ETSI GR NFV-REL 007 V1.1.1 (2017-09)

GROUP REPORT

Network Function Virtualisation (NFV); Reliability; Report on the resilience of NFV-MANO critical capabilities

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

DGR/NFV-REL007

Keywords

management, MANO, NFV, orchestration, resilience

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

Important notice

The present document can be downloaded from:
http://www.etsi.org/standards-search

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the only prevailing document is the print of the Portable Document Format (PDF) version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx

If you find errors in the present document, please send your comment to one of the following services:
https://portal.etsi.org/People/CommiteeSupportStaff.aspx

Copyright Notification

Contents

Intellectual Property Rights

Essential patents

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.

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.

Foreword

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

Modal verbs terminology

In the present document "**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 **investigates the building of** a resilient NFV-MANO functional block from the reliability/availability perspective. **In order to achieve this objective**, the present document:

1) Identifies critical NFV-MANO capabilities required to provide reliable services to the VNFs and the NSs.

2) Maps the resiliency requirements, e.g. established in ETSI GS NFV-REL 001 [i.2], with existing NFV-MANO capabilities as listed in up to Release 2 GSs.

3) Studies specific needs and constraints for the identified capabilities.

The work reports on possible mechanisms that enable high-availability within the different entities of NFV-MANO to render the identified capabilities dependable and trustworthy.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

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] ETSI GS NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[i.2] ETSI GS NFV-REL 001: "Network Functions Virtualisation (NFV); Resiliency Requirements".

[i.3] ETSI GS NFV-IFA 010 (V2.2.1): "Network Functions Virtualisation (NFV); Management and Orchestration; Functional requirements specification".

[i.4] ETSI GS NFV-IFA 005: "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Or-Vi reference point - Interface and Information Model Specification".

[i.5] ETSI GS NFV-IFA 006: "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Vi-Vnfm reference point - Interface and Information Model Specification".

[i.6] ETSI GS NFV-IFA 008: "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Ve-Vnfm reference point - Interface and Information Model Specification".

[i.7] ETSI GS NFV-IFA 011: "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; VNF Packaging Specification".

[i.8] ETSI GS NFV-SWA 001: "Network Functions Virtualisation (NFV); Virtual Network Functions Architecture".

[i.9] ETSI GS NFV-REL 003: "Network Functions Virtualisation (NFV); Reliability; Report on Models and Features for End-to-End Reliability".

[i.10] M. Chiosi et al.: "Network Functions Virtualisation - An Introduction, Benefits, Enablers, Challenges & Call for Action", SDN and OpenFlow World Congress, Darmstadt, Germany, October 22-24, 2012.

[i.11] M. Chiosi et al.: "Network Functions Virtualisation - Network Operator Perspectives on Industry Progress", SDN and OpenFlow World Congress, Frankfurt, Germany, October 15-17, 2013.

[i.12] M. Chiosi et al.: "Network Functions Virtualisation - Network Operator Perspectives on Industry Progress", SDN and OpenFlow World Congress, Dusseldorf, Germany, October 14-17, 2014.

[i.13] ETSI GS NFV-MAN 001: "Network Functions Virtualisation (NFV); Management and Orchestration".

[i.14] ETSI GS NFV-IFA 014: "Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Network Service Templates Specification".

[i.15] ETSI GS NFV-SEC 013: "Network Functions Virtualisation (NFV) Release 3; Security; Security Management and Monitoring Specification".

[i.16] ETSI GR NFV-IFA 021: "Network Functions Virtualisation (NFV); Management and Orchestration; Report on Management of NFV-MANO and Automated Deployment of EM and Other OSS Functions".

[i.17] IETF RFC 4412 (2006): "Communications Resource Priority for the Session Initiation Protocol (SIP)".

[i.18] IETF RFC 4594 (2006): "Configuration Guidelines for DiffServ Service Classes".

[i.19] IETF RFC 5865 (2010): "A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic".

[i.20] IETF RFC 4090 (2005): "Fast Reroute Extensions to RSVP-TE for LSP Tunnels".

[i.21] Recommendation ITU-T E.412 (2003): "Network management controls".

3 Abbreviations

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

(D)DoS	(Distributed) Denial of Service
API	Application Programming Interface
BFD	Bidirectional Forwarding Detection
BSS	Business Support System
COTS	Commercial Off The Shelf
CPU	Central Processor Unit
DF	Deployment Flavour
EM	Element Manager
FCAPS	Fault, Configuration, Accounting, Performance and Security
HA	High Availability
ID	Infrastructure Domain
IFA	(ETSI ISG NFV) Interface and Architecture (Working Group)
IMS	IP Multimedia Subsystem
IO	Input Output
IT	Information Technology
KPI	Key Performance Indicator
L2	Layer 2
L3	Layer 3
LSA	Link State Advertisement
MANO	Management and Orchestration
NF	Network Function
NFP	Network Forwarding Path
NFV	Network Functions Virtualisation

NFVI	NFV Infrastructure
NFVO	NFV Orchestrator
NS	Network Service
NSD	NS Descriptor
NSR	NS Record
OS	Operating System
OSPF	Open Shortest Path First
OSS	Operations Support System
P-CSCF	Proxy - Call Session Control Function
PNF	Physical Network Functions
SAL	Service Availability Level
S-CSCF	Serving - Call Session Control Function
SDN	Software Defined Networking
SDO	Standard Developing Organization
SLA	Service Level Agreement
TD	Tenant Domain
TTM	Time To Market
VDU	Virtualisation Deployment Unit
VIM	Virtualised Infrastructure Manager
VLD	Virtual Link Descriptor
VLR	Virtual Link Record
VM	Virtual Machine
VN	Virtual Network
VNF	Virtualised Network Function
VNFC	VNF Component
VNFD	VNF Descriptor
VNFFG	VNF Forwarding Graph
VNFFGD	VNFFG Descriptor
VNFFGR	VNFFG Record
VNFM	VNF Manager
VNFR	VNF Record
WG	Working Group

4 Introduction

Based on NFV resiliency requirements resulting from [i.2], the present document aims to verify that the current MANO capabilities as specified in [i.3] cover all these requirements. It also proposes, with respect to network services' reliability and availability, a classification of those resiliency requirements in two categories: requirements of normal importance and critical requirements. Through a mapping of the MANO capabilities vs. the resiliency requirements, it identifies the missing capabilities for some resiliency requirements. It finally lists some mechanisms which could be exploited to enhance the identified critical capabilities' resiliency.

To this end, the items processed are the followings:

- listing of the NFV resiliency characteristics;
- identification of resiliency requirements related to MANO;
- alignment of the MANO resiliency requirements within its functional blocks;
- resiliency requirements classification;
- mapping of MANO capabilities vs. resiliency requirements;
- identification of and recommendations for missing MANO capabilities following the mapping;
- proposals for resiliency mechanisms related to the identified critical capabilities.

5 NFV resiliency and its recommendations towards MANO

5.1 Resiliency and its application to NFV

Resiliency - as defined in [i.1] - provides the ability "*to limit disruption and return to normal or at a minimum acceptable service delivery level in the face of a fault, failure, or an event that disrupts the normal operation*".

The definition above mentions fault as distinct from failure. [i.2] has introduced the fault as leading to a failure chain. In a PNF environment, redundancy, the core mechanism of dependability, can be exploited in the following way. In a redundant system composed of two servers, one active and one standby, if the active fails (i.e. a fault occurs), the service is still rendered thanks to the standby one which takes over. It is noteworthy to mention that, in some cases (e.g. the standby server is less performant), the service may be degraded after the failover. Moreover, if the failed server is not replaced in a timely fashion and the new active one also fails, the failure is observed at the service level.

The resiliency engineering for a network service consists of the following phases:

- definition of the resiliency requirements, e.g. availability which can be done through different service availability levels and grades of the service [i.2];
- design for resiliency - this important step mainly targets failure prevention, e.g. with the use of a redundant architecture;
- operations - once the network service is deployed, three other tasks help to maintain the service resiliency:
    - continuous monitoring, e.g. using health check, watchdog mechanisms;
    - detection of abnormal situations, e.g. threshold exceeded;
    - potential actions based on anomalous events - with the use of, e.g. fault correlation, root cause analysis:
        - launch corrective actions, e.g. prevent failures to happen;
        - prevent propagation of a breakdown, i.e. failure containment - needless to say, as not all failures are equal, the study of failure severity, e.g. number of users impacted, is generally done in the design phase;
        - diagnose that an entity is deteriorating, i.e. failure prediction.

During the operations, the failures are diverse: they can come from the infrastructure level, the VNF layer, or can concern the MANO software (see annex A for a reminder of MANO). As such, the MANO components NFVO, VNFM and VIM play critical roles in ensuring VNFs resiliency from their instantiation through their operation and recovery from failures. With respect to some life cycle management tasks (e.g. instantiation, scaling, migration), MANO is responsible for adhering to affinity and anti-affinity rules to ensure there is no single point of failure. Actually, the application of redundancy in the NFV framework implies the creation of secondary VNF instances (for those VNF considered as critical), but this duplication is useless if both instances run on the same server or NFVI-PoP (i.e. a single point of failure). Anti-affinity rules thus impose the use of different infrastructures for those instances. On the other hand, performance constraints may need that different VNF instances run in the same location (affinity rules).

Based on the type of failure, the MANO components provide the appropriate recovery or remediation such as VNF re-instantiation/re-creation, VNF scaling (i.e. adding a VNF instance to provide more capacity so that the service is not degraded), VNF migration (i.e. moving the VNF to another server during maintenance or hardware failures), or failure containment (i.e. keep it from propagating further and negatively impact other VNFs), while ensuring service continuity.

NOTE: For the sake of clarity, the present document has adopted a simplistic picture of VNFs as composed of a single VM wherever it was possible. It is explicitly noted when this is not the case.

The MANO components, particularly the NFVO, can have an impact beyond a single VNF, as they are able to provide orchestration within a service chain or between geographically separated VNFs or between multi-vendor VNFs. The next clause will provide details on the recommendations and the role of the individual MANO components.

5.2 Resiliency recommendations towards MANO

5.2.1 MANO-related resiliency requirements classification

Annex B shows the resiliency requirements identified in [i.2]. Most of these items are related to MANO and are thus analysed in this clause. The ones which have been excluded from the analysis are treated in clause 5.2.2.

This clause classifies the [i.2] requirements related to MANO: the classification is based on the pertinent MANO component i.e. NFVO, VNFM, VIM. The requirements which involve more than one MANO component are gathered under the MANO umbrella.

Note that any "item" not expressed as such in [i.2], but implied by a requirement, is set in italic. As an example, the requirement 10.8.13 in [i.2] states that "The VNF shall inform the VNFM in the event that it cannot remediate a failure that results in a degradation of the overall capacity of the VNF". This "item" thus states that "The VNFM receives remediation failure information from VNFs and acts accordingly" and it is listed in italic in reference to the requirement 10.8.13 in [i.2]. The resiliency keywords and main actions are underlined for ease of reading.

NFVO

The NFV Orchestrator has two main responsibilities:

(i) orchestration of NFVI resources across multiple VIMs;

(ii) lifecycle management of network services.

The NFVO is also responsible for activities ensuring that the VNFs and the network services they support meet any reliability and availability requirements by passing on resiliency requirement data or failure information to other MANO components and/or taking corrective actions on its own:

- interpret **resiliency requirement data** contained in NSD and VNFD (e.g. NFVI geo-redundancy) and transfer it to VIM and potentially to VNFM as appropriate (see [i.2], Req. 4.2.2);
- take **corrective actions** for other VNFs of a service chain in case a VNF failure is detected (see [i.2], Req. 10.8.18).

VNFM

The VNF Manager supports operations of VNF instances lifecycle management, e.g. health monitoring and recovery upon failure, which includes the following:

- provide an **interface** to VNFs for health checking (see [i.2], Req. 10.8.17);
- **restart** a VNF having exceeded its health check response timeout (see [i.2], Req. 10.8.20);
- *receive remediation **failure information** from VNFs and act accordingly (see [i.2], Req. 10.8.13);*
- **request additional resources** in case VNF encounters failures (see [i.2], Req. 10.8.14).

NOTE: In case of VNF scaling, the VNF may consist of more than one VM.

VIM

The Virtualised Infrastructure Manager (VIM) is responsible for controlling and managing the NFVI compute, storage and network resources. In that capacity if the VIM detects a possible impending resource shortage it can initiate pre-emptive actions to provide additional resources before the shortage occurs, and/or launch admission mechanisms to limit the access to resources. If an NFVI-related failure occurs, it can take corrective actions such as migrating VMs to other servers:

- receive **resiliency requirement data** derived from NSD and VNFD (e.g. NFVI geo-redundancy) sent by NFVO and act accordingly (see [i.2], Req. 4.2.2);
- apply anti-affinity policies for preventing **single point of failure** (see [i.2], Req. 9.5.1 and 9.5.2);
- hardware resource **monitoring** (see [i.2], Req. 5.4.10, 10.8.5 and 10.8.11);

- get real-time **resource** (e.g. CPU, memory) **usage information** and act accordingly (see [i.2], Req. 9.3.1);
- get **health information** from the infrastructure and act accordingly (see [i.2], Req. 9.3.6);

NOTE: VNFs are not able to report to the VIM nor the VIM is aware of the VNFs.

- get indication of hardware and environmental events from the NFVI layer and **migrate** pro-actively the VMs impacted by such faulty events (see [i.2], Req. 10.8.15);
- receive **failure detection information** from the hypervisor (and/or its underlying host) and act accordingly (see [i.2], Req. 10.8.4);
- take **corrective actions** following NFVI failures (see [i.2], Req. 10.8.16);
- assign dedicated resources to a VM for failure **containment** (i.e. propagation avoidance) (see [i.2], Req. 9.2.1):
  - in conjunction with the hypervisor (see [i.2], Req. 9.2.2);
  - by defining maximum virtual resources allocated to VMs, e.g. for attack containment (see [i.2], Req. 9.2.3).
- transport network (e.g. virtual network) redundancy (see [i.2], Req. 9.5.3).

MANO

The MANO components work together to ensure service availability and continuity for VNFs and adherence to the terms laid out in SLAs. In support of life cycle management operations, the MANO components make sure that the operations are carried out successfully, or provide support for the recovery of any failed actions. During the instantiation of VNFs, anti-affinity rules have to be followed such that there is no single point of failure. This rule is enforced by MANO:

- no **single point of failure** (see [i.2], Req. 4.2.13), e.g. redundant deployment (see [i.2], Req. 10.8.21) such as geo-redundancy (see [i.2], Req. 10.8.22) → NFVO, VNFM, VIM.

The MANO components that are configured to manage multiple VNFs have to be able to manage each of them separately and, in each case, according to their requirements:

- resiliency for **multi-vendor** environment (see [i.2], Req. 4.2.14) → NFVO, VNFM, VIM.

When priorities are assigned to VNFs, the MANO components are responsible for complying with those priorities by ensuring that when resources are reaching capacity limits, sufficient resources are allocated as per those priorities (i.e. VNFs with higher priorities are satisfied first even when VNFs of lower priorities are only partially or not satisfied at all):

- VNFs **priority handling** for resource allocation (see [i.2], Req. 5.4.7) → NFVO, VNFM;
- support of service availability levels and grades of service (e.g. following the SLA expressed in VNFD, NFVO/VNFM will order VIM to allocate the appropriate resources) (see [i.2], Req. 7.3.1, 7.3.2 and 7.3.3) → NFVO, VNFM, VIM;
- *analysis of load information (from VNFs, hypervisor, …), e.g. determine if scaling or migration or … is needed (see [i.2], Req. 9.4.2)* → VNFM, VIM;
- support **watchdog** functionality whose characteristics include reliable triggering mechanisms and low latency notifications (see [i.2], Req. 10.8.10) → VNFM, VIM;
- support of a means to indicate VNFs' ability to support external health check (see [i.2], Req. 10.8.19) → NFVO, VNFM;
- ensure **service continuity** (at session/service establishment and during service relocation) (see [i.2] Req. 4.2.11) by respecting the acceptable **disruption time** (see [i.2], Req. 5.4.2) defined in the SLA (see [i.2], Req. 5.4.3) → NFVO, VNFM, VIM;

- support (receive, analyse) **failure notification** at run time and act accordingly (see [i.2], Req. 4.2.5), e.g. data from NFVI (see [i.2], Req. 4.2.6) → NFVO, VNFM, VIM;
- support of VNF **scaling** → NFVO, VNFM:
  - on site and/or offsite (see [i.2], Req. 7.3.8);
  - within the VNF's limits, limit the number of users in order to minimize the failures impact (see [i.2] Req. 4.2.10);
  - including efficient load distribution/balancing (see [i.2], Req. 5.4.9).
- appropriate handling of malicious traffic to avoid excessive resource consumption (see [i.2], Req. 7.3.9) → NFVO, VNFM, VIM;
- support of VNF **migration**, e.g. following a hardware, a software failure, or resource shortage (see [i.2] Req. 5.4.1) → NFVO, VNFM, VIM (support of VM migration when requested):
  - on site and/or offsite (see [i.2], Req. 7.3.8);
  - maintaining pre-migration communication between VNFs (see [i.2], Req. 5.4.4);
  - the pre-migration process (which may be activated during the initial instantiation of the VNF) includes the availability of VNFD (to retrieve deployment conditions and/or VNF constraints) (see [i.2], Req. 5.4.6).
- support of automated fail-over on the NFVI level (see [i.2], Req. 4.2.12) → NFVO, VNFM, VIM.

NOTE 1: Automated fail-over of compute resources hosting stateful VNFs with internal redundancy is undesirable as it may result in fail-overs at both NFVI and VNF levels resulting in a conflicting situation.

- **pre-emption** capability (VNFM) and support of this capability (VIM), e.g. in case of resource shortage (see [i.2], Req. 5.4.8) → VNFM, VIM;
- overload control to be considered in the implementation (see [i.2], Req. 9.4.3) → NFVO, VNFM, VIM;
- *handling of fields error **severity** (see [i.2], Req. 9.3.7, 10.8.8) and failure cause (see [i.2], Req. 9.3.8) received* → NFVO, VNFM, VIM;
- provide visibility (into specific classes of hardware failures) to the VNFs (see [i.2], Req. 10.8.2) → NFVO, VNFM, VIM:
  - send hardware failure notifications (see [i.2], Req. 10.8.3) → NFVO, VNFM, VIM.

NOTE 2: The exploitation of hardware failure information by the VNFs is for further study.

- enabling to provide a **degraded service** under abnormal conditions (see [i.2], Req. 7.3.4) → NFVO, VNFM, VIM (support of virtualised resource management when requested);
- **failure prediction** framework:
  - false alarm filtering to avoid triggering unnecessary prevention procedure for anomaly alarming (see [i.2], Req. 9.3.3) → VNFM, VIM (involved as alarm manager);
  - trend identification, period or seasonal variations, and randomness analysis of data on resource usage to predict unhealthy state (e.g. resource exhausting) (see [i.2], Req. 9.3.4) → VNFM, VIM;
  - diagnose/verify which entity is progressing towards a failure and which VNFs might be affected (see [i.2], Req. 9.3.5) → VNFM, VIM (support of fault management related to virtualised resources).
- distributed **fault correlation** processing: avoiding propagating large number of failure notifications to a centralized entity by sending locally correlated reports (see [i.2], Req. 10.8.7) → NFVO, VNFM, VIM;
- automatic VNF **re-instantiation** following a failure, e.g. VM failure (see [i.2], Req. 4.2.4) → NFVO, VNFM:
  - within specified time intervals (see [i.2], Req. 7.3.7);

  - retrieve VNF state information (previously stored) before VNF recreation, e.g. after a failure (see [i.2], Req. 4.2.9).

- service **admission control** capability (see [i.2], Req. 7.3.5) → NFVO, VNFM, VIM;

- MANO failure:

  - not affecting existing VNF instances and no change in existing VNFs service level (see [i.2] Req. 10.8.23) → NFVO, VNFM, VIM;

  - integrity preserved (e.g. during orchestration) thanks to the atomicity of internal state data change (see [i.2], Req. 10.8.25) → NFVO, VNFM, VIM;

  - recover the state of the environment by means of restoring persistent storage and audit the environment for determining its state, without interruption to the in service VNF instances (see [i.2], Req. 10.8.24) → NFVO, VNFM, VIM.

Figure 5.1 shows the resiliency requirements analysed in this clause and their positioning within the MANO functional blocks.

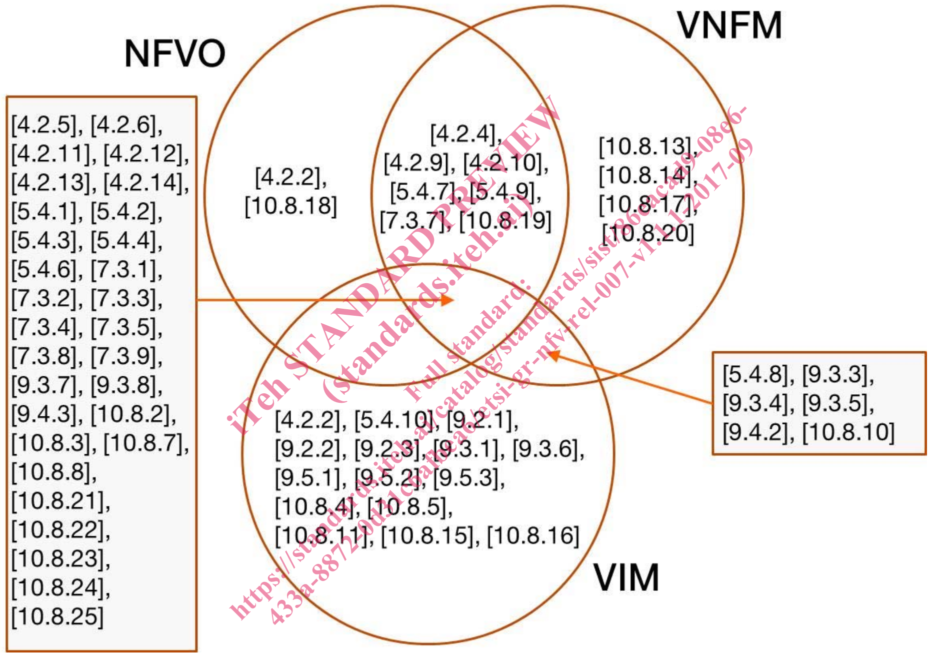

Figure 5.1: Positioning of resiliency requirements within the MANO functional blocks

Table 5.1 shows the resiliency requirements with respect to the phases of resiliency engineering in the 'Life cycle vs. MANO functional blocks' matrix. Note that the failure prediction and fault correlation tasks are regrouped under 'Analysis'.