ETSI GS NFV-SEC 012 V3.1.1 (2017-01)

Network Functions Virtualisation (NFV)
Release 3;
Security;
System architecture specification
for execution of sensitive NFV components

Disclaimer

The present document has been produced and approved by the Network Functions Virtualisation (NFV) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference
DGS/NFV-SEC012

Keywords
architecture, NFV, security

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Network Functions Virtualisation (NFV).

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document defines requirements for host system elements on which sensitive workloads are to be run. The present document defines requirements to ensure isolation of sensitive workloads from non-sensitive workloads sharing a platform. The present document discusses a wide range of different technologies which aim to increase the security of a host system for the workloads which will be executing on it.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

[1] ETSI TS 133 310: "Universal Mobile Telecommunications System (UMTS); LTE; Network Domain Security (NDS); Authentication Framework (AF) (3GPP TS 33.310)".

[2] ETSI TS 133 210: "Digital cellular telecommunications system (Phase 2+) (GSM); Universal Mobile Telecommunications System (UMTS); LTE; 3G security; Network Domain Security (NDS); IP network layer security (3GPP TS 33.210)".

[3] ISO/IEC 18031:2001: "Information technology -- Security techniques -- Random bit generation or equivalent specification".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] NIST Publication (SP) 800-90B: "Recommendation for the Entropy Sources Used for Random Bit Generation".

[i.2] NIST Publication (SP) 800-88 revision 1: "Guidelines for Media Sanitization".

[i.3] ETSI GS NFV-SEC 009: "Network Functions Virtualisation (NFV); NFV Security; Report on use cases and technical approaches for multi-layer host administration".

[i.4] Greg Hoglund, Gary McGraw (2007): "Exploiting Online Games: Cheating Massively Distributed Systems", Addison-Wesley, New Jersey.

[i.5] ETSI TS 103 487 "CYBER; Baseline security requirements regarding sensitive functions for NFV and related platforms".

[i.6] ETSI TR 103 309: "CYBER; Secure by Default - platform security technology".

[i.7] NIST SP800-123: "Guide to General Server Security".

[i.8] NIST SP800-125: "Guide to Security for Full Virtualization Technologies".

[i.9] ISO/IEC 15408: "Information technology -- Security techniques -- Evaluation criteria for IT security -- Part 1: Introduction and general model".

[i.10] ETSI GR NFV-SEC 003: "Network Functions Virtualisation (NFV); NFV Security; Security and Trust Guidance".

[i.11] ETSI GS NFV-INF 004 (V1.1.1): "Network Functions Virtualisation (NFV); Infrastructure; Hypervisor Domain".

[i.12] TCG: "Virtualized Trusted Platform Architecture Specification", Version 1.0, Revision 0.26.

[i.13] NIST SP 800-162: "Guide to Attribute Based Access Control (ABAC) Definition and Considerations".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

host system: collection of hardware, software and firmware making up the system which executes workloads

NOTE 1: When the host system is part of the NFVI, it is the "hypervisor" and "host" as defined by ETSI NFV-INF 004 (V1.1.1) [i.11]. In the case of virtualisation of workloads within the MANO domain, there is no corresponding definition available.

NOTE 2: The definition in ETSI NFV-INF 004 [i.11] specifically excludes containers, but the present document does not.

workload: component of the NFV architecture that is virtualised in the context of a particular deployment

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ABAC	Attribute-Based Access Control
DH	Diffie-Hellman
DHE	Diffie-Hellman Exchange
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie-Hellman
ECDHE	Elliptic Curve Diffie-Hellman Exchange
ECDSA	Elliptic Curve Digital Signature Algorithm
ECP	Elliptic Curve modulo a Prime
GMAC	Galois Message Authentication Mode
HBRT	Hardware-Based Root of Trust
HMEE	Hardware-Mediated Execution Enclave
ICV	Integrity Check Value
HSM	Hardware Security Module
IOMMU	Input-Output Memory Management Unit
MANO	MANagement and Orchestration
MODP	More mODular exPonential
NIST	National Institute of Standards and Technology
PKI	Public Key Infrastructure
PRF	Pseudo-Random Function
RNG	Random Number Generator
RSA	Rivest-Shamir-Adleman

SSL	Secure Sockets Layer
TCB	Trusted Computing Base
TCG	Trusted Platform Group
TLS	Transport Layer Security
TPM	Trusted Platform Module

4 Principles

4.1 Introduction

Trust, as defined in ETSI GR NFV-SEC 003 [i.10], is an important component of security. One weakness of software as opposed to hardware, is that software can be copied in whole or in part. Trust that is rooted in software may be less reliable than trust rooted in hardware, quickly, easily, and any number of times. For the particular case of sensitive workloads that have to be trusted, only the highest assurance in the root of trust is considered acceptable, thus for the purposes of the present document the root of trust shall be provided in hardware.

There is, however, a concomitant concern that when a device is subject to black box testing, it is impossible to determine if the responses to interrogation come from hardware or software. To counter this, a NFVI vendor shall be able to provide evidence on demand that the root of trust is a hardware element. The means by which the vendor provides such evidence is not considered in the present document but should be mutually agreed between the vendor and operator.

A vendor shall be able to provide evidence on demand to authorized parties of the security claims for the root of trust. The means by which the vendor provides such evidence is not considered in the present document, but should be mutually agreed between the vendor and operator. An examples of 3rd a party assurance programme is Common Criteria (defined in ISO/IEC 15408 [i.9]).

The host system, acting as a black box (closed) environment, shall provide access to authorized external entities only to those capabilities identified in the authorization agreement.

5 Platform requirements

5.1 Core hardware requirements

1) The host system shall implement a Hardware-Based Root of Trust (HBRT) as Initial Root of Trust with the following requirements:

   - The HBRT shall be both physically and electronically tamper-resistant.
   - The HBRT shall be both physically and electronically tamper-evident.
   - The HBRT physical and software interfaces between the HBRT and other hardware components of the host system to which it directly communicates shall be protected from eavesdropping, manipulation, replay or similar attacks.
   - The level of resistance against attacks of the HBRT shall be verifiable and trustable using a certification process.
   - It shall be possible to restrict the booting procedure if assistance from the HBRT is not available or the HBRT currently does not contain valid cryptographic material.
   - Any tampering to the HBRT should lead to detectable degradation of its function.
   - The HBRT shall be physically protected such that any attempts to remove or replace the HBRT shall cause physical damage to both the HBRT and host system hardware to which the HBRT is attached, rendering both inoperable.

- The HBRT shall be (physically and/or logically) bound to the host system, so that any attempt to remove the HBRT will be detected and prevent normal operation of the host system.
- The HBRT shall include an Immutable Unique Identification value physically linked to the physical root of trust that can be used as identification of the platform. This value shall be stored in a shielded location protected from unauthorized use and disclosure.
- The HBRT shall provide capabilities to allow itself to be part of an attestation function.
- The host system shall have a mechanism to discover the tampered/non-tampered status of the HBRT.
- The host system shall have an interface to provide authorized external services with information about the tampered/non-tampered status of the HBRT.
- The host system shall provide a mechanism to report to authorized external services when tamper events occur.
- The HBRT shall implement a key management function with the requirements in the following bullet 2.

2) The host system shall implement a key management system which includes key generation, key storage, key deletion and cryptographic processing with the following requirements:

- The cryptographic material shall be stored in a shielded location, protected against eavesdropping and physical and environmental tampering.
- The key generation processing shall be protected against eavesdropping and physical and environmental tampering.
- The key management system shall include an access right management to the sensitive data.
- The key management system shall ensure a complete deletion of outdated keys under deletion request.
- The key management system shall be scalable and ensure a high availability service.
- The key management system shall be remotely manageable to allow evolution, security strengthening, and countermeasure deployment of the system.

The host system shall provide cryptographically separated secure environments to different applications.

5.2 Core software requirements

The following core software requirements are defined within the present document:

- Secure logging
- OS-level access control
- Logical authentication controls
- Communications security (e.g. Confidentiality, Integrity, Availability, Non-repudiation)
- Secure firmware (e.g. BIOS) upgrade
- Secure remote management of keys, cryptographic algorithms and security services offered by the platform to ensure ability of evolution, security strengthening, and countermeasure deployment

It shall be possible to restrict the booting procedure by preventing the running of workloads if assistance from the HBRT is not available or the HBRT currently does not contain valid cryptographic material. The intent of this requirement is to stop VNFs/VNFCIs being loaded onto possibly compromised hardware and to allow appropriate mitigations to be put in place.

6 Lifecycle

6.1 Trusted Computing Base

The Trusted Computing Base (TCB) comprises those components of the system - hardware, software and processes - that need to be trusted by default: it is on this foundation that the host system operates and on which the workload can operate with defined levels of trust in the overall security of the system. An example of a simplified boot scheme diagram for a TCB which utilizes at TPM as its Hardware-Based Root of Trust (HBRT) is provided in figure 1.

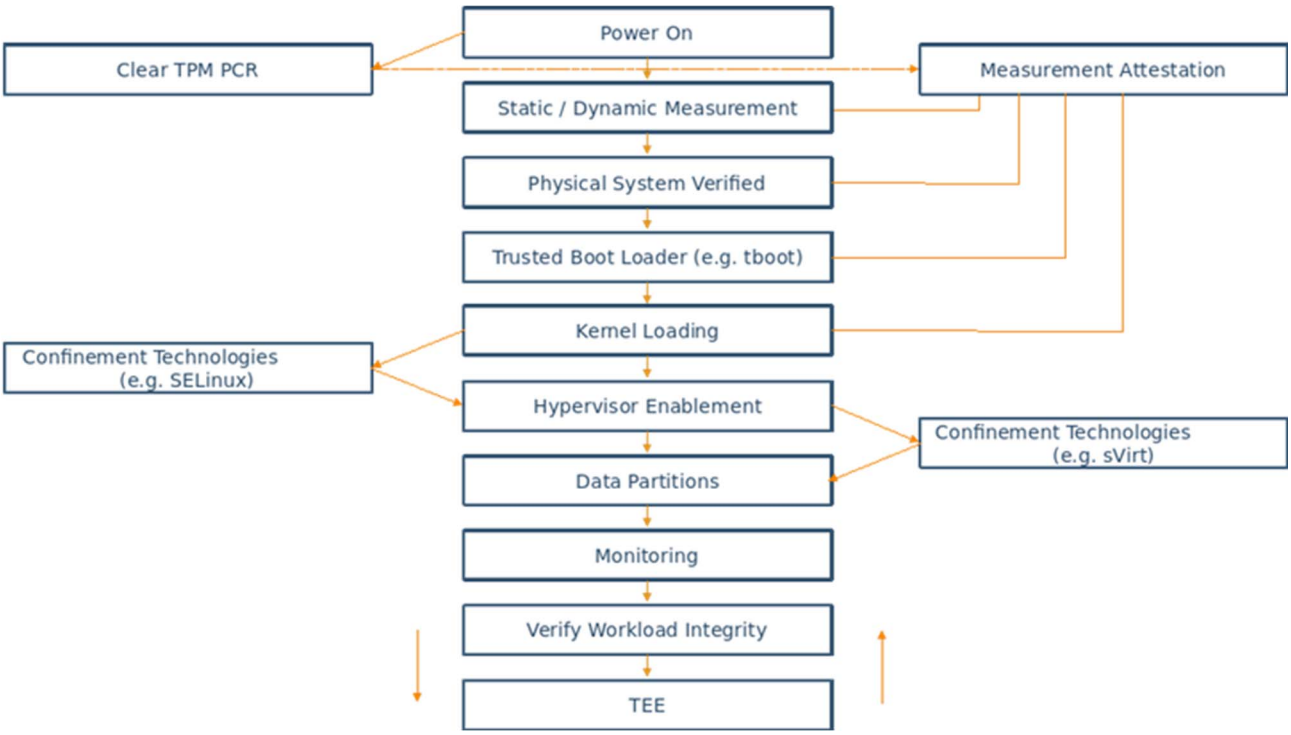

Figure 1: Example of a simplified boot scheme diagram using a Trusted Platform Module (TPM)

The detailed steps for building a TCB will be NFVI-vendor dependent, and are beyond the scope of the present document. One example of detailed guidance is Virtualised Trusted Platform Architecture Specification, Version 1.0, Revision 26 (TCG) [i.12].

The host system shall support the use of a service providing remote attestation.

Although the scope of the present document does not allow for requirements to be imposed on systems external to the host system, the attestation server should be implemented as a "bare-metal" deployment, rather than as a virtualised workload. This is because the attestation server needs to serve as one of the fundamental roots of trust of the MANO domain, and from there to the NFVI domain.

The measures discussed in the present document provide various protections for the host system and the workloads which execute on it. Vulnerabilities may exist which allow attackers using a compromised workload to "break out" to its host system. While the measures in clause 8.2, when correctly implemented, can mitigate against such compromises, a serious compromise of a host system may have implications beyond that single host system. This is especially true where explicit or implicit trust relationships exist between host systems in, for example, virtualised computing clusters. Although out of the scope of the present document, it is important that, when considering a deployment, the implications of explicit and implicit trust relationships are considered.

6.2 Workload provisioning

The host system shall have an interface to provide authorized external services with information about its ability to prohibit host or hypervisor memory deduplication techniques that allow for sharing of memory pages between workloads.

The host system shall provide a mechanism to disable host or hypervisor memory deduplication techniques that allow for sharing of memory pages between workloads.

The host system shall disable host or hypervisor memory deduplication techniques that allow for sharing of memory pages between workloads by default. Where a capability is disabled, it shall not be possible re-enable it without a host system reboot.

The host system shall allow appropriately authorized parties to specify that certain memory types or locations (e.g. volatile vs non-volatile, on-blade vs off-blade) are not used for particular workloads.

The host system should allow appropriately authorized parties to specify that only certain memory types or locations (e.g. volatile vs non-volatile, on-blade vs off-blade) are used for particular workloads

The host system shall provide the ability to prohibit local caching of binary images for workloads.

The host system shall have an interface to provide authorized external services with information about its ability to prohibit binary image caching.

The host system shall have an interface to provide authorized external services with information about its ability to provide perform secure provisioning of workloads.

The host system shall have an interface to provide authorized external services with information about its ability to provide secure de-provisioning of workloads.

The host system shall have an interface to provide authorized external services with information about its ability to block migration of workloads.

The host system shall provide secure provisioning of workloads.

The host system shall provide secure de-provisioning of workloads.

6.3 Runtime checks

One of the most challenging issues around the integrity of a "non-sealed" system is the making assertions around the integrity of a running system beyond its initial stages of provisioning, boot and software-loading. Techniques that can be employed on the host system include:

- Integrity checking of running processes by local agents.
- Periodic checking of executable and binary file integrity by local agents.

Agents, being software themselves, are also vulnerable to compromise [i.4]. What is more, both of these approaches are vulnerable to privilege escalation by non-authorized entities which may hide activities or reduce the privileges of agents, and there may be some processes whose operation should not be visible to agents of particular types, so use of these techniques requires a clear definition - and enforcement of - trust domains within the host. Note that this may be enforced by such mechanisms as Hardware-Mediated Execution Enclaves.

External techniques include behavioural monitoring of both host systems and workloads. These techniques may be combined - where trust models permit, to allow correlation of information. External techniques are outside the scope of the present document.

6.4 Entropy and random numbers

The host system should meet a minimum entropy standard such as that described in NIST SP800-90b [i.1].

The host system shall provide a means by which the designed availability (quality and available bandwidth) of entropy on the system can be queried by an authorized party.

The host system should provide a means by which the actual availability (quality and/or available bandwidth) of entropy on the system can queried by an authorized party.

The host system shall implement a Random Number Generator. The RNG function shall meet or exceed the quality performances specified in ISO/IEC 18031:2011 [3].

6.5 Cryptographic primitives

The ETSI-NFV platform is intended to be the basis for network services such as 3GPP services. The cryptographic primitives used by a 3GPP platform as described in ETSI TS 133 310 [1] and ETSI TS 133 210 [2] shall then be supported.

Some of the algorithms currently in common use are not considered safe for extended use due to projected improvements in computing power and techniques (e.g. RSA, DSA, DH, ECDH and ECDSA) but are broadly deployed and need to be supported for legacy reasons until new algorithms are available and adopted.

The host system shall provide access to the following cryptographic primitives:

- Hashing algorithms:
  - SHA-256
  - SHA-384
  - AES128-GMAC
  - HMAC-SHA128
  - HMAC-SHA256
  - HMAC-SHA384
  - MD5, MD2, SHA1, shall not be used
- Encryption algorithms:
  - AES-CBC-128
  - AES-GCM-128 (16 octet ICV)
  - AES-CBC-256
  - AES-GCM-256 (16 octet ICV)
- Signature
  - RSA 2048
  - RSA 3072
  - RSA 4096
  - ECDSA-256 (secp256r1)
  - ECDSA-384 (secp384r1)
- Public Key Infrastructure (PKI)
  - RSA 2048,
  - RSA 3072,
  - RSA 4096
  - id-ecPublicKey (secp256r1)
- Key Exchange
  - DH group 14 (2048-bit MODP)
  - DH group 19 (256-bit random ECP group)
  - DH group 20 (384-bit random ECP group)

  - ECDHE secp256r1 (P-256)
  - DHE groups of at least 2 048 bits
- Pseudo Random Function (PRF)
  - PRF-HMAC-SHA2-256
  - PRF-HMAC-SHA2-384

To increase the resistance to emerging and anticipated threats, it is highly recommended that the host system supports largest key length within an algorithm family.

- The host system shall not make use of MD5, MD2, SHA1 and these algorithms shall not be present on the host system in software or firmware.
- The host system shall not make use of PRF-HMAC-SHA1 and this algorithm shall not be present on the host system in software or firmware.
- The host system shall not provide capabilities to allow workloads to make use of MD5, MD2, SHA1.
- The host system shall not provide capabilities to allow workloads to make use of PRF-HMAC-SHA1.
- The host system MAY use hardware acceleration for one or more of the algorithms listed above for the purposes of NFVI-level management.
- The host system MAY provide hardware accelerated access to one or more of the algorithms listed above to workloads.

To assist longevity of the NFV platform the cryptographic algorithms should be implemented with a crypto agility requirement.

6.6 Installed software and configurations on host system

While outside the scope of the present document, it is important that a record be maintained of all installed and/or running software, configurations and versions, which should include debug status of any component of the host system.

While outside the scope of the present document, it is important that a time-stamped, confidentiality-protected and integrity-protected record and history of changed/updated software with reasons for changes be maintained, which should be protected from unauthorized access.

While outside the scope of the present document, it is important that a record be maintained with a list of authorized users and accounts for each host.

The host shall provide a means by which current versions of software and configurations can be queried by an authorized party.

6.7 De-provisioning workloads

See clause 4.4.7.2 "Secured wipe" of ETSI NFV-SEC 003 [i.10] for informative discussion of this topic.

The host system shall provide the capability to perform secure wipe of storage at the request of authorized external services.

The host system shall provide secure wipe capabilities meeting at least those specified in NIST SP800-88r1 [i.2].

The host system shall provide a mechanism by which authorized external services can confirm completion of the secure wipe operation.

The host system shall provide a mechanism to ensure that storage which is in the process of a secure wipe cannot be re-allocated until that operation is successfully completed.

The host system shall provide a mechanism to perform a secure wipe, at the time of de-provisioning, of any and all files associated with a workload.

6.8 Dealing with failure

6.8.0 General points

The host system shall be booted with debug options off by default.

The host system may provide a capability that makes it impossible to turn on debug options for the host system. When this capability is provided, the host system shall have an interface to provide authorized external services with information about this capability.

The host system shall make a record when debug options are turned on, and a specific log entry shall be created. This record shall be unalterable without a power off or reboot of the host system.

The host system shall have an interface to provide authorized external services with information about the state of its debug options, including historical state since boot.

The host system shall provide a mechanism to report to authorized external services when a change in debug status occurs.

6.8.1 Requirements relating to failure conditions

The host system shall have an interface to provide authorized external services with information related to failures or replacement of its components.

The host system shall provide a mechanism to report to authorized external services when failures occur.

7 External dependencies

The present document examines the system host, rather than supporting systems such as those in the MANO domain (though it should be noted that some MANO workloads may themselves be sensitive and therefore have requirements applied to them from the present document). Other systems that may be required, depending on what security profile is applied, include:

- Attestation authority.
- Certificate Authority.
- Remote logging server.

8 Architecture section

8.0 System hardening techniques

The following clauses address specific techniques to ensure an appropriate security posture for the host system and also those technologies to be made available (optionally, in some cases) to workloads. The intention is not to provide an exhaustive list of measures but to supplement best practice in the industry.

NOTE: Many of the technologies discussed in this clause are defined in ETSI GS NFV-SEC 009 [i.3].

Other useful references include:

- ETSI TS 103 487 [i.5].
- ETSI TR 103 309 [i.6].
- NIST SP800-123 [i.7].
- NIST SP800-125 [i.8].

8.1 Secure logging

The host system shall log system events to a remote location.

The host system shall encrypt log entries during transmission to the remote location.

The host system shall encrypt locally stored log entries.

The host system shall provide integrity protection for the list comprising locally stored log entries.

The host system shall preserve the time order of locally stored log entries. The host system shall have an interface to provide authorized external services with access to log entries and associated proofs of integrity.

The host system shall provide a mechanism to report to authorized external services when logging errors occur.

The host system shall monitor communications with the remote logging location and shall write a record in an internal error log entry if communications fail. The host system shall continue to log errors locally and shall attempt to send these to the remote location if communications are re-established.

It shall not be possible to change logging level(s) of any component on the host system without proper authorization.

The host system shall provide measures to reduce the possibility of inference attacks from traffic generated by logging and management events.

8.2 OS-level access and confinement control

The host system shall have OS-level access and confinement controls (e.g. SELinux, sVirt) configured to enforce capability controls on system processes.

The host system shall have OS-level access controls (e.g. SELinux, sVirt) configured to enforce capability controls on workloads.

The host system shall have logging configured for all OS-level access controls.

8.3 Physical controls and alarms

The host system shall provide software alarms to expose tampering of physical hardware (e.g. BIOS alarms).

The host system shall have an interface to provide authorized external services information to hardware-generated events.

8.4 Authentication controls

The host system shall provide authentication for all user accounts.

The host system shall ensure that all authentication credentials are presented in an encrypted fashion.

The host system shall not store any authentication credentials in clear text or unencrypted.

The host system shall disable all non-essential functions for user accounts (e.g. administrative privilege, change root capabilities, input/output (physical and virtual)).

The host system shall disable all non-essential functions for service ("non-user") accounts (e.g. shell access, remote login, input/output (physical and virtual)).

8.5 Access controls

The host system shall implement mandatory Attribute-Based Access Control (ABAC) - as defined in NIST-SP800-162 [i.13].

The host system shall extend ABAC to restrict the capabilities available to the super user/root administrative user.

8.6 Communications security

The host system shall use one or more of the following methods for communications security:

- TLS (minimum version 1.2), employing cryptographic primitives specified in clause 8.5, with client and server authentication required.
- IPSec employing cryptographic primitives specified in clause 8.5.
- The host system shall not use any of the following methods for communications security: SSL (any version), TLS 1.0, TLS 1.1.
- The host system shall not employ anonymous TLS.
- The host system should, when employing TLS or IPSec, use the latest approved version non-draft available.

8.7 Boot

The boot process is an integral part of the Trusted Compute Base: see clause 8.1 for a discussion of this topic.

8.8 Attestation

The host system shall support static root-of-trust measurement for hardware-based remote attestation.

The host system shall support dynamic root-of-trust measurement for hardware-based remote attestation.

8.9 Hardware-mediated execution enclaves

A hardware-mediated execution enclave (HMEE) is described in ETSI GS NFV-SEC 009, clause 6.16 [i.3].

The host system shall provide workloads access to hardware-mediated execution enclaves.

The host system shall make use of hardware-mediated execution enclaves when protecting its own sensitive processes

The host system shall provide the ability for authorized actors to perform a secure wipe of sections of memory in the HMEE.

The host system shall provide workloads with isolated enclaves.

8.10 Hardware-Based Root of Trust (HBRT)

The host system shall include an HBRT. The HBRT shall be based on hardware-based TPM or equivalent hardware root of trust (e.g. Secure Element including TPM functionalities, HSM including TPM functionalities).

The host system HBRT shall be able to provide isolated instances of the HBRT capabilities for individual workloads.

The host system HBRT shall include a hardware based compute engine to be used by the workloads for cryptographic and security functionality.

8.11 Self-encrypting storage

The host system should provide self-encrypting storage for use by workloads.

The host system should make use of self-encrypting storage, however the use of self-encrypting storage is a mechanism of last resort, and not relied on to protect data.

The host system shall be able to identify whether a storage entity (e.g. a hard disk) is employing self-encrypting storage.

The host system shall have an interface to provide authorized external services with information about whether a storage entity is employing self-encrypting storage.

The host system shall be shall be able to refuse to utilize any storage media which does not comply with the self-encrypting policies of the host system.

Workloads shall be responsible for providing sufficient confidentiality protection for data and processes under their control.

8.12 Direct access to memory

The host system shall be able to deny direct access to memory to particular hardware resources.

The host system may specifically implement direct access to memory to hardware resources as a general capability.

The host system may specifically implement direct access to memory to hardware resources by particular workloads.

The host system shall implement hardware memory access controls (IO-memory management, e.g. IOMMU, VT-d, AMD-Vi).

8.13 Hardware Security Modules

The host system MAY provide a Hardware Security Module (HSM) for key storage and cryptographic operations of sensitive applications.

If the host system provides an HSM:

- The host system shall provide an interface to workloads for their use of HSM resources.
- The security level of HSM shall be verifiable by sensitive applications.
- The host system MAY use the HSM for its own cryptographic operation.
- If the HSM is used by the host system, the HSM shall ensure isolation between its operation and those of workloads.
- If workloads are using the HSM, the HSM shall ensure isolation between different workload operations.
- The host system MAY use the HSM as a HBRT. In this case the HSM shall be compliant to the HBRT requirements in clauses 7 and 8.10.

8.14 Software integrity protection and verification

The host system shall verify the provenance and integrity of all instances and versions of software components before installing them.

The host system shall refuse to install all software which fails verification against the policies held by the host system.

The host system should verify the integrity of software components before execution.

History

Document history		
V3.1.1	January 2017	Publication