ETSI GS NFV-SEC 024 V1.1.1 (2026-04)

GROUP SPECIFICATION

Network Functions Virtualisation (NFV);
Security;
Security Management

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-SEC024

Keywords
cyber security, network monitoring, NFV,
policy management, security,
security management, threat analysis,
threat intelligence

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 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are 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 IPR online database.

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, 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.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™**, **LTE™** and **5G™** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**® and the GSM logo are trademarks registered and owned by the GSM Association.

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 describes the whole-system security framework required to manage NFV-based virtualised networks securely. The present document provides an architecture and capabilities for security management, which includes MANO, NFVI (including the underlying compute hardware infrastructure), the virtualised function application layer (e.g. 5G) and PNFs. The security management architecture addresses all network and VNF lifecycle stages from VNF onboarding, instantiation, VNF instance runtime and post-VNF instance teardown cleanup.

The present document considers both baseline security requirements and policies which need to be applied across all network functions and additional requirements that are applicable to sensitive network functions.

The present document is intended to include, update and replace NFV Security Management and Monitoring concepts that were defined in ETSI GS NFV-SEC 013 [i.2].

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 in the ETSI docbox.

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 GS NFV-IFA 026 (V3.2.1): "Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Architecture enhancement for Security Management Specification".

[2] ETSI GS NFV-SEC 025: "Network Functions Virtualisation (NFV)Release 4; Security; Secure End-to-End VNF and NS management specification".

[3] ETSI GS NFV-SEC 021: "Network Functions Virtualisation (NFV) Release 5; Security; VNF Package Security Specification".

[4] ETSI GS NFV-SEC 026: "Network Functions Virtualisation (NFV) Release 5; Security; Isolation and trust domain specification".

[5] ETSI TS 104 000: "Lawful Interception (LI); Internal Network Interface X0".

[6] ETSI TS 104 007: "Lawful Interception (LI); Lawful Interception Architecture".

[7] NIST SP 800-88 Rev.2: "Guidelines for Media Sanitization".

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 may be useful in implementing an ETSI deliverable or add to the reader's understanding, but are not required for conformance to the present document.

[i.1] ETSI GR NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[i.2] ETSI GS NFV-SEC 013: "Network Functions Virtualisation (NFV) Release 3; Security; Security Management and Monitoring specification".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in ETSI GR NFV 003 [i.1] and the following apply:

Security Agent: distributed security function performing security monitoring/management with a local actionable behaviour

Security Management: functionality that applies security policy to a virtualised network based on both predefined default policy and active analysis of information provided through security monitoring

NOTE: Security management actions will consist of both passive default security policy automatically applied by NFV-MANO (including through VNFDs or of vendor / CSP configuration) and active real-time security management actions where the Security Management system actively updates or overrides default passive policy.

Security Monitoring: functionality that collects and performs analysis of relevant events from across the virtualised network, which allow the Security Management and Monitoring system to make informed security management decisions

NOTE: Security monitoring is not restricted to real-time (or near real-time) collection and analysis of network events. Virtual network-wide monitoring will include security analysis of longer-term logging, AI data set analysis and human intelligence to predict and update monitoring criteria.

3.2 Symbols

For the purposes of the present document, the symbols given in ETSI GR NFV 003 [i.1] apply.

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in ETSI GR NFV 003 [i.1] and the following apply:

ABAC	Attribute Based Access Control
AI/ML	Artificial Intelligence / Machine Learning
A-SA	Adjunct Security Agent
CA	Certificate Authority
CSP	Communication Service Provider
DLP	Data Loss Prevention
EOL	End-Of-Life
E-SM	Embedded Security Agent
FQDN	Fully Qualified Domain Name
GUID	Globally Unique Identifier
HA	High Availability
HMEE	Hardware-Mediated Execution Enclave
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
I-SA	Infrastructure Security Agent
M-SA	MANO Security Agent
OSSM	OSS/BSS Security Manager

SA	Security Agent
SIEM	Security Information and Event Manager
SM	Security Manager
SMA	Security Monitoring & Analysis
SO	Security Orchestrator
TDE	Trust Decision Engine
TLS	Transport Layer Security
VBS	VNF Bootstrapping Service

4 NFV Security Management and Monitoring Overview

4.1 General

The security of any system or network is ultimately governed by the weakest link. Attackers will seek out points or interfaces within a system or network that provide the easiest way in and those points of entry or exploit do not necessarily need to be associated with the ultimate network resource or goal that the attacker wishes to access or obtain. Therefore, security shall be implemented to an equally strong level throughout the system or network, from the underlying NFVI, management systems and to the application layer containing end user services (e.g. 5G).

Similarly, while approaches such as Transport Layer Security (TLS) can provide a high level of security protection against person in the middle attacks on network links or against unauthenticated attempts to access interfaces, they provide limited protection if an attacker gains access to an insecure trusted endpoint and spreads their attack over the encrypted TLS connections. All the encryption does is stop security monitoring systems from preventing or detecting the attack. This is especially true in virtualised systems, as the encrypted virtual links and virtual endpoints all exist in the same logical memory space. Therefore, if an attacker can access the resources used to implement the virtual connection, they can also access the resources of the endpoints and management system.

Considering this from a system or network-wide monitoring perspective, the same applies. Focusing security defences or threat detection in isolation at the application layer, NFV-MANO layer or NFVI will not allow the detection of threats or attacks which cross the layers. Similarly, taking security mitigations at one layer without considering the big picture of an attack across all layers may result in attack amplification or an attacker being able to force the use of specific network resources which have been compromised. Security monitoring is described further in clause 7.

Therefore, it is necessary to consider security management using a system-wide architecture that is able to detect threats in user services, network management, virtual network functions and underlying NFVI hardware, in order to apply system-wide policies and take security actions which minimize both attack propagation and service disruption. Security management is described in further detail in clauses 5 and 6.

Similarly, application layer functions do not magically appear for discovery in a service-based architecture and NFVI server resources do not get allocated for use with a VNF without the OSS/BSS having requested (either directly or indirectly through delegated policy) that the VNF was or resources were required. Security monitoring and management needs to have a holistic view across the entire chain of events.

Conceptually, this should also include physical security systems used to protect the data centres. However, such physical security is out of scope of the present document.

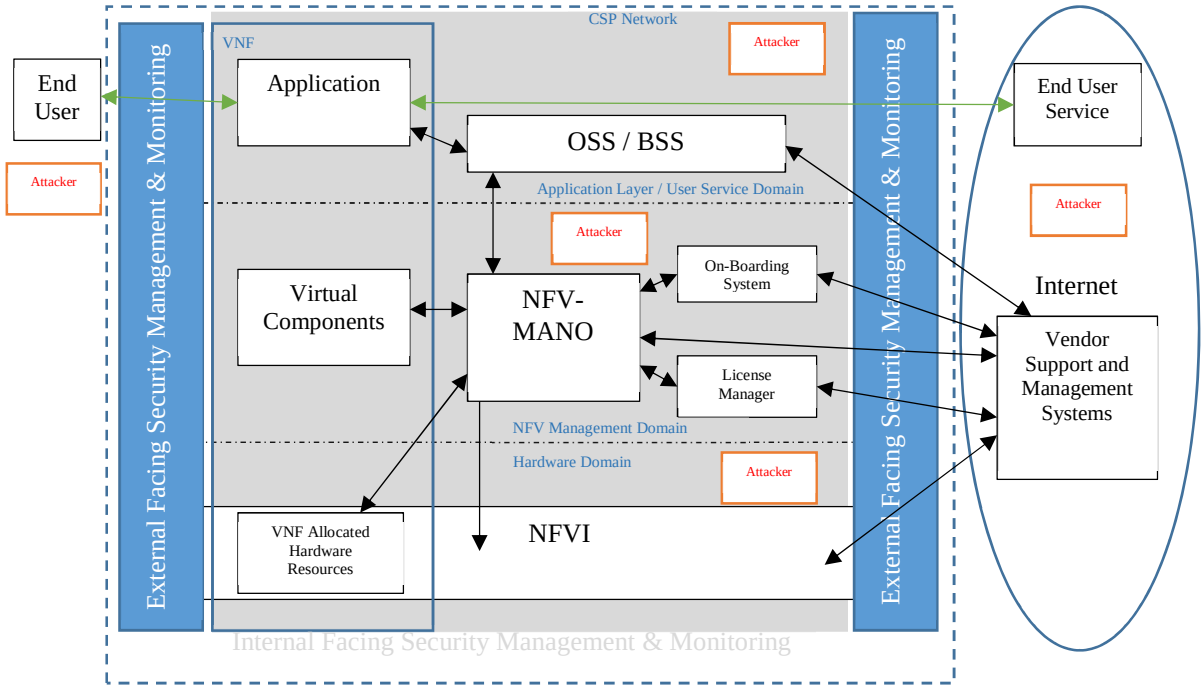

Figure 4.1-1: Simplified High-Level Network Wide Security Management and Monitoring Architecture

Figure 4.1-1 shows a simplified high-level security management and monitoring architecture. The key concept here is that security monitoring needs to be considered across the whole virtualised network (including any legacy PNF components). Figure 4.1-1 makes a distinction between internal and external facing security, as there are differences in functionality, actions and policies applied to external vs internal facing security functions. The 5 distinct "Attackers" in Figure 4.1-1 illustrate that attacks can occur both internally and externally to the virtualised network, and to any layer of the virtualised network. Furthermore, the attackers may be either external to functions within a given domain or internal (i.e. attacker has legitimate credentials to perform certain actions not in the role of an attacker). Further consideration of attack scenarios is described in ETSI GS NFV-SEC 025 [2].

4.2 Whole System Management and Monitoring Lifecycle Overview

There are several philosophical reasons for implementing security as an overlay over the network and the need for multiple sources of data to verify NFV-MANO reported information, especially in the case of compromise:

1) Holistic Protection: Security should be comprehensive and all-encompassing. By overlaying security measures on top of the network infrastructure, a layered defence mechanism is created. This approach aligns with the principle of "defence in depth", where multiple layers of security are employed to protect against various threats. It is akin to having multiple lines of defence to safeguard critical assets.

2) Adaptability and Flexibility: Security should be adaptable and flexible, capable of evolving to counter emerging threats. By having security as an overlay, it is easier to adapt to changing threat landscapes. This aligns with the philosophy that security should not be static but should continuously evolve to address new vulnerabilities and attack vectors.

3) Risk Mitigation: From a philosophical standpoint, risk mitigation is a fundamental principle in cybersecurity. Overlaying security allows for risk management by creating virtual boundaries and control points. It is akin to setting up checkpoints in a fortress to detect and stop potential threats. In essence, this approach aligns with the philosophy of minimizing risk exposure.

Multiple sources of data are required to verify NFV-MANO-reported information, especially in the case of compromise, so that:

1) Trust but Verify: While NFV-MANO systems are essential for the orchestration and management of network resources and VNFs, it is crucial not to rely solely on their reports. Multiple sources of data provide independent verification, reducing the risk of blind spots or deception in the event of a compromise.

2) Resilience and Redundancy: By having multiple sources of data, ensures that even if one source is compromised or manipulated, others can provide a more accurate picture of the network's state.

3) Transparency and Accountability: Multiple data sources enhance transparency by making it more difficult for malicious actors to manipulate or cover up their activities within the network. This aligns with the philosophy that accountability should be a cornerstone of any security system.

Implementing security as an overlay over the network aligns with the principles of comprehensive protection, adaptability, and risk mitigation. Additionally, relying on multiple sources of data for NFV-MANO verification reflects the concepts of trust but verify, resilience, and transparency in the realm of cybersecurity. These principles help create a robust and reliable security posture in an ever-evolving threat landscape.

A comprehensive approach to NFV security management and monitoring is critical to effectively safeguard VNFs throughout their lifecycle. This begins with the onboarding of VNFs, when security policies, controls, and trust requirements are defined. During this phase, the security attributes of the VNF package shall be verified (see ETSI GS NFV-SEC 021 [3] and ETSI GS NFV-SEC 025 [2], clause 5.3). This onboarding sets the foundation for consistent enforcement and monitoring of security as the VNF transitions into an active service.

During the instantiation phase, the security properties and policies of the VNF shall be taken into account, this shall include requirements for VNF isolation and any trust boundaries which exist within the Network Service and VNF components, as well as within the NFV infrastructure (see ETSI GS NFV-SEC 026 [4] for more detail on isolation and trust domains).

Once deployed and running, VNFs require ongoing security management that adapts to dynamic operational conditions. Policy-driven enforcement, including the dynamic application, adjustment, and revocation of security controls in response to lifecycle events such as scaling, migration, or modification is required. Security monitoring is equally crucial: continuous health and anomaly monitoring across VNF, management, and infrastructure layers provides near real-time insights and enables prompt mitigation of threats, unauthorized changes, or failures. The lifecycle perspective ensures that monitoring is not static, but keeps pace with the evolving VNF ecosystem by integrating with NFV orchestrators and security managers to enforce, verify, and, if necessary, remediate both security policy and posture as the system evolves.

As VNFs are eventually decommissioned or retired, the security management lifecycle concludes with the secure removal of instances, revocation of privileges, and cleanup of associated resources, configurations, and policies. This step is vital to prevent residual data exposure and potential exploitation of stale entities. Adhering to a whole-lifecycle security management approach offers assurance that security is both proactive and adaptive, addressing risk from the earliest onboarding stage through to decommissioning, for the full duration of the VNF's presence in the network.

5 Security Management Framework Architecture

5.1 Security Manager Architecture

5.1.0 Introduction

The logical security management architecture in the context of NFV-MANO and OSS/BSS is depicted in Figure 5.1.0-1.

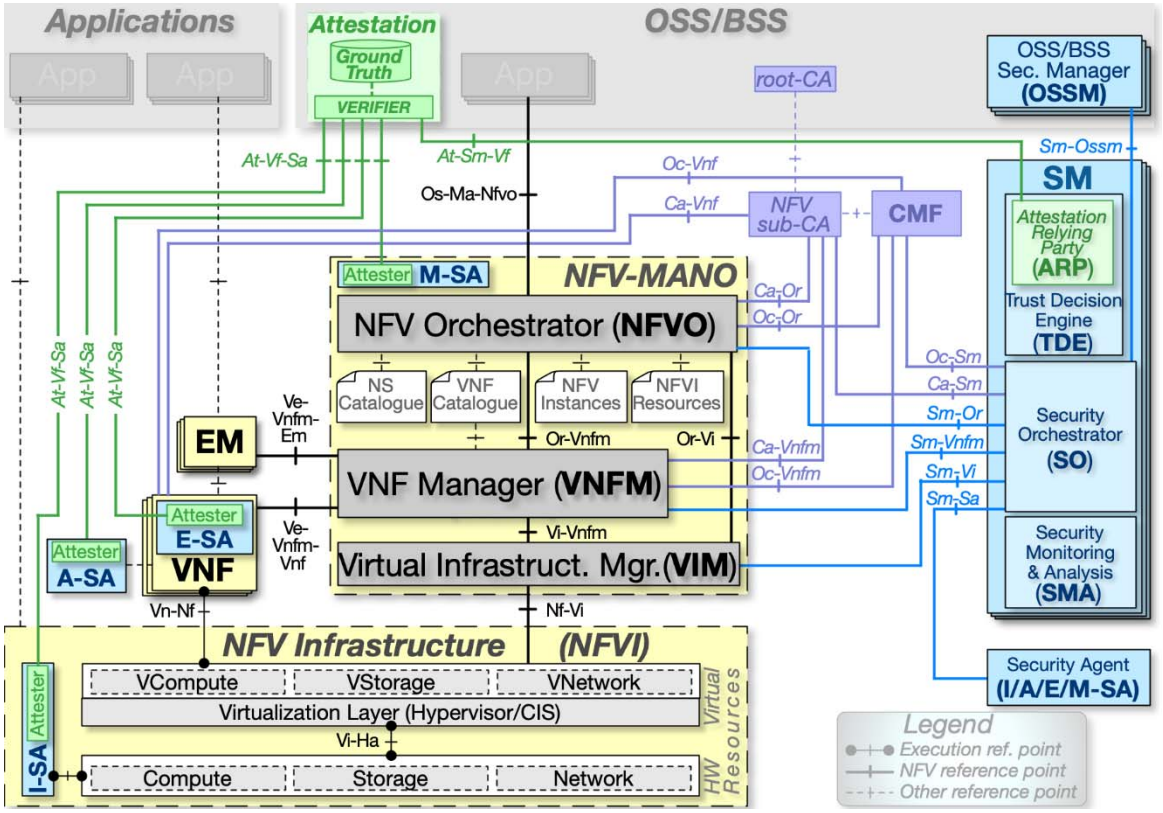

Figure 5.1.0-1: Security Management Architecture

This architecture defines the following functions and interfaces:

- Security Manager (**SM**) - A function that centralizes security concerns system-wide, for one, and only one, security trust domain.
- Security Agent (**SA**) - I/A/E/M-SAs security function performing security monitoring/management with a local actionable behaviour. For a given implementation, one or more SAs may be associated with or embedded in VNFs, NFVI and MANO. Some SAs will have the ability to attest (as Figure 5.1.0-1 depicts), others will not. The SM should therefore consider their output accordingly.
- Sm-Ossm Interface (**Sm-Ossm**) - An interface that passes security-related information between the virtualisation and the application layer within one or more trust domains.
- Sm-Sa Interface (**Sm-Sa**) - An interface that passes security-related information between the SMs and the SAs. In the present document, only the generic types of information to be reported by SAs to SMs are specified, i.e. those agnostic to specific SA applications.

All dotted line interfaces are outside the scope of the present document. Interfaces Sc-Or, Sc-Vnfm, and Sc-Vi are described in ETSI GS NFV-IFA 026 [1].

5.1.1 Security Manager

5.1.1.1 Security Orchestrator

The Security Orchestrator (SO) acts as a secure proxy between the sensitive Trust Decision Engine (TDE) and the rest of the network, facilitating provisioning and other security events. It holds the network security state.

5.1.1.2 Trust Decision Engine

The TDE is the beating heart of the system. It manages the network-wide security primitives (keys, nonces, salts, etc.) needed by the network functions. It contains the Relying Party in the attestation layer.

The element of the SM responsible for the configuration of SAs contains a configuration repository function and the attestation relying party.

5.1.1.3 Security Monitoring & Analysis

The security monitoring and analysis functionality of the SM incorporates such functions as the Security Information & Event Management (SIEM) and threat hunting activities. As the amount of information consumed by the SM increases, it is likely that Artificial Intelligence / Machine Learning (AI/ML) will be needed to aid in the analysis and alerting of potential indicators of compromise.

5.1.2 Security Agent

5.1.2.1 Introduction

The SA is a security function performing security monitoring/management with a local actionable behaviour. In an NFV-based environment, the SAs communicate with the corresponding SM in their security trust domain based on configurable security policies. The SA may provide security data both to the SM and to the OSSM (via the SM), some of which is only intelligible to the OSSM (see clause 5.5) at the application layer.

Preconfigured data for initial configuration is assumed to be injected or loaded at SA instantiation (e.g. from NFV-MANO). An API for runtime configuration could also be available. Data which would need to be preconfigured includes the likes of the SM to register/communicate with, authentication credentials and any required identifiers/identities. The configuration of this data can be done via various out-of-band methods, for example, using CloudInit, Infrastructure as Code templates, etc.

If an SA is allowed to take specific actions in its security trust domain, the access rights shall be explicitly authorized in the security policy. SA management should be proven capable of guaranteeing SA security trust domain isolation.

NOTE: The security policy may cover two aspects:

a) *how* the SA agent is secured (e.g. where the SA is to be located, what trust domain it belongs to, what/where are the keys); and

b) *what* the SA should do, e.g. enforce monitoring, filtering, collection options.

While different variants of SA can be envisioned (virtualised SA, NFVI SA, SA service, etc.), the present document describes the differences in the management of such SAs with their NFV-MANO/SM interactions.

5.1.2.2 Embedded SA

The Embedded SA is an SA integrated with the VNF/VNFC that it monitors and/or protects. The E-SA contains information relevant to the VNF. The E-SA will typically be supplied by the VNF supplier and is included in the VNF package. The E-SA implicitly follows the NFV-MANO LCM of the VNF/VNFC it is embedded into and relies on NFV-MANO to perform secure instantiation and to provide the SA initial application configuration.

5.1.2.3 Adjunct SA

The Adjunct SA (A-SA) is not embedded in a VNF/VNFC but can be associated with one or several VNF/VNFC. The A-SA can monitor the functions of the VNF/VNFC exposed outside of the VNF/VNFC but has no access to the internal architecture or operational state of the VNF/VNFC.

The A-SAs may form an overlay network of SAs and may be orchestrated by the SM. A-SAs could operate in a security trust domain that is different from the VNF. The security policy shall state that the A-SAs trust domain is allowed to monitor the VNF trust domain.

An SA associated with the VNF accesses the tenant network, including this VNF. Accessing NFVI from the VNF layer is likely to require opening, for example, the OpenStack REST API or local orchestration to the SA. If the Operations, Administration and Management (OAM) is deployed as a separate tenant, different from the tenant operating the VNF and the SA, then the SA would also access the OAM network. Network isolation is therefore broken. If such network or tenant boundary crossing is accepted by the operator, this shall be explicitly stated in the security policy and the risk in case of a compromised SA shall be mitigated. As a result, the SA needs to be defined as a security asset to protect for SM.

5.1.2.4 Infrastructure SA

This SA functionality is associated with, and with concerns limited to one NFVI / CIS cluster. I-SAs communicate with the corresponding SM for their given security trust domain. These SAs contain information relevant to the NFVI/CIS cluster layer. They may be lifecycle managed by the NFV layer, or separately managed.

The type of information that the ISA will report is the existence or non-existence of HMEEs on a particular host, geographical location for use in geofencing, etc.

5.1.2.5 MANO SA

The M-SA functionality is associated with the NFV-MANO components (NFVO, VNFM, VIM, CISM, CCM). Although Figure 5.1-1 depicts a single M-SA for the whole of NFV-MANO it is expected that each NFV-MANO component will have one or possibly more M-SAs. From an SM perspective NFV-MANO is simply another workload which needs to be monitored. Like the I-SAs, the M-SA follows the NFV-MANO infrastructure blocks lifecycle.

It is expected that NFV-MANO entities have a long lifetime. The type of information MANO SAs should provide to the SM is therefore left for implementation.

NOTE: How M-SA is implemented in a PaaS environment is for future study.

5.2 Security Manager Modes

The SM and NFV-MANO shall support three modes of operation:

- Passive: SM is able to subscribe to applicable lifecycle management events provided by NFV-MANO, but the SM does not take any active part in the lifecycle management of the VNFs.
- Semi-Active: SM analyses applicable lifecycle management events passed to it by NFV-MANO. The SM may provide security policies (e.g. geographical restrictions) to NFV-MANO as part of VNF lifecycle management, but the SM takes an otherwise passive part in VNF lifecycle management. The SM is able to request NFV-MANO, or other entities (e.g. other SMs, OSSM, OSS/BSS), to undertake security mitigation actions (e.g. terminate a VNF instance, or surrounding it with firewalls without affecting its lifecycle). NFV-MANO (or the other entities) can refuse to comply with the request.
- Fully-Active: NFV-MANO passes applicable VNF lifecycle events to the SM and requests approval from the SM. The SM authorizes, modifies, or rejects NFV-MANO requests per applicable security policy. The SM is also able to instruct NFV-MANO to take security mitigation actions (e.g. immediately terminate a VNF instance). In fully active mode, the SM applied policy will supersede MANO-level attributes (e.g. in the VNFD).

5.3 Multiple Trust Domains and Security Managers

5.3.1 Introduction

In networks with multiple trust domains or where a CSP wishes to achieve security role separation, there may be one or more SMs. Each SM may operate in Passive, or Semi-Active or Fully Active mode as described in clause 5.2.

It shall be possible for the SMs to act independently of each other, or for SMs to operate in a hierarchical arrangement, where one SM may be able to issue VNF termination instructions across trust domains of one or more sub SMs.

In hierarchy terms, a sub SM is an SM which is overseen or controlled by another higher security level SM. For example, a sub SM in Semi-Active Mode may be subservient to a network-wide Fully Active SM. In this case, the sub SM is able to fulfil its role autonomously, but the higher-level SM would be able to overrule it at any time. NFV-MANO needs to be able to support such hierarchical models and provide interface instance isolation for such sub SM to SM relationships.

5.3.2 Trust Domains

5.3.2.1 Trust Domain Definition

A trust domain is a collection of functions that share the same set of administrative and security concerns and policies (particularly access control). It is expected that compartmentalization of trust domains spans the vertical stack of functionality in the CSP network, ideally from hardware to the application layer. Administrators in the CSP network are assigned to one or more trust domains based on CSP criteria, using state of the art Attribute Based Access Controls (ABAC). In virtual networks, multiple trust domains should be considered by CSPs during the deployment phase, where a CSP wishes to achieve security role and management separation, security isolation, separation between sensitive and non-sensitive components, etc.

5.3.2.2 Trust domain isolation

In this architecture, trust in the network is based on technical procedures and requirements placed on functions in an attempt to replace the assumed trust of functions on the basis of their location only. Because the network functions where the SAs reside are dynamically created, this architecture is designed to be able to integrate the orchestration and management of network functions and the SAs. Consequently, the SM shall contain functions that handle the trust establishment of SAa, as they are instantiated and as they evolve through their life cycle. In such context, this architecture contains a logical separation of trust domains and requirements on how information exchange between trust domains is to be handled to reduce the risk of compromise propagation.

While maximal sub-trust domain separation is clearly a good principle to drive implementation, this approach is costly both in terms of network and personnel resources. An implementation may choose to collapse sub-trust domains, but the risks of doing so can only be quantified if the implementation of the present document takes the most restrictive approach. Every decision to collapse (sub) trust domains shall be accompanied by a thorough security analysis, coupled with explicit security mitigations.

5.4 Security Domain Bootstrapping

5.4.1 General Introduction

In order for a new VNFI/VNFCI containing an SA to be configured for use, the VNFI/VNFCI needs to be able to establish communication with the SM. This presents an issue, as instance-specific SA configuration data and keys cannot be provided in the generic VNF image.

Clauses 5.4.2 and 5.4.3 consider two scenarios for establishing initial communication with the SM/OSSM. Where practical, the trusted MANO is considered to be the preferred option. However, the aim of both clauses is to arrive at the same secure running LI implementation.

5.4.2 Low criticality deployments

In the low criticality deployment scenarios (e.g. consumer retail services), MANO is trusted by the OSSM to issue initial SA VNFI/VNFCI identities and communications certificates. Standard MANO lifecycle management is used for OSSM. In this deployment, there will be minimal to no segregation and there is a high risk of OSSM bypass or compromise. This deployment scenario is not covered in the present document.

5.4.3 Medium criticality deployments

In the medium criticality deployment scenario (e.g. enterprise services), MANO is considered sufficiently trusted by the OSSM to issue initial SA VNFI/VNFCI identities and communications certificates. In addition, the OSSM acts as a sub-CA for the security domain, under a common operator root CA, which is shared by both MANO and the OSSM.

When MANO instantiates a new VNFI/VNFCI containing an SA function, or a new VNFI/VNFCI with an SA function associated with it, MANO will allocate a MANO-level identity to the VNFI/VNFCI. MANO performs initial configuration based on the VNFD and other SA configuration information provided to MANO by the SM. MANO will also provide the VNFI/VNFCI with initial identity verification and communication certificates.

Once configured, the VNFI/VNFCI will initiate communications with the SM using the initial MANO-issued certificates to perform authentication and IPSEC or TLS tunnel establishment. Once the secure configuration tunnel has been established, the SM will configure the VNFI/VNFCI as an SA, ready for use by the OSSM. This will include the provision of security domain-specific OSSM sub-CA certificates for use with the application layer security interfaces. The MANO level certificates would be retained and used for MANO level maintenance and mobility of the SA associated with the particular VNFI/VNFCI, etc.

This scenario is considered to be the recommended approach for initial provisioning and communications establishment as SAs can be instantiated by MANO using the same basic identity and security procedures used for other VNFs. In addition, the OSSM is able to explicitly trust the MANO-issued certificates as both the OSSM and MANO trust the same operator root CA.

5.4.4 High criticality deployments

In the high-criticality deployment scenario (e.g. national critical infrastructure), the OSSM does not rely on MANO to manage the initial communications establishment of SAs with the SM. In this scenario, the OSSM root certificate is not provided by a common operator root CA. The OSSM will act as the dedicated root CA for the security domain, independent of other security-domain functions and MANO.

Since MANO is not fully trusted, the SM and a newly initiated VNFI/VNFCI SA need to be able to establish trust independently of MANO, and an initial secure communication channel over which the SA instance-specific keys and configuration can be downloaded. It is assumed that MANO is trusted to apply a unique name to the new VNFI/VNFCI SA according to the normal MANO naming scheme and that MANO will provision the VNFI/VNFCI SA with certificates for the purpose of managing the VNFI/VNFCI SA at a MANO level (e.g. mobility and scaling).

In this scenario, the SM may be able to verify that the VNFI/VNFCI SA is a valid SA function through verification of the signing applied to the SAs when they are stored into the VNF catalogue, but cannot use the MANO certificates for initial connection as they are considered untrusted. The mechanism required to allow establishment of initial connection and trust in the low-trust MANO scenario is not defined in the present document.

5.5 OSSM, VNFI/VNFCI and SA Connectivity Tracking

5.5.1 General

In a legacy network, the OSSM and SAs are implicitly configured to know from where to obtain the information required to enable security functionality (e.g. firewalls), and the relationships between network elements required to allow the OSSM to correlate the information being received from the CSP network.

With NFV, both the number of VNFIs/VNFCIs and associated SAs, and the number of interconnecting VNFIs/VNFCIs are potentially highly dynamic. In order to reliably perform security monitoring and management, and provide necessary correlation information to the application layer, the OSSM and SM both need to be able to derive the total number of SAs at any point in time. In addition, the OSSM and SM need to understand the SDN level interconnectivity between the various VNFIs/VNFCIs.

At the cost of increased visibility of the SAs by MANO, it may be possible for the SAs to dynamically adapt to the VNFIs/VNFCIs interconnected around them utilizing default MANO procedures. However, elements of the correlation information required by the OSSM are likely to require the OSSM and SM to generate and maintain a real-time service level (and potentially NFV level) network map. Similarly, requiring the SAs to be overtly visible to MANO in order to utilize standard MANO procedures likely violates SA confidentiality requirements across security trust domains.

The adjunct (non-embedded) SA scenario is likely to be even more difficult, as a change in the number or relationships around a VNFI/VNFCI for which external SA targeting is being applied (e.g. number of peers or SDN links), may result in the external SA no longer being in the correct place, or no longer monitoring 100 % of the traffic, unless the OSSM or SM is able to constantly adapt to network changes.

5.5.2 OSSM VNFI/VNFCI Tracking

The OSSM and SM are responsible for signing embedded VNF SAs as part of the software catalogue onboarding process if required by CSP policy. Therefore, the OSSM/SM needs to be aware of all VNF/VNFC types that require or contain SAs. This information is either acquired through the onboarding process, or by other means outside the scope of the present document. However, the OSSM or SM needs to understand the meaning of all VNF/VNFC types within the MANO VNF catalogue in order to understand the function of any new VNFI/VNFCI when it is instantiated.

A manual process/naming scheme or a standard automated process may be required in order for the OSSM or SM to understand the meaning or function of VNFs in the MANO catalogue.

MANO is required to be able to notify the SM the start of every VNFI/VNFCI instantiation request and subsequent confirmation of the success or failure. This provides the SM and OSSM with the basic information required to construct a list of running VNFIs/VNFCIs.

In order to maintain the list, MANO also needs to provide notifications of VNF de-instantiation, so that VNFIs/VNFCIs can be removed from the OSSM's running list.

To support error recovery and auditing, MANO should be able, when requested by the OSSM or OS, to provide a complete list of all currently running VNFIs/VNFCIs. The OSSM should be able to establish control over already running SAs following a restart of the associated VNFI/VNFCI. The OSSM or OS should be able to dynamically re-establish control and security associations without requiring a re-instantiation of running SAs.

5.5.3 OSSM VNFI/VNFCI Connectivity Tracking

Assuming the OSSM is able to maintain a running list of VNFIs/VNFCIs, as per clause 5.5.2, then the OSSM also needs to maintain a network map of the interconnectivity relationships between the VNFIs/VNFCIs.

At a minimum, MANO needs to report to the SM the interface connectivity requirements for each VNFI/VNFCI as included in the VNFD for that VNF type, and, for a successful VNFI/VNFCI Instantiation, the subsequent IP address/FQDN naming applied to those interfaces (as known by MANO). For a full NFV network this may be sufficient to determine the full network service connectivity map.

However, where the SDN applies NAT between VNFI/VNFCI connections, or where there is other network routing information which the OSSM cannot resolve from the VNFI/VNFCI connectivity information supplied by MANO, then the OSSM may need additional network service layer information. Such information provision to the OSSM is outside the scope of NFV. This is also likely to be an issue in part NFV, part legacy, mixed deployments.

As with clause 5.5.2, the OSSM/SM should be able to obtain sufficient information to re-establish the current VNFI/VNFCI connectivity map following an OSSM/SM restart, or where security monitoring/management is enabled in the network (SAs started) after the first VNFIs/VNFCIs are instantiated.

5.5.4 VNFI Scaling/Migration

A VNFI/VNFCI's scaling or migration will have a number of potential impacts on the OSSMs VNFI connectivity map:

1) If a VNFCI is added to an existing VNFI, the number of, or bandwidth, of the SDN links may change. The OSSM should be able to receive sufficient information over the Sm-Ossm interface from MANO via the SM to understand such changes, as part of the scaling process.

2) If a VNFI is migrated from one location to another, the OSSM needs to receive sufficient information over the Sm-Ossm interface from MANO via the SM to understand such changes and update the VNFI/VNFCI list and VNFI/VNFCI connectivity map accordingly.

NOTE: Extensions to existing MANO procedures required to provide sufficient information to the OSSM/SM as a result of VNFI migrations or scaling are for future study.

3) If a VNFCI is migrated from one location/host to another, any key pair associated with the VNFCI will have to be moved or revoked and recreated. If the key pair is migrated with the VNFCI, the private key will be placed at risk during any transfer. In addition, if the VNFCI uses certificate management in direct mode (see ETSI GS NFV-IFA 026 [1], clause 5.2.3), any linkage between the key pair and the hardware on which it was generated will be broken, preventing attestation of the key pair's creation.

6 Security Procedures and Policy Management

6.1 Instantiation/Boot Time Concerns

6.1.1 General

Clause 6.1 describes the instantiation/boot time information flows which can be dynamically executed without user intervention. These are described in detail in Figures 6.2.1-1 and 6.2.2-1. The figures do not describe specific implementation messages or control signals, but rather a sequence of pseudo-code-like functional events for provisioning, deploying and activating Security Management.

If an SA is somehow compromised, the SM is responsible for remediation. One option could be to perform a runtime inspection of SAs, and/or terminate and create a new SA. SAs can provide failure detection and remediation options, such that remedies can be made configurable and automated.

6.1.2 Secure VNF Bootstrap Protocol

Trust relationships shall be established during the bootstrapping phase. The relationships between the Security Manager (SM) and the Security Agents (E/A/I/M-SA) are analogous to the relationships between the Lawful Intercept (LI) Administration Function (ADMF) and Element of Lawful Intercept (ELI), where the Security Manager takes the place of the ADMF and the Security Agents replace the ELI. Within the LI context the trust relationship is established over the X0 interface as detailed in ETSI TS 104 000 [5]. Within the NFV architecture, the trust relationship between the SM and SA's can also be established using an X0 interface.

When the VNF is instantiated, a multi-phase process to build trust shall take place. To achieve this, the instantiation process shall follow the same procedures as the LI X0 interface phases 0 through 6 as detailed in clauses 5.1 and 5.2 of ETSI TS 104 000 [5] and further depicted in ETSI TS 104 007 [6], clause 7.

6.2 Run-Time Concerns

6.2.1 Initial Personalization and Policy Provisioning

SAs will each perform a secure protocol exchange with the SM to enable each SA to establish secure connections with their respective VNFM and EMs which have secure access to the personalization data. The VNFM and EM will personalize their SA(s), for instance, performing set name, security policy groups, per-tenant policies, and any additional configuration parameters and necessary state information.

Subsequent to the successful secure launch and instantiation of a VNF, and the successful completion of the VNF Secure Bootstrapping protocol as defined in the previous clauses, the VNF is activated. As discussed earlier, these VNF protocols are applicable to SAs, VNFs, VNFCs, and application workloads. This VNF now needs to be personalized, which includes secure configuration with an initial set of parameters, provisioned with appropriate metadata and state, establishing state for inter-VNF connections as in Service Function Chains, and seeded with any other VNF-specific information. Personalization also includes secure delivery of VNF vendor-specific private information, including security credentials, as identified in ETSI GS NFV-IFA 026 [1].

VNF personalization is followed by the secure delivery of a start-up set of policy and behavioural parameters for that VNF. These VNF policies and parameters are usually dynamic and dependent on deployment security policy and procedures. Policy can be delivered at the initial post-secure bootstrap time and also during active execution of the VNF. Secure policy updates to VNFs are necessary to maintain a consistent state across the entire system deployment, including NFV and traditional networks.

For NFV systems, the personalization and policy protocols are dynamic and automated. In an orchestration-driven system, VNFs will be dynamically and securely installed, and expected to follow the protocols defined in Figure 6.2.1-1.

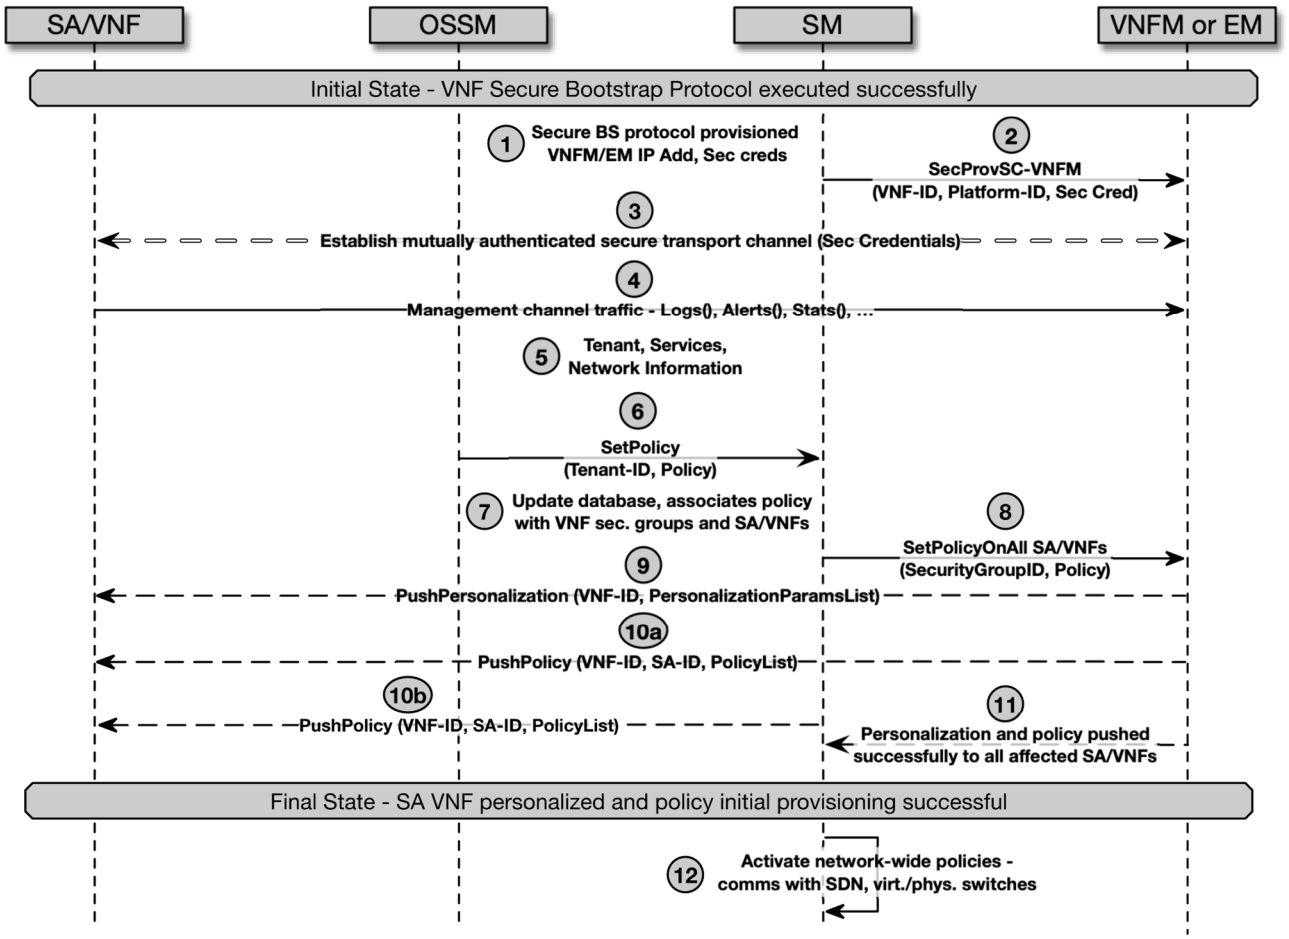

Figure 6.2.1-1: Secure VNF/SSA Personalization and Policy Protocol: Initial Provisioning

Personalization and Policy Protocol for Initial Provisioning is described in this clause. A necessary prerequisite is the prior secure and mutually authenticated establishment of encrypted, integrity and replay protected communication channels between the various management entities: OSSM ⇔ SM; SM ⇔ VNFM/EM. The secure bootstrap protocol enables a secure channel between the dynamically instantiated VNFs and their VNFM and EMs:

1) At the successful execution of the Secure Bootstrap protocol, the VNF has been secure instantiated and provisioned with the IP address (or any other reachability mechanism) of its VNFM and EM, and the initial root security credential(s) that uniquely identifies and would be used in a mutual authentication between this VNF and its VNFM/EM.

2) In a secure, OOB provisioning step, the SM will provision the VNFM and EM with the VNF_ID, Platform ID, and the same Security Credential that the SM provisioned into the VNF. This step is a prerequisite for a secure and trusted established communication channel between the VNF and its VNFM.

3) The VNF and the VNFM/EM perform a mutually authenticated key exchange procedure using the security credentials provisioned in the steps above. This choice of mutual authentication security protocol will be established by the security policy of the NFV deployment, and in most scenarios, may be TLS.

4) Once the secure and trusted channel is setup, the SA/VNF opens a management channel to its VNFM/EM. This channel is now used by the VNF to communicate management information (logs, alerts, statistics, etc.) with its VNFM/EM.

5) At the time a new tenant, network, service or capability is provisioned within the OSS/BSS, the OSS/BSS system communicates this new tenant, service, or network provisioning to the OSSM.

6) The OSSM is responsible for securely communicating the new tenant or service information to the SM, which will associate a set SAs with the tenant, service or network. This is based on the security policy of this new tenant, service, or network designed by the SM. Policies are identified with Tenant ID, and associated with configurable security groups or other programmable structures or databases established in the NFV deployment. In most cases, a tenant domain administrator may also assign and deploy their appropriate SAs to their workloads for NS-level security management. It is expected that the NFVO and the SM have a secure mutually authenticated channel. Establishing this secure channel is outside the scope of this procedure.

7) The SM updates the appropriate internal database with the new policy for the tenant, service or network, and the associated security group(s) or other security policy information. This allows long-term retention and access of sensitive security policy information. The contents of this database may optionally be encrypted, based on NFV deployment security policy. The policies will exist in the OSSM system as well. It is envisioned that the OSSM and SM are broad systems that maintain separate databases for policies, including formats, storage encryption requirements, different access controls, etc. Over time, the OSSM and SM policies may merge into a unified DB.

8) The SM distributes the new security policy across to all VNFMs and EMs that are responsible for managing the SA/VNFs. This policy distribution shall happen over a reliable protocol, with the SM maintaining state and ensuring policy consistency and delivery assurance across the entire network. In some tenant administrative domain cases, the Tenant's SM delivers policy to their VNFMs and EMs that manage their Physical Network Functions or Hybrid Network Functions. All messages in this step shall be signed and integrity protected, such that the receiving entities only enforce policies that are verified as integral, and originating from the SM.

9) The VNFM(s) and EMs push the VNF personalization data to all their VNFs that are affected by the new tenant, service, and/or network provisioning. This personalization data distribution has to happen over a reliable protocol with the VNFMs/EMs maintaining state and ensuring that all VNFs/SAs have received the personalization data. Personalization data includes secure configuration data, initial set of VNF parameters, metadata, connection information about other VNFs, vendor-specific information, performance, traffic engineering, and QoS parameters and policies, etc. These are the data and components that the newly instantiated VNF needs to get ready for processing workload traffic.

10) (a) The VNFM(s) or EMs push the tenant, service, and/or network policy to all their SA/VNFs. This policy distribution shall happen over a reliable protocol with the VNFMs maintaining state and ensuring that all VNFs have received the policy data. The policy list includes tenant-specific security processing policy, security traffic policy, security groups, network services processing policy, etc.

(b) In some cases, in which the SAs associated with a VNF run in a different trust zone than the VNF, the SM may push policies directly to the SAs, without passing through the VNFM.

NOTE: It is a deployment decision whether to perform steps (a), (b), or (a and b), based on the purpose of the SAs (one policy to all, or bespoke policies to each SA).

11) On successful completion of all Personalization and Policy updates, the VNFM(s)/EM(s) report back to the SM. The VNFM(s)/EMs will also report back all failures and errors that might have occurred. The SM will process all responses and issue alerts and status updates to the Security administrator, including securely logging.

12) Upon receiving all successful responses, the SM will activate network wide security policy for the traffic. To enable and activate traffic for the tenant, the SM will securely communicate with network-wide traffic switching elements, including SDN Controllers, Openstack Neutron Plugins, various virtual and physical Switch/Router Managers, etc.

SAs in the NFVI virtualisation/base OS layer and in physical devices will follow the VNF Secure Bootstrap protocol and the VNF Personalization and Policy protocols as described herein. In addition, these protocols will be followed by other tenant workloads.

In normal NFV operation, security management VNFs will be instantiated and triggered with personalization and policy information as per protocols defined above. In a dynamic and automated NFV deployment, it is expected that existing services, tenants, networks policies will be updated and would have to be securely and consistently pushed across the NFV deployment, including to any physical/hybrid network functions, as dictated by the security policy.

This security and management update can occur in many scenarios, including the addition or removal of tenants, tenant workload migration, update tenant's SLA and QoS, geo-based or regulatory requirements updates, failures or High Availability (HA) configurations of current VNFs, etc. Security policy updates are also expected as a remediation action of the automated security response within an NFV deployment, especially for security threat mitigation, malware response, network DoS attacks, and other such threat remediation scenarios.

The SM is expected to orchestrate the new security policy across the NFVI administrative domain, and within each tenant's virtualised and physical deployments. Security management policies are part of administrative domain of NFVI hence are orchestrated by the NFV SC. Tenants also have their security management policies that can be orchestrated by the tenant's SMs to their SAs.

6.2.2 Runtime Personalization and Policy Updates

Policy associations between VNFs and SAs are provisioned, for instance, when tenants or new services (VNFs) are provisioned. Tenant VNFs can be put into groups based on various criteria and those groups can be assigned services exposed by SA security management devices. The SM can then push down security policy associated with the VNFs groups to the SAs. These policies that are then propagated across all SAs. This enables a set of SAs to apply the security functions' policies uniformly across the NFV deployment, no matter which SA processes that data, management and/or control traffic.

SAs may interact with the virtual switch and the physical network infrastructure to translate that mapping of a VNF group to a policy before handing off traffic to the SA. A VNF group may be based on traffic type (e.g. VLAN, MPLS, GTP, etc.) and their associated security policies that will be enforced by the SA. These mappings may change due to any number of reasons including new tenants joining, tenants leaving, traffic engineering, workload balancing, etc. When this happens, the OSS/BSS system conveys these changes to the NFV SC, which then automates the distribution of the new policy to the SAs.

This process is described in Figure 6.2.2-1.

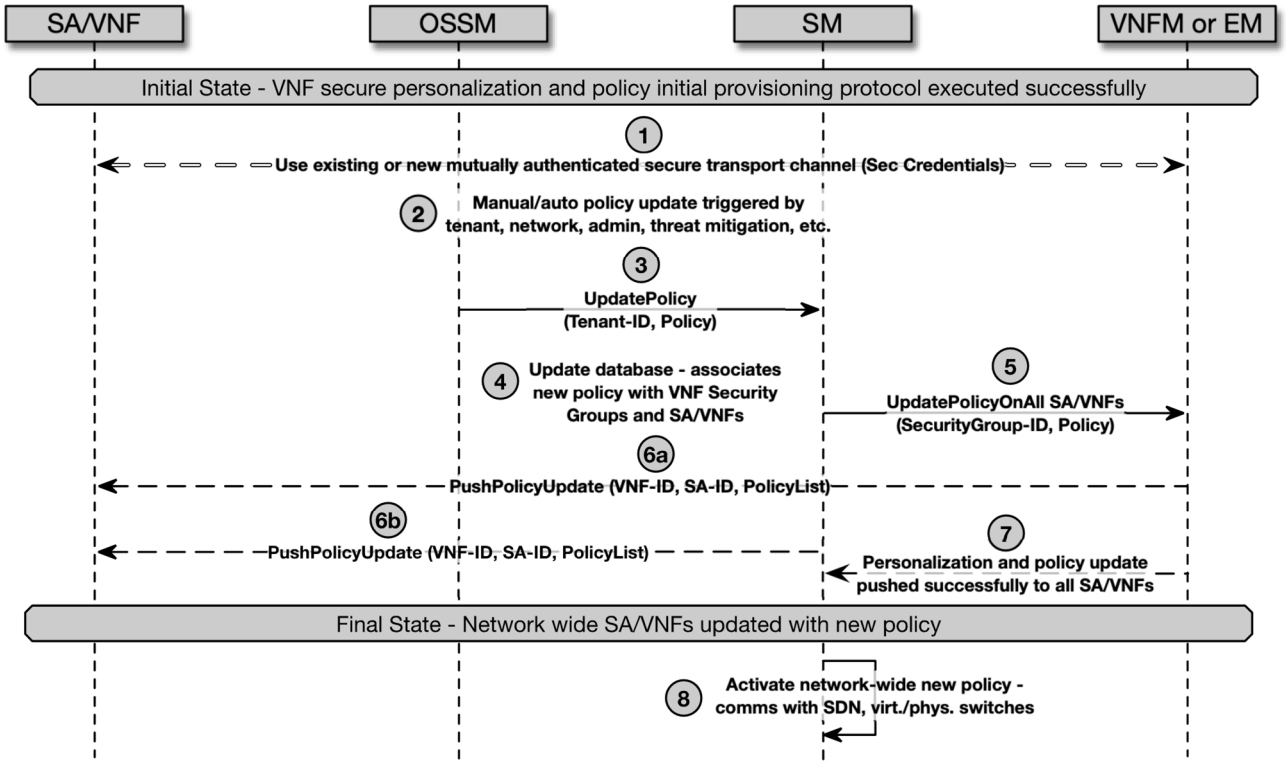

Figure 6.2.2-1: Secure VNF Personalization and Policy Protocol: Provisioning Update

It is assumed that the Provisioning and Policy Protocol - Initial Provisioning, as described above, has successfully been executed:

1) The SAs and the VNFM/EM have a secure, mutually authenticated channel, as described in earlier flows. The same channel may be used, or as per the NFV security deployment policy, a new secure channel may be set up for this procedure to ensure security credentials are still current and active.

2) A system event may trigger the Secure VNF Personalization and Policy update procedure. As described above this trigger may be add/remove tenants, response to a network or system threat, site security policy updates, etc. These system events may be automated, manual, time-driven, and/or dynamic.

3) The new security policy is added by the SM triggered by events from the OSS/BSS and OSSM, and security management analytics system.

4) The SM will update the new security policies, security groups, and other associated configurations into the appropriate SM database. Update of this secure database will follow the same security procedures as were used for Initial Provisioning steps of this protocol.

5) The SM will send a secure policy update to all affected VNFMs/EMs. The SM tracks the VNFMs/EMs for all SAs in all deployments. All messages in this step shall be signed and integrity-protected, such that the receiving entities only enforce policies that are verified as integral, and originating from the SM.

6) The VNFM(s) and EM(s) will push new policy updates to all affected SA/VNFs. This update will follow the same secure and reliable delivery that was used in the Initial Provisioning steps of this protocol. It is important that the SAs behaviour follow the prescribed security enforcement procedure. This means that based on the pre-set security policy, update priority or flags, an SA may immediately abandon its execution, gracefully exit, block all traffic, continue execution until another trigger, etc.

   (b) In some cases, in which the SAs associated with a VNF run in a different trust zone than the VNF, the SM may push policies directly to the SAs, without passing through the VNFM.

NOTE: It is a deployment decision whether to perform steps (a), (b), or (a and b), based on the purpose of the SAs (one policy to all, or bespoke policies to each SA). The VNFM(s) and EM(s) will report success back to the SM, similar to the Initial Provisioning steps of this protocol. At this point, the new security policies have been pushed into the entire network (including NFV, Physical and/or hybrid networks).

7) Based on the security policy of NFV deployment, the new policies are now activated by the SM. This follows the same procedure as specified in the initial provisioning steps of this protocol.

Successful execution of this update protocol should ensure a more secure, consistent new state of the NFV deployment, including mixed deployments with physical security functions.

6.3 NFV Security Management Principles

Secure management and administration of the NFVI and NFV-MANO is critical for the security of a virtualised network. The following describes the basic principles for such secure management:

a) Best practice for network administration is applied to the administration of the NFVI.

b) Administration of the NFVI is only available over mutually authenticated, encrypted and integrity-protected channels or APIs.

c) All channels or APIs are separated from each other and use separate credentials.

d) The number of privileged accounts for the NFVI is constrained to a minimal manageable number to meet the CSP's needs.

e) NFV-MANO and NFVI administrators do not have any privileged rights to other services within the CSP.

f) NFV-MANO and NFVI administrators are only provided with the privileges and accesses required to carry out their role.

g) NFV-MANO and NFVI administrators do not have access to workloads running within the virtualised environment.

h) NFV-MANO and NFVI administration access is limited to best practice configuration methods (e.g. authorized API calls).

i) Internal components within VNFs are not able to directly connect to entities or management functions outside of the network trust domain, except via interfaces that are explicitly part of the VNF security design.

j) NFV-MANO and NFVI administration is automated wherever possible.

k) Manual administration of the NFVI is by exception and raises a security alert.

l) Functions that manage the administration and security of the NFVI (e.g. MANO) are physically separate and do not run on the same NFVI as the NFs they manage.

m) Functions that support the administration and security of the NFVI are treated as security-critical functions.

7 Security Monitoring and Analysis

7.1 Introduction

The security of any system or network is ultimately determined by its weakest link, with attackers actively seeking the easiest points or interfaces to exploit, which may not always be directly associated with their ultimate target. In NFV environments, where network services and functions are dynamically created, updated, and terminated across multiple distributed NFVI Points of Presence, traditional security monitoring techniques are insufficient and will not scale.

Historically, in non-virtualised deployments, many interfaces between functional components were standardized and exposed, allowing for the use of traditional active or passive probes to monitor packets, flows, configurations, and metadata across management, data, and control planes. However, the advent of NFV introduces several critical challenges that necessitate a more sophisticated and integrated approach to security monitoring:

- **Dynamic Nature of VNFs**: The dynamic lifecycle management of Virtual Network Functions (VNFs) and Network Services (NSs) means that traditional physical security devices become less effective due to their lack of visibility into changes of virtualised functions, service chains, and inter-VM traffic.
- **Multi-Trust Domains**: NFV deployments inherently involve multiple trust domains, each potentially with distinct and separate monitoring responsibilities. For instance, an Infrastructure Security Monitoring domain, managed by the infrastructure provider, focuses on the NFVI, while a Tenant's administrative domain is confined to the Tenant's VNFs/VNF Components (VNFCs) and network. A Tenant does not possess knowledge of the underlying NFVI or other Tenants, necessitating clear separation and independence.
- **Layer-Crossing Attacks**: Focusing security defences or threat detection in isolation at the application layer, NFV-MANO layer, or NFVI layer will not enable the detection of threats or attacks that span across multiple layers. Similarly, security mitigations applied at one layer without considering the broader attack picture across all layers may lead to attack amplification or the forced use of compromised resources.
- **Increased Insider Attack Surface**: The NFV environment requires a significantly greater level of Security Monitoring than traditional deployments, partly due to the diminishing effectiveness of physical security devices and the need to minimize the insider attack surface. Attacks can occur both internally and externally to the virtualised network, impacting any layer.
- **Data Access and Privacy**: Passive security monitoring inspects real network traffic, raising privacy concerns when inspecting the payload directly. Meanwhile, active monitoring consumes network bandwidth and may disrupt operations.
- **Compromise Risks**: Security monitoring agents or the Security Monitoring & Analysis (SMA) function, if compromised, could give attackers control over the entire virtual network or NFV environment.

Given these complexities, it is necessary to consider Security Monitoring & Analysis using a system-wide architecture that can detect threats across user services, network management, virtual network functions, and underlying NFVI hardware. This allows for the application of system-wide security actions to minimize attack propagation and service disruption. The overall goal is to achieve an equivalent or higher level of security than existing non-virtualised networks, by continuously monitoring and managing the security posture of the NFV infrastructure and Network Services (NS).

8 End-to-end lifecycle

End-Of-Life (EOL) handling for VNFs/VNFCs shall be treated as part of the security-controlled lifecycle, not just an operational cleanup task. When a workload is shut down or destroyed, residual data, credentials, logs, and configuration artefacts can persist in storage, memory snapshots, backups, images, and orchestration metadata, and may be reintroduced or harvested by attackers if not properly controlled. When a VNFCI, whether VM or OS-Container-based reaches End-Of-Life (EOL), controls shall address compute, storage, identity, and management-plane artefacts.

Virtual disks, attached volumes, ephemeral storage, and container writable layers should be sanitized according to NIST SP 800-88 [7], using cryptographic erasure (key destruction) for encrypted volumes or secure wipe/purge/destroy for unencrypted media, with logs recording who/which system initiated and verified the action.

The NFV-MANO shall ensure that snapshots, templates, and backups containing the VNFCIs data or secrets are either brought under a retention policy (for legal/regulatory reasons) or sanitized or destroyed using equivalent methods and recorded in configuration management and ticketing systems. Secrets (e.g. API keys, tokens, SSH keys, certificates) associated with the VNFCI shall be revoked or rotated; this includes identifiers, service accounts, and issued certificates to avoid "orphaned" credentials from being reused in spoofed workloads. Within NFV-MANO, the VNFCI shall be removed from inventory systems so that it cannot be accidentally re-instantiated.

Monitoring, logging, and SIEM pipelines should retain security logs in accordance with the service providers policy while ensuring that local log files or debug traces on disks are sanitized alongside the data volumes.

SDN and micro-segmentation policies associated with the workload (e.g. security groups, firewall rules, container network policies) shall be cleaned up or reset to avoid leaving permissive rules tied to non-existent workloads. Configuration management databases and asset registers should be updated to reflect the VNFCI's retired status, in line with lifecycle and asset-management expectations.

Annex A (informative): Change history

Date	Version	Information about changes
Nov 2019	V0.0.0	First draft
Dec 2019	V0.0.1	Agreed output of NFVSEC#156
Apr 2020	V0.0.2	Agreed output of NFVSEC#162 (incl. revisions of NFVSEC(20)000027,31)
Apr 2020	V0.0.3	Agreed output of NFVSEC#163 (NFVSEC(20)000032r1)
Jul 2020	V0.0.4	Agreed output of NFVSEC#169 (NFVSEC(20)000041r6 - Includes base imports from SEC 013). Also includes drafting rule and formatting corrections
Nov 2020	V0.0.5	Agreed output of NFVSEC#175 (NFVSEC(20)000089r1)
Apr 2021	V0.0.6	Agreed output of NFVSEC#185 (NFVSEC(21)000017r2,18)
March 2022	V0.0.7	Agreed output of NFVSEC#205 (NFVSEC(22)000013r2)
Sept 2023	V0.0.8	Addition of contribution NFVSEC(23)000207r1
October 2024	V0.0.9	Addition of contribution NFVSEC(24)000195 following drafting session at NFVSEC#272
June 2025	V0.0.10	Addition of contribution NFVSEC(25)000042r1
December 2025	V0.0.11	Addition of contributions NFVSEC(25)000090, NFVSEC(25)000097 & NFVSEC(25)000161
December 2025	V0.0.12	Addition of contribution NFVSEC(25)000171r1
December 2025	V0.0.13	Addition of contribution NFVSEC(25)000182r1
March 2026	V0.0.14	Updates following Remote Consensus feedback

History

Version	Date	Status
V1.1.1	April 2026	Publication