrupted Code

1054 Execution of unauthorized or corrupted firmware on a device could damage the device, inject  
1055 malware in the system, or otherwise impact the security functions and capacities of a device or

1056 encompassing system. The following guidelines describe mechanisms to verify the integrity of  
1057 firmware during the boot process using the Root of Trust for Detection (RTD), specified in  
1058 Section 4.1.3. While cryptographic integrity checks are preferred, either by the device itself or a  
1059 host device, some hardware devices, such as FPGAs or CPLDs, may use other mechanisms to  
1060 detect corruption in their code and programmable logic.

1061 For these detection mechanisms to be effective, the design of the device needs to ensure that the  
1062 RTD remains trustworthy in the event of a successful attack on the firmware itself.

- 1063 1) A successful attack which corrupts the active critical data or the firmware image, or  
1064 subverts their protection mechanisms, **shall not** in and of itself result in a successful  
1065 attack on the RTD or the information necessary to detect corruption of the firmware  
1066 image.
- 1067 2) One or more of the following techniques **shall** be used by the RTD or CTD to validate  
1068 firmware code:
- 1069 a) Integrity verification, using an approved digital signature algorithm or cryptographic  
1070 hash, of device firmware code prior to execution of code outside the RTD.

1071 *Note: Integrity verification **may** also be performed at runtime. These mechanisms*  
1072 *may or may not be anchored in the RTD.*

- 1073 b) Symbiont devices **may** rely on a host device to perform detection. If the symbiont  
1074 device boots independently from the host device, integrity verification of the  
1075 symbiont's device firmware **shall** be performed prior to execution of code outside the  
1076 host CTD. In such cases, the following additional requirements apply:
- 1077 i) The symbiont's firmware **shall** be protected according to the requirements in  
1078 Section 8.2.1.
- 1079 ii) The host device **should** be capable of immediately triggering a recovery of the  
1080 symbiont's firmware, followed by a restart of the device, in cases where  
1081 corruption is detected.
- 1082 c) Certain hardware devices (e.g., FPGAs, CPLDs) may have field-upgradable logic  
1083 rather than firmware code, often referred to as configuration bitstream. If these  
1084 devices do not have the ability to support cryptographic verification or the ability to  
1085 measure and report in conformance with a) or b), they **shall** use hardware-based  
1086 mechanisms to detect device load failures.
- 1087 d) Other techniques (e.g., watchdog timers) **may** be used in combination with  
1088 cryptographic integrity checks to detect other problems with the initialization process  
1089 of platform devices.
- 1090 3) If corruption is detected, the RTD or CTD **shall** be capable of starting a recovery process  
1091 to restore the device firmware code back to an authentic version.
- 1092 4) The detection mechanism **should** be capable of creating notifications of corruption.
- 1093 5) The detection mechanism **should** be capable of logging events when corruption is  
1094 detected

- 1095 6) The detection mechanisms **may** be capable of using policies set by the Platform  
 1096 Administrator which define the actions taken by the RTD/CTD in the above guidelines.

### 1097 **4.3.2 Detection of Corrupted Critical Data**

1098 This section describes mechanisms to detect invalid or corrupted critical data in platform  
 1099 devices. As noted above, invalid critical data could render a device inoperable or disable certain  
 1100 critical security functionality. Verifying critical data is challenging, because data is often  
 1101 intended to be user-configurable. The guidelines in this section recommend either directly  
 1102 verifying critical data contents or implementing other mechanisms to look for symptoms of data  
 1103 corruption.

- 1104 1) The RTD or CTD **shall** perform integrity checks on the critical data prior to use. Integrity  
 1105 checks may take the form, for example, of validating the data against known valid values  
 1106 or verifying the hash of the data storage.
- 1107 2) Either as an alternative to integrity checks (for devices that cannot support such a  
 1108 capability) or in addition to those checks, the RTD or CTD **may** use watchdog timers to  
 1109 detect potential corruption of critical data.
- 1110 3) If corruption of critical data is detected, the RTD or CTD **shall** be capable of starting a  
 1111 recovery process to restore the device's critical data.
- 1112 4) The detection mechanism **should** be capable of logging events when corruption is  
 1113 detected
- 1114 5) The RTD or CTD **should** be capable of creating notifications of corruption.
- 1115 6) The RTD or CTD **may** be capable of forwarding notifications of corruption.

## 1116 **4.4 Recovery**

1117 This section describes mechanisms for restoring platform firmware and critical data to a valid  
 1118 and authorized state in the event that any such firmware or critical data are detected to have been  
 1119 corrupted, or when an administrator initiates a manual recovery process.

### 1120 **4.4.1 Recovery of Mutable Code**

1121 The firmware recovery guidelines in this section specify mechanisms to recover firmware to a  
 1122 locally-stored backup or to a recovery image downloaded from another source. In either case,  
 1123 these guidelines specify using an Authenticated Update Mechanism (Section 4.2.1.1) to verify  
 1124 the integrity and authenticity of the image prior to recovery.

- 1125 1) A successful attack which corrupts the active critical data or the primary firmware image,  
 1126 or subverts their protection mechanisms, **shall not** in and of itself render the device's  
 1127 firmware code unrecoverable.
- 1128 a) The RTRec, CTRec and authentic recovery firmware image **should** be protected  
 1129 independently of the running firmware.
- 1130 2) The RTRec or CTRec **shall** be capable of obtaining an authentic device firmware image.

- 1131 a) If the authentic firmware image is stored locally in non-volatile memory, the image  
1132 **shall** be protected from unauthorized modification.
- 1133 3) Updates to a locally stored authentic firmware image **shall** be by way of an Authenticated  
1134 Update Mechanism (Section 4.2.1.1) or a secure local update (Section 3.5.3).
- 1135 4) Non-local recovery mechanisms **shall** use an Authenticated Update Mechanism (Section  
1136 4.2.1.1) to verify the integrity and authenticity of recovery images prior to restoring them.
- 1137 5) If the authentic firmware image is stored remotely, the RTRec or CTRec **shall** be  
1138 configurable with the location of this image or via authorized administrator's directed  
1139 locations as established by recovery policies.
- 1140 6) If the device (a *symbiont device*) relies upon another platform device (a *host device*) to  
1141 provide its RTRec or CTRec, then the host device's RTRec or CTRec **shall** invoke the  
1142 host/symbiont Authenticated Update Mechanism during recovery operations.
- 1143 7) The device **shall** either implement its own recovery capability, or that device (a *symbiont*  
1144 *device*) and another platform device (a *host device*) **shall** together implement the  
1145 symbiont device's recovery capability.
- 1146 8) The recovery mechanism **should** be capable of logging and reporting events when  
1147 recovery is performed.
- 1148 9) The recovery mechanism **should** be capable of providing notifications of recovery events  
1149 and actions.
- 1150 10) The recovery mechanism **may** be capable of performing the recovery action without  
1151 notification or intervention by the user or system administrator.
- 1152 11) The recovery mechanism **may** request approval from the user or system administrator to  
1153 perform a recovery action.
- 1154 12) The platform administrator **should** be able to initiate recovery of mutable code. Devices  
1155 **should** provide a method for Platform Administrators to force recovery. Devices **may**  
1156 receive platform-level authorization to force recovery through a chain of one or more  
1157 trusted devices, the first of which **shall** verify platform-level authorization prior to  
1158 instructing downstream devices to recover.
- 1159 13) The recovery process **should** protect against unauthorized recovery to an earlier firmware  
1160 version that contains a security weakness. The overall recovery process **should** facilitate  
1161 recovery to a recent firmware version. These **may** be implemented as a multi-stage  
1162 recovery process.

#### 1163 4.4.2 Recovery of Critical Data

1164 This section describes mechanisms to recover critical data in platform device in the event that the  
1165 device or administrator has reason to believe it has been corrupted. Because critical data can be  
1166 user-configurable, recovery requires the availability of trusted, known-good backup copies of  
1167 critical data. These backup copies may be stored on the device itself or by some other host  
1168 device. Because these backups may also be vulnerable to attack, these guidelines also specify  
1169 that devices also provide a means to restore to known-good factory defaults.

- 1170 1) A successful attack which corrupts the active critical data or the Primary Firmware  
 1171 Image, or subverts their protection mechanisms, **shall not** in and of itself render the  
 1172 device or its critical data unrecoverable.
- 1173 2) The device **shall** provide a method to backup a known good copy (or copies) of the active  
 1174 critical data to another location or locations. The protections on those locations **shall** be  
 1175 at least as good as that for the active critical data. The protections **should** be better than  
 1176 those for the active critical data.
- 1177 a) A symbiont device that cannot backup its own critical data **shall** make its critical data  
 1178 available to its host device. In this case, the host device shall backup the symbiont's  
 1179 data.
- 1180 b) If a symbiont provides its critical data to a host device so the host device may backup  
 1181 the data, then the symbiont **shall** be able to consume the recovered critical data.
- 1182 3) The device **may** determine that its critical data is "known good" by using that data as part  
 1183 of a successful reboot.
- 1184 4) The device **shall** backup the critical data either automatically or when instructed by the  
 1185 user or when instructed to do so by another trusted device.
- 1186 5) The RTRec or CTRec **shall** be capable of recovering critical data to factory defaults.
- 1187 6) The RTRec or CTRec **should** be capable of recovering to last known good critical data.
- 1188 7) A device **shall not** use policies stored as critical data by that device to recover its own  
 1189 critical data. However, a symbiont **may** rely on policies which are provided by a host  
 1190 device.
- 1191 8) If multiple back-ups are available, the RTRec or CTRec **may** allow a choice of which  
 1192 back up to use.
- 1193 9) If detection of corruption of critical data is automatic, the RTRec or CTRec **may** gain  
 1194 approval from a host device or the user before replacing the current critical data.
- 1195 10) In the absence of the RTD or CTD triggering recovery actions, the platform administrator  
 1196 **should** be able to initiate recovery of critical data. Devices **should** provide a method for  
 1197 Platform Administrators to force recovery. Devices which receive authorized requests to  
 1198 force recovery **may** then instruct other devices with which there are established trust  
 1199 relationships to force recovery, either directly or through chains of trusted devices.

1200 While it is outside the scope of this document to define what constitutes an appropriate state for a  
 1201 platform to recover to (other than a state of integrity), examples include any of the following:

- 1202 • Recover to a last known good state
- 1203 • Reset to factory defaults
- 1204 • Update to the newest firmware image
- 1205 • Perform a partial 'repair' operation
- 1206 • Recover to an enterprise-defined 'starting point'
- 1207 • Any combination of the above

1208 Note that recovery processes may require multiple stages before normal operation is restored.  
1209 For example, devices may initially restore factory-default configuration data prior to recovering  
1210 last-known-good configuration data.

1211 When considering how to determine which of the above states of integrity to recover a device to,  
1212 using a policy-based approach necessitates the use of critical data in order to store/maintain those  
1213 policies. However, if that critical data becomes corrupted, then the recovery process may either  
1214 not be able to happen or it may happen incorrectly. As such, vendors should carefully consider  
1215 the algorithm used to recover. As an example, a straightforward mechanism might be the use of  
1216 a simple Most Recently Used (MRU) algorithm, e.g., the algorithm might first try to recover to  
1217 the last known good state; if that data is not valid, then it might next try recovering to a prior  
1218 saved state; if that is not available, then it might try recovering from a remote enterprise storage  
1219 location; if that is not available, then it might try resetting to factory defaults. Using an  
1220 algorithmic approach in this manner eliminates the need to rely on critical data being in a state of  
1221 integrity during the recovery operation.

1222 **Appendix A Acronyms**

1223 Selected acronyms and abbreviations used in this paper are defined below.

BIOS	Basic Input/Output System
CoT	Chain of Trust
CPLD	Complex Programmable Logic Device
CTD	Chain of Trust for Detection
CTRec	Chain of Trust for Recovery
CTU	Chain of Trust for Update
FPGA	Field-Programmable Gate Array
ROM	Read Only Memory
RoT	Root of Trust
RTD	Root of Trust for Detection
RTRec	Root of Trust for Recovery
RTU	Root of Trust for Update
UEFI	Unified Extensible Firmware Interface

1224



1225

**Appendix B Glossary**

<b>Active Critical Data</b>	The copy of critical data that is used to initialize or configure a device. <i>Note: active critical data does not include back-ups of critical data.</i>
<b>Add-in Card</b>	A generic term used to refer to any device which can be inserted or removed from a platform through a connection bus, such as PCI. Add-in cards are typically inserted within a platform's physical enclosure, rather than residing physically external to a platform. An add-in card will have its own devices and associated firmware, and may have its own Expansion ROM Firmware.
<b>Authenticated Update</b>	An update which uses an authenticated update mechanism.
<b>Authenticated Update Mechanism</b>	An update mechanism which ensures that an update firmware image has been digitally signed and that the digital signature can be verified using one of the keys in the key store of the Root of Trust for Update (RTU) before updating the firmware image.
<b>Authorized Update</b>	An update which uses an authorized update mechanism.
<b>Authorized Update Mechanism</b>	An update mechanism that checks for approval before installation. Approval could consist of checking possession of a credential, confirmation by a physically present person or similar means.
<b>Boot Firmware</b>	Generic term to describe any firmware on a platform executed to boot (start-up, initialize) a device. This may include, but is not limited to, initialization of device memory, initialization of device registers, initialization of connectivity interfaces, health checks of the device, etc. The general purpose of boot firmware is to prepare a device or platform for normal operational use.
<b>Chain of Trust (CoT)</b>	A Chain of Trust (CoT) is a sequence of cooperative elements which are anchored in a Root of Trust (RoT) that extend the trust boundary of the current element by conveying the same trust properties to the next element when it passes it control. The result is both elements are equally able to fulfill the trusted function as though they were a single trusted element. This process can be continued, further extending the chain of trust. Once control is passed to code which is not, or cannot be, verified then the Chain of Trust has ended. This is also referred to as passing control to a non-cooperative element.
<b>Code</b>	Instructions directly executed by a processor or like device (FPGA, CPLD, etc.). Also included are instructions that are interpreted by a program such as UEFI EBC and JVM. <i>Source code</i> is the human-readable instructions that are translated into code and then executed.

<b>Corruption</b>	Corruption is a loss of integrity of firmware code, or an error or unexpected value in critical data, which could be the result of any one of a number of different causes, including but not limited to, malicious activity (e.g. an attacker), poorly written code (e.g. buffer overflows, algorithmic error), accidental events (e.g. inadvertent user action), failure to install a security patch, unauthorized changes, or hardware-induced (e.g. signal integrity, alpha particles, power failures).
<b>Critical Data</b>	Critical data is mutable data which persists across power cycles and must be in a valid state that has been authorized by the Platform Administrator in order for the recovery and/or booting of the platform to securely and correctly proceed.
<b>Critical Platform Firmware</b>	The collection of all Platform Firmware that either (a) performs the functions of protection, detection, recovery and update of any Platform Firmware, (b) maintains the security of Critical Data or (c) implements interfaces for Critical Data that are non-bypassable.
<b>Device</b>	A generic term used to refer to any computing or storage element or collection of computing or storage elements on a platform. Examples of devices include a central processing unit (CPU), applications processing unit (APU), embedded controller (EC), baseboard management controller (BMC), Trusted Platform Module (TPM), graphics processing unit (GPU), network interface controller (NIC), hard disk drive (HDD), solid state drive (SSD), Read Only Memory (ROM), flash ROM, etc.
<b>Device Firmware</b>	The collection of non-host processor firmware and Expansion ROM firmware that is only used by a specific device. This firmware is typically provided by the device manufacturer.
<b>Expansion ROM Firmware</b>	Peripheral Component Interconnect (PCI) term for firmware executed on a host processor which is used by an add-in device during the boot process. This includes Option ROM Firmware, UEFI applications, and UEFI drivers. Expansion ROM Firmware may be bundled as part of the Host Processor Boot Firmware, or may be separate (e.g., from an add-in card). In this document, we will use the term Expansion ROM Firmware when referring to either Option ROM Firmware or UEFI Drivers and Applications generically.
<b>Firmware</b>	Generic term to describe any code stored in a chip that either resides at the reset vector (or equivalent) of the corresponding processor or which is provided as extensions to other firmware (such as Expansion ROM Firmware). General purpose operating systems that

are stored in chips are generally not considered firmware in the scope of this document.

**Host Device**

A device that assists a symbiotic device to establish the roots and chains of trust required to meet the protection, detection, and/or recovery guidelines in this document. The allocation of functionality between the host and symbiotic devices is implementation dependent. The Host Device must, itself, meet the guidelines for its own firmware on its own. Note that this should not be confused with a Host Processor. A Host Processor may serve as a Host Device, but a Host Device is not necessarily a Host Processor.

**Host Processor**

A host processor is the primary processing unit in a platform, traditionally called a Central Processing Unit (CPU), now also sometimes referred to as an Application Processing Unit (APU), or a System on Chip (SoC). This is the processing unit on which the primary operating system (and/or hypervisor), as well as user applications run. This is the processor that is responsible for loading and executing the Host Processor Boot Firmware.

**Host Processor Boot Firmware**

Generic term used to describe firmware loaded and executed by the Host Processor which provides basic boot capabilities for a platform. This class of firmware includes Legacy BIOS, System BIOS and UEFI, as well as other implementations. Where the distinction between Legacy BIOS and UEFI is not important, the term Host Processor Boot Firmware will be used. Where the distinction is important, it will be referenced accordingly. Expansion ROM firmware may also be considered as part of the Host Processor Boot Firmware. Expansion ROM Firmware may be embedded as part of the Host Processor Boot Firmware, or may be separate from the Host Processor Boot Firmware (e.g., loaded from an add-in card). Host Processor Boot Firmware includes firmware which may be available during runtime.

**Immutable**

Unchangeable. In the context of this document, this refers only to the inability to make changes in the field through manufacturer intended mechanisms and/or defined interfaces. Note that a platform or device manufacturer may still be able to make changes through manufacturing or service tools directly connected to a locally (physically) present platform or device.

**Legacy BIOS (Basic Input/Output System)**

One form of Host Processor Boot Firmware used on x86 platforms which uses a legacy x86 BIOS structure. This form of host processor boot firmware has been or is being replaced by UEFI.

<b>Mutable</b>	Changeable. In the context of this document, this refers only to the ability to make changes in the field through manufacturer intended mechanisms and/or defined interfaces. Such mechanisms may require cryptographic mechanisms or unambiguous physical presence.
<b>Non-critical Data</b>	Non-critical data is mutable data which persists across power cycles but is not critical to the booting, operating, or recovery of the device.
<b>Non-Host Processor</b>	A non-host processor is a generic term used to describe any processing unit on a platform which is not a host processor (e.g. a microcontroller, co-processor, etc.).
<b>Non-Host Processor Firmware</b>	Non-host processor firmware is a generic term used to describe firmware used by any processing unit on a platform which is not a host processor.
<b>Option ROM Firmware</b>	Legacy term for boot firmware typically executed on a host processor which is used by a device during the boot process. Option ROM firmware may be included with the host processor boot firmware or may be carried separately by device (such as an add-in card).
<b>Peripheral (aka external device)</b>	A peripheral (also known as an external device) is a device which resides physically external to a platform and is connected to a platform, either wired or wirelessly. A peripheral is comprised of its own devices which may have their own firmware. While conceptually the principles and guidelines in this document could apply equally to peripherals, they are outside of the scope of this document.
<b>Platform</b>	A platform is comprised of one or more devices assembled and working together to deliver a specific computing function, but does not include any other software other than the firmware as part of the devices in the platform. Examples of platforms include a notebook, a desktop, a server, a network switch, a blade, etc.
<b>Platform Administrator Privilege</b>	Privileges required to manage the firmware and critical data on a platform devices. In particular, this privilege may be needed to authorize firmware updates, change firmware configuration settings, and firmware recovery operations.  This privilege must be carefully controlled as compromise of Platform Administrator Privileges can reduce the dependability of a platform and its devices.

<b>Platform Administrator</b>	An entity with Platform Administrator Privileges.
<b>Platform Firmware</b>	The collection of all device firmware on a platform. In this document the term Platform Firmware may be used when making references to the collection of all firmware used by devices on a platform.
<b>Primary Firmware Image</b>	The executable code stored on the device. Different parts of the primary firmware image may be protected differently.
<b>Read Only Memory (ROM)</b>	A memory device that once it has been initially set, cannot be overwritten through any mechanism, making that memory immutable (unchangeable).
<b>Root of Trust (RoT)</b>	An element that forms the basis of providing one or more security-specific functions, such as measurement, storage, reporting, recovery, verification, update, etc. A RoT is trusted to always behave in the expected manner because its misbehavior cannot be detected and because its proper functioning is essential to providing its security-specific functions.
<b>Runtime Firmware</b>	Generic term to describe any firmware on a platform active/functional or available for use during runtime (after boot has completed).
<b>Software</b>	Software is comprised of elements required to load the operating system, the operating system, and all user applications and user data subsequently handled by the operating system. Refer to Figure 1 for a graphical depiction.
<b>Symbiont Device</b>	A symbiont device is a device that relies wholly or partially on another device (a host device) to establish the roots and chains of trust required to meet the protection, detection, and recovery guidelines in this document. The allocation of functionality between host device and symbiont device is implementation dependent. For example, the host device verify updates and back up critical data while the symbiont device is responsible for meeting all other guidelines. The symbiont property can be transitive: A device may be a symbiont device to a host device while the two may then serve as the host to another symbiont device and so on.
<b>System</b>	A system is the entirety of a computing entity, including all elements in a <i>platform</i> (hardware, firmware) and <i>software</i> (operating system, user applications, user data). A system can be thought of both as a logical construct (e.g. a software stack) or physical construct (e.g. a

notebook, a desktop, a server, a network switch, etc.). Refer to Figure 1 for a graphical depiction.

**UEFI (Unified Extensible Firmware Interface)**

One form of Host Processor Boot Firmware which uses a Unified Extensible Firmware Interface (UEFI) structure (as defined by the UEFI Forum).

**UEFI Drivers**

Standalone binary executables in PE/COFF format which are loaded during the boot process to handle specific pieces of hardware.

**Unambiguous Physical Presence**

Indicates authorization by a local person that cannot be impersonated by malware.

1226

1227

**Appendix C References**

- [1] David Cooper, et al., *BIOS Protection Guidelines*, NIST Special Publication (SP) 800-147, National Institute of Standards and Technology, Gaithersburg, Maryland, April 2011, 26pp. <https://doi.org/10.6028/NIST.SP.800-147>
- [2] Andrew Regenscheid., *BIOS Protection Guidelines for Servers*, NIST Special Publication (SP) 800-147B, National Institute of Standards and Technology, Gaithersburg, Maryland, August 2014, 26pp. <https://doi.org/10.6028/NIST.SP.800-147B>
- [3] Specifications, Unified Extensible Firmware Interface Forum [Web site], <http://www.uefi.org/specifications> [accessed 4/21/17]
- [4] TPM Main Specification, Trusted Computing Group [Web site], <https://trustedcomputinggroup.org/tpm-main-specification/> [accessed 4/21/17]
- [5] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, RFC 2119, Internet Engineering Task Force, March 1997, 2pp, <https://doi.org/10.17487/RFC2119>
- [6] U.S. Department of Commerce. *Secure Hash Standard*, Federal Information Processing Standards (FIPS) Publication 180-4, August 2015, 36pp. <https://doi.org/10.6028/NIST.FIPS.180-4>
- [7] U.S. Department of Commerce. *Digital Signature Standard*, Federal Information Processing Standards (FIPS) Publication 186-4, July 2013, 130pp. <https://doi.org/10.6028/NIST.FIPS.186-4>
- [8] Elaine Barker, *Recommendation for Key Management, Part 1: General*, NIST Special Publication (SP) 800-57 Part 1, Rev. 4, National Institute of Standards and Technology, Gaithersburg, Maryland, July 2016, 160pp. <https://doi.org/10.6028/NIST.SP.800-57pt1r4>
- [9] Elaine Barker, *Recommendation for Obtaining Assurance for Digital Signature Applications*, NIST Special Publication (SP) 800-89, National Institute of Standards and Technology, Gaithersburg, Maryland, November 2006, 38pp. <https://doi.org/10.6028/NIST.SP.800-89>

1228