ETSI GR NFV-EVE 021 V5.1.1 (2023-09)

GROUP REPORT

Network Functions Virtualisation (NFV) Release 5; Evolution and Ecosystem; Report on energy efficiency aspects for NFV

Reference

DGR/NFV-EVE021

Keywords

energy efficiency, management, NFV

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 Web server (https://ipr.etsi.org/).

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™** and **LTE™** 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 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 aspects of NFV that have an impact on power consumption and energy efficiency, including but not limited to:

- the design and runtime characteristics of the NFV framework, including VNF, NFVI, NFV-MANO, etc. which can influence energy consumption;
- the deployment configuration of NFV-MANO systems and their optimization in various scenarios such as small footprint but highly distributed deployments; and
- the mechanisms that enable the collection of energy related information, smart energy usage and decision making in NFV considering NFV capabilities such as orchestration and automation.

The present document also documents potential solutions, and where applicable, it also provides recommendations for enhancements to the NFV architectural framework and its functionality aiming to provide further support to address energy efficiency objectives.

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

[i.2] ITU-T Development Sector: "Measuring digital development: Facts and figures 2020".

NOTE: Available at https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2020.pdf. Last access on 12.10.21.

[i.3] Wikipedia: "Internet traffic".

NOTE: Available at https://en.wikipedia.org/wiki/Internet_traffic. Last access on 12.10.21.

[i.4] International Energy Agency: "Data Centres and Data Transmission Networks," 2023, Paris.

NOTE: Available at https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks Last access on 15.09.23.

[i.5] GSMA Future Networks: "Energy Efficiency: An Overview".

NOTE: Available at https://www.gsma.com/futurenetworks/wiki/energy-efficiency-2/#. Last access on 12.10.21.

[i.6] European Commission: "2030 climate & energy framework".

NOTE: Available at https://ec.europa.eu/clima/eu-action/climate-strategies-targets/2030-climate-energy-framework_en. Last access on 12.10.21.

[i.7] United Nations: "Sustainable Development Goals: goal 7 Energy".

NOTE: Available at https://sdgs.un.org/goals/goal7. Last access on 12.10.21.

[i.8] ETSI EN 303 471 (V1.1.1): "Environmental Engineering (EE); Energy Efficiency measurement methodology and metrics for Network Function Virtualisation (NFV)".

[i.9] ETSI GR NFV 001: "Network Functions Virtualisation (NFV); Use Cases".

[i.10] ETSI GS NFV 004: "Network Functions Virtualisation (NFV); Virtualisation Requirements".

[i.11] M. Chiosi et al.: "Network Functions Virtualisation - Introductory White Paper".

NOTE: Available at https://portal.etsi.org/NFV/NFV_White_Paper.pdf. Last access on 26.11.21.

[i.12] 3GPP TR 21.866 (V15.0.0): "3GPP; Technical Specification Group Services and System Aspects; Study on Energy Efficiency Aspects of 3GPP Standards (Release 15)".

[i.13] ETSI TR 132 972 (V17.0.0): "5G; Telecommunication management;Study on system and functional aspects of energy efficiency in 5G networks (3GPP TR 32.972 version 17.0.0 Release 17)".

[i.14] ETSI TS 128 304 (V17.0.0): "Digital cellular telecommunications system (Phase 2+) (GSM); Universal Mobile Telecommunications System (UMTS); LTE; 5G; Control and monitoring of Power, Energy and Environmental (PEE) parameters Integration Reference Point (IRP); Requirements (3GPP TS 28.304 version 17.0.0 Release 17)".

[i.15] ETSI TS 128 305 (V17.0.0): "Digital cellular telecommunications system (Phase 2+) (GSM); Universal Mobile Telecommunications System (UMTS); LTE; 5G; Control and monitoring of Power, Energy and Environmental (PEE) parameters Integration Reference Point (IRP); Information Service (IS) (3GPP TS 28.305 version 17.0.0 Release 17)".

[i.16] ETSI TS 128 306 (V17.0.0): "Digital cellular telecommunications system (Phase 2+) (GSM); Universal Mobile Telecommunications System (UMTS); LTE; 5G; Control and monitoring of Power, Energy and Environmental (PEE) parameters Integration Reference Point (IRP); Solution Set (SS) definitions (3GPP TS 28.306 version 17.0.0 Release 17)".

[i.17] ETSI TS 128 310 (V17.5.0): "LTE; 5G; Management and orchestration; Energy efficiency of 5G (3GPP TS 28.310 version 17.5.0 Release 17)".

[i.18] 3GPP TR 28.813 (V1.2.0): "3GPP; Technical Specification Group SA; Telecommunication management; Study on new aspects of Energy Efficiency (EE) for 5G (Release 17)".

[i.19] ETSI TS 128 552 (V17.10.0): "5G; Management and orchestration; 5G performance measurements (3GPP TS 28.552 version 17.10.0 Release 17)".

[i.20] ETSI EN 303 470: "Environmental Engineering (EE); Energy Efficiency measurement methodology and metrics for servers".

[i.21] ETSI EN 305 200-2-2: "Access, Terminals, Transmission and Multiplexing (ATTM); Energy management; Operational infrastructures; Global KPIs; Part 2: Specific requirements; Sub-part 2: Fixed broadband access networks".

[i.22] ETSI EN 305 200-2-3: "Access, Terminals, Transmission and Multiplexing (ATTM); Energy management; Operational infrastructures; Global KPIs; Part 2: Specific requirements; Sub-part 3: Mobile broadband access networks".

[i.23] Recommendation ITU-T L.1316 (11/2019): "Series L: Environment and ICTS, climate change, e-waste, energy efficiency; Construction, installation and protection of cables and other elements of outside plant; Energy efficiency framework".

[i.24] Recommendation ITU-T L.1330 (03/2015): " Series L: Construction, installation and protection of cables and other elements of outside plant; Energy efficiency measurement and metrics for telecommunication networks".

[i.25] Recommendation ITU-T L.1331 (01/2022): " Series L: Environment and ICTS, climate change, e-waste, energy efficiency: Construction, installation and protection of cables and other elements of outside plant; Energy efficiency, smart energy and green data centers; Assessment of mobile network energy efficiency".

[i.26] Recommendation ITU-T L.1332 (01/2018): "Series L: Environment and ICTS, climate change, e-waste, energy efficiency; Construction, installation and protection of cables and other elements of outside plant; Total network infrastructure energy efficiency metrics".

[i.27] Recommendation ITU-T L.1350 (10/2016): "Series L: Environment and ICTS, climate change, e-waste, energy efficiency; Construction, installation and protection of cables and other elements of outside plant; Energy efficiency metrics of a base station site".

[i.28] Recommendation ITU-T L.1361 (11/2018): "Series L: Environment and ICTS, climate change, e-waste, energy efficiency; Construction, installation and protection of cables and other elements of outside plant; Measurement method for energy efficiency of network functions virtualization".

[i.29] DMTF Redfish DSP2051 (1.0.0): "Redfish Telemetry White Paper", published 2020-03-07.

[i.30] ATIS-0600015.09.2015: "Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting of Base Station Metrics".

[i.31] ETSI ES 203 237: "Environmental Engineering (EE); Green Abstraction Layer (GAL); Power management capabilities of the future energy telecommunication fixed network nodes".

[i.32] ETSI ES 203 539: "Environmental Engineering (EE); Measurement method for energy efficiency of Network Functions Virtualisation (NFV) in laboratory environment".

[i.33] ETSI ES 203 682: "Environmental Engineering (EE); Green Abstraction Layer (GAL); Power management capabilities of the future energy telecommunication fixed network nodes; Enhanced Interface for power management in Network Function Virtualisation (NFV) environments".

[i.34] DMTF Redfish DSP2046 (2021.4): "Redfish Resource and Schema Guide", published 2021-12-02.

[i.35] DMTF Redfish DSP0266 (1.15) "Redfish Specification", published 2021-12-02.

[i.36] UEFI Forum ACPI Specification (Version 6.4): "Advanced Configuration and Power Interface (ACPI) Specification", published January 2021.

NOTE: Available at https://uefi.org/sites/default/files/resources/ACPI_Spec_6_4_Jan22.pdf. Last access on 26.05.22.

[i.37] ETSI TS 128 554 (V17.10.0): "5G; Management and orchestration; 5G end to end Key Performance Indicators (KPI) (3GPP TS 28.554 version 17.10.0 Release 17)".

[i.38] ETSI GR NFV-IFA 041 (V4.1.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Report on enabling autonomous management in NFV-MANO".

[i.39] Libvirt virtualization API.

NOTE: Available at https://libvirt.org/index.html. Last access on 21.11.22.

[i.40] Libvirt virtualization API: "Module libvirt-domain from libvirt".

NOTE: Available at https://libvirt.org/html/libvirt-libvirt-domain.html. Last access on 21.11.22.

[i.41] Kubernetes® Power Manager.

NOTE: Available at https://github.com/intel/kubernetes-power-manager. Last access on 21.11.22.

[i.42] ETSI GS NFV-IFA 010 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Functional requirements specification".

[i.43] ETSI GS NFV-IFA 005 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Or-Vi reference point - Interface and Information Model Specification".

[i.44] ETSI GS NFV-IFA 006 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Vi-Vnfm reference point - Interface and Information Model Specification".

[i.45] ETSI GS NFV-IFA 040 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Requirements for service interfaces and object model for OS container management and orchestration specification".

[i.46] ETSI GS NFV-IFA 007 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Or-Vnfm reference point - Interface and Information Model Specification".

[i.47] ETSI GR NFV-EVE 022 (V5.1.1): "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on VNF configuration".

[i.48] ETSI GS NFV-IFA 011 (V4.3.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; VNF Descriptor and Packaging Specification".

[i.49] ETSI GR NFV-EVE 019 (V4.1.1): "Network Functions Virtualisation (NFV) Release 4; Architectural Framework; Report on VNF generic OAM functions".

[i.50] ETSI GS NFV-IFA 049 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Architectural Framework; VNF generic OAM functions specification".

[i.51] ETSI GS NFV-IFA 036 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Requirements for service interfaces and object model for container cluster management and orchestration specification".

[i.52] ETSI GR NFV-IFA 023 (V3.1.1): "Network Functions Virtualisation (NFV); Management and Orchestration; Report on Policy Management in MANO; Release 3".

[i.53] ETSI GS NFV-IFA 048 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Policy Information Model Specification".

[i.54] ETSI GS NFV-IFA 031 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Requirements and interfaces specification for management of NFV-MANO".

[i.55] ETSI GS NFV-SOL 009 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Protocols and Data Models; RESTful protocols specification for the management of NFV-MANO".

[i.56] ETSI GR NFV-IFA 042 (V4.1.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Report on policy information and data models for NFV-MANO".

[i.57] ETSI GS NFV-IFA 027 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Performance Measurements Specification".

[i.58] ETSI GS NFV-IFA 008 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Ve-Vnfm reference point - Interface and Information Model Specification".

[i.59] ETSI GS NFV-IFA 013 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Os-Ma-nfvo reference point - Interface and Information Model Specification".

[i.60] SERT® Suite: "Server Efficiency Rating Tool", Standard Performance Evaluation Corporation.

NOTE: Available at https://www.spec.org/sert/.

[i.61] ETSI EN 303 804 (V0.0.3): "Environmental Engineering (EE); Energy efficiency metrics and measurement methods for data storage equipment".

[i.62] Linux® manual page: "time(7) - Linux manual page".

NOTE: Available at https://man7.org/linux/man-pages/man7/time.7.html.

[i.63] Scaphandre.

NOTE: Available at https://github.com/hubblo-org/scaphandre.

[i.64] Kepler: "Kubernetes Efficient Power Level Exporter".

NOTE: Available at https://sustainable-computing.io/.

[i.65] ETSI GS NFV-IFA 053: "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Requirements and interface specification for Physical Infrastructure Management".

[i.66] Recommendation ITU-T L.1310 (09/2020): "Series L: Environment and ICTS, climate change, e-waste, energy efficiency; Construction, installation and protection of cables and other elements of outside plant; Energy efficiency metrics and measurement methods for telecommunication equipment".

[i.67] ATIS-0600015.03.2013: "Energy Efficiency for Telecommunications Equipment: Methodology for Measurement and Reporting for Router and Ethernet Switch Products".

[i.68] Commission Regulation (EU) 2019/424 of 15 March 2019 laying down ecodesign requirements for servers and data storage products pursuant to Directive 2009/125/EC of the European Parliament and of the Council and amending Commission Regulation (EU) No 617/2013.

NOTE: Available at https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32019R0424.

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:

NOTE: A term defined in the present document takes precedence over the definition of the same term, if any, in ETSI GR NFV 003 [i.1].

infrastructure power consumption: power consumption of physical resources and subsystems in the NFVI.

EXAMPLE: Power consumption of physical compute servers and NFVI subsystems such as cooling systems, racks mounting the compute, storage and network resources.

NFV power consumption: power consumption associated to objects managed by NFV-MANO.

NOTE: Objects managed by NFV-MANO include virtualised resources, containerized workloads, VNF and NS.

EXAMPLE: Power consumption incurred by and associated to a VNF instance.

3.2 Symbols

Void.

3.3 Abbreviations

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

ACPI Advanced Configuration and Power Interface
BMC Baseboard Management Controller
DMTF Distributed Management Task Force

ECA	Event-Condition-Action
EE	Energy Efficiency
EPA	Environmental Protection Agency
IEA	International Energy Agency
IRP	Integration Reference Point
IS	Information Service
ITU-T	ITU Telecommunication standardization sector
MDA	Management Data Analytics
NMI	Non-Maskable Interrupt
NUMA	Non-Uniform Memory Access
OPEX	Operational Expenditure
PAP	Policy Administration Point
PCIe	Peripheral Component Interconnect express
PDU	Power Distribution Unit
PF	Policy Function
PEE	Power, Energy and Environmental
PIM	Physical Infrastructure Management
RAN	Radio Access Network
SDG	Sustainable Development Goal
SERT	Server Efficiency Rating Tool
SPEC	Standard Performance Evaluation Corporation
TSG SA	Technical Specification Group Service and System Aspects
UEFI	Unified Extensible Firmware Interface

4 Introduction and overview

4.1 Introduction to energy efficiency and global initiatives

The digitization of the society and communication among people, things and machines and the demand for data and digital services relies on distributed and extensive deployed worldwide networks and compute resources. Since 2010, the number of Internet users worldwide has doubled [i.2]. It is estimated that from 1990 until 2017 the IP traffic has increased from a mere 0,001 PB/month to around 122 000 PB/month [i.3], about 12-fold more traffic from 2010 at around 30 % increase per year. These trends are driving exponential growth in demand of data and communication services. To cope with the increase of traffic, networks and the Internet have grown tremendously with many more network and compute devices being installed and operated year over year, which naturally account for the consumption of more energy.

However, computer technology is changing and evolving fast, and becoming more efficient regarding the amount of energy consumed per processed/transferred bit. Once a network is operating, the amount of power it uses does not increase linearly with the amount of data that are being transmitted. Furthermore, the energy used in the data centre is not necessarily directly proportional to the amount of data transferred because it depends on the amount of processing that the servers need to do when processing each request. According to International Energy Agency's (IEA) summary report on "Data Centres and Data Transmission Networks" [i.4], for every bit of data that is forwarded on the network from data centres to end users, another five bits of data are transmitted within and among data centres. The same report shows that data centres account for about 1 % of global electricity demand, and that if current trends in the efficiency of hardware and data centre infrastructure can be maintained, the energy consumption could remain nearly flat despite a 60 % increase in service demand.

In spite of the optimistic estimates that newer and more energy efficient technologies become a reality, still energy consumption is one of the highest operating costs that many network operators face. According to GSMA Future Networks report on "Energy Efficiency: an Overview" [i.5], energy consumption constitutes between 20 % and 40 % of Operational Expenditure (OPEX) for the network operator.

It is also increasingly important for network operators to address energy consumption due to the climate change and sustainability considerations and to adhere to national and worldwide mandates. For instance, the EU's 2030 climate and energy framework [i.6] stipulates that at least a 40 % cut in greenhouse gas emissions is expected to be achieved by 2030 compared to the levels in 1990. The scale of actions regarding energy consumption is also considered by the United Nations (UN) as part of the Sustainable Development Goals (SDG). Target #7 on "Energy" aims to ensure access to affordable, reliable, sustainable and modern energy for all, and more precisely, target #7.3 aims that by 2030, the global rate of improvement in energy efficiency is expected to double [i.7].

Driving towards these targets, network operators and equipment providers can increase the network energy efficiency and decrease energy consumption costs with actions falling into three major groups:

- increase the use of alternative energy sources;
- optimize the network load towards policies reducing the energy consumption; and
- invest on and deploy infrastructure making use of more energy efficient network, storage and compute technologies.

In the present document, energy efficiency is considered as specified in ETSI EN 303 471 [i.8], which states that energy efficiency is the relation between the useful output (e.g. telecom service, etc.) and energy consumption.

4.2 NFV and energy efficiency

Aspects about energy efficiency have been considered since long in telecommunication network environments. As introduced in clause 4.1, energy consumption constitutes a major factor in the total OPEX amount of a network operator. An expectation for deploying telecommunication networks based on NFV is to reduce energy consumption based on the key principles indicated in clause 4.1 (i.e. optimization of the network deployment and operation to reduce the energy consumption and make use of infrastructure that is more energy efficient).

As introduced in the first "NFV white paper" published in 2012 [i.11], reduction of energy consumption or increasing energy savings is a key driver for the virtualization of telecommunication networks, which could be achieved by exploiting power management features in commercial-off-the-shelf servers and storage, as well as with workload consolidation and location optimization. Indeed, the virtualization of the telecommunication networks results in more possibilities of disaggregation of specific network, storage and compute resources, thus resulting in a different distribution of power consumption of the network which can be further optimized.

ETSI GR NFV 001 [i.9] describes diverse telecommunication network use cases in the context of NFV, and some of them introduce relevant aspects of energy efficiency and consumption, such as:

- Use case #17 on "Network Slicing" in clause 6.2.7 of ETSI GR NFV 001 [i.9] indicates that energy saving key performance indicators (KPI) could be associated to network slices, which from a resource fulfilment perspective can be deployed as one or more Network Services (NS).
- Use case #6 on "virtualization of mobile base station" in clause 6.3.2 of ETSI GR NFV 001 [i.9] illustrates the expectation to lower footprint and energy consumption coming from dynamic resource allocation and traffic load balancing in the virtualization of RAN.
- Use case #8 on "fulfilment of virtual content delivery network (vCDN)" in clause 6.3.4 of ETSI GR NFV 001 [i.9] refers to the disadvantages of non-virtualised CDN, such that the capacity of the dedicated hardware appliances are designed for peak hours and during certain periods of time these are mainly unused but still consuming energy and generating heat.
- Use case #9 on "fixed access network functions virtualization" in clause 6.3.5 of ETSI GR NFV 001 [i.9] remarks the specifics of fixed access network systems which are typically deployed in remote nodes located in the street or in buildings, and which are expected to include novel powering mechanisms to make them more energy efficient. In this context, access network virtualization would be expected to further improve the power consumption of remote nodes in two main areas: at the link level, by optimizing the complex trade-offs involved with dynamic adaptation of different operational parameters, service levels, traffic, delay, etc., and at the network level, where energy can be reduced by selectively forcing nodes and/or links to switch to "sleep" modes.

Energy efficiency related business-level requirements for the NFV framework were also specified in ETSI GS NFV 004 [i.10]. The requirements specified in clause 5.10 of ETSI GS NFV 004 [i.10] indicate the expectation that an NFV framework could bring in reducing the energy consumption of large-scale network infrastructures. Three specific requirements were specified with the following scopes:

- Requirement [EE.1] related to capabilities of the NFV framework to perform workload (sets of VNF) consolidation, so that unused resources (compute and storage) could be placed into a power conserving state.
- Requirement [EE.2] related to capabilities of the NFV framework to enable authorized entities to control and optimize energy consumption on demand, e.g. by being able to manage power states as needed.
- Requirement [EE.3] related to capabilities of the NFV framework to provide information about timeframes of compute resources, hypervisors and/or VNFs to return to normal operating mode after leaving a specific power-saving mode, so that proper operation and scheduling of infrastructure resources can be determined.

4.3 Other organizations work

4.3.1 3rd Generation Partnership Project (3GPP)

3GPP, and more specifically, the TSG SA WG5 (SA5) on management, orchestration and charging aspects, has studied and specified aspects related to energy efficiency and applicability to 3GPP systems.

Table 4.3.1-1 summarizes the main technical reports and specifications of 3GPP in the domain of energy efficiency and provides some concluding remarks about the relationship with NFV.

Table 4.3.1-1: Summary of 3GPP technical reports and specifications about energy efficiency

Document	Scope (see note)	Relationship with NFV
3GPP TR 21.866 [i.12]	The document identifies and studies key issues and potential solutions in defining energy efficiency Key Performance Indicators (KPI) and energy efficiency optimization operations in existing and future 3GPP networks. Potential solutions for an "energy efficiency control framework" are also introduced.	The document provides explicit references to ETSI NFV's framework for the management of 3GPP systems that include VNFs. No specific technical conclusion is derived apart from indicating that further detailed analysis of EE KPIs for specific VNFs in 3GPP system is suggested.
ETSI TR 132 972 [i.13]	The document studies energy efficiency aspects of 5G networks at both NF level and 3GPP system level, and it covers: KPI definitions, measurement methods, potential solutions and definition of energy efficiency control framework.	The document makes specific reference of energy efficiency KPIs that can be applied at various levels including telecommunication sites, of which data centres are also regarded as an example. References to the virtualization of network elements, thus use of NFV and VNF are explicitly stated. Clause 4.2.3 of ETSI TR 132 972 [i.13] describes EE KPIs for VNFs. The document also highlights the expectation that in the context of VNF's EE KPIs, the NFV-MANO functional blocks are involved in the collection of relevant measurements.
ETSI TS 128 304 [i.14]	The document specifies the requirements for the control and monitoring of Power, Energy and Environmental (PEE) parameters in the following types of Radio Access Network (RAN): GSM, UTRAN, and E-UTRAN.	No specific relationship with NFV is documented. The specification focuses on "radio access network" aspects, and on management of power, energy and environmental parameters of base stations.

Document	Scope (see note)	Relationship with NFV
ETSI TS 128 305 [i.15]	The document specifies the control and monitoring of PEE parameters Integration Reference Point (IRP) Information Service (IS). It specifies the semantics and behaviour of operations, notifications and their parameters visible across Itf-N in a protocol and technology neutral way. The document specifies attributes related to power consumption, temperature, voltage, current and humidity.	The document provides the stage 2 specification of ETSI TS 128 304 [i.14], thus no specific relationship with NFV is documented.
ETSI TS 128 306 [i.16]	The document specifies the solution set definitions for the control and monitoring of PEE parameters IRP.	The document provides the stage 3 specification of ETSI TS 128 305 [i.15].
ETSI TS 128 310 [i.17]	The document specifies concepts, use cases, requirements and solutions for the energy efficiency assessment and optimization for energy saving of 5G networks.	The document highlights that energy efficiency KPIs can also be applicable not only to the network itself, but also to the telecommunication sites, of which data centres are considered to be an element.
3GPP TR 28.813 [i.18]	The document investigates about the definition of new potential energy efficiency KPIs and new energy saving solutions, by identifying a series of key issues and evaluating potential solutions.	Key issue #2a in clause 4.2a of 3GPP TR 28.813 [i.18] describes issues and potential issues for estimating the resource efficiency KPI assuming the scenario of a fully virtualised 5G Core Network (5GCN). For the computation of the resource efficiency, methods for estimating the useful output and resource consumption terms of the resource efficiency formula are also described. Key issue #5 in clause 4.5 of 3GPP TR 28.813 [i.18] proposes potential solutions for the estimation of the 5GC when the NFs are virtualised and run on NFVI. Key issue #8 in clause 4.8 of 3GPP TR 28.813 [i.18] describes potential solutions to derive EE KPI at the network slice level, which can in turn be composed of NFs, fully or partially virtualised.
ETSI TS 128 552 [i.19]	The document specifies performance measurements for 5G networks including network slicing. Within the specified set of performance measurements, measurements related to PEE are also defined.	No specific relationship is established with NFV. However, clause 5.1.1.19 of ETSI TS 128 552 [i.19] specifies PEE measurements valid for 5G PNF, including "PNF power consumption", "PNF energy consumption", "PNF temperature", "PNF voltage", "PNF current" and "PNF humidity", which provides hints about the type of PEE KPIs expected to be gathered from an NF, which might further be of relevance for deriving KPIs in the VNF case.
ETSI TS 128 554 [i.37]	The document specifies end-to-end KPIs for the 5G network and network slicing, including KPIs related to energy efficiency.	The document has relationship with NFV. Energy consumption and related EE KPIs for VNFs are defined. The energy consumption is estimated, rather than actually measured.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

4.3.2 ETSI

4.3.2.1 Environmental Engineering (EE)

ETSI EE has published a number of technical specifications on energy efficiency, including the specification of energy efficiency in NFV environments.

Table 4.3.2.1-1 summarizes the main technical reports and specifications of ETSI EE.

Table 4.3.2.1-1: Summary of ETSI EE technical reports and specifications about energy efficiency

Document	Scope (see note)	Relationship with NFV
ETSI EN 303 471 [i.8]	The document defines an energy efficiency measurement methodology and metrics for Network Functions Virtualisation (NFV).	The document references explicitly ETSI NFV and defines KPIs for Energy Efficiency (EE): KPIEE-bit transfer and KPIEE- packet transfer.
ETSI EN 303 470 [i.20]	The document defines an energy efficiency measurement methodology and metrics for servers.	No reference to ETSI NFV. However, it defines metrics for servers which might be useful for determining energy efficiency according to ETSI NFV's framework.
ETSI ES 203 237 [i.31]	This deliverable specifies an architectural framework and an interface that provide access to the power management capabilities of network nodes to adapt their energy consumption to the load variations.	This deliverable does not reference NFV but serves as a basis for ETSI ES 203 682 [i.33], which focuses on NFV.
ETSI ES 203 539 [i.32]	This deliverable defines metrics and measurement methods for assessing energy efficiency in NFV systems. The measurement methods are intended to be used in lab testing and pre-deployment testing, not in operational NFV environments.	This deliverable focuses on NFV.
ETSI ES 203 682 [i.33]	This deliverable proposes an evolved version of the Green Abstraction Layer formulation capable of operating within NFV environments.	This deliverable extends ETSI ES 203 237 [i.31] by considering NFV requirements.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

4.3.2.2 Access, Terminals, Transmission and Multiplexing (ATTM)

ETSI ATTM has published a number of specifications on energy efficiency that might be of interest when considering energy efficiency for virtualised fixed broadband networks.

Table 4.3.2.2-1 summarizes the main technical reports and specifications of ETSI ATTM.

Table 4.3.2.2-1: Summary of ETSI ATTM technical reports and specifications about energy efficiency

Document	Scope (see note)	Relationship with NFV
ETSI EN 305 200-2-2 [i.21]	This deliverable specifies the requirements for a Global KPI for energy management and their underpinning Objective KPIs addressing the following objectives for the Fixed Access Networks (FANs) of broadband deployment: • energy consumption; • task effectiveness; and • renewable energy.	No reference to ETSI NFV. However, the document defines KPIs for fixed broadband, which might be useful for determining the energy efficiency of virtualised fixed broadband systems.
ETSI EN 305 200-2-3 [i.22]	The document specifies the requirements for a Global KPI for energy management and their underpinning Objective KPIs addressing the following objectives for the mobile access networks of broadband deployment: • energy consumption; • task effectiveness; and • renewable energy	No reference to ETSI NFV. However, the document defines KPIs for mobile broadband, which might be useful for determining the energy efficiency of virtualised mobile broadband systems.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

4.3.3 ITU-T

ITU-T, in particular SG5, has published a number of technical recommendations on energy efficiency, in particular the following listed in table 4.3.3-1.

Table 4.3.3-1: Summary of ITU-T technical recommendations about energy efficiency

Document	Scope (see note)	Relationship with NFV
Recommendation ITU-T L.1316 [i.23]	The document contains a list of main standards on energy efficiency.	No reference to ETSI NFV, but it relates indirectly to NFV by referencing ETSI EN 303 471 [i.8].
Recommendation ITU-T L.1330 [i.24]	The document provides a set of metrics for the assessment of energy efficiency of telecommunication mobile networks, together with proper measurement methods.	No reference to ETSI NFV. However, the document provides measurements methods which might be useful for defining energy efficiency according to ETSI NFV's framework.
Recommendation ITU-T L.1331 [i.25]	The document is an evolution of Recommendation ITU-T L.1330 [i.24] with new requirements for radio sites.	No reference to ETSI NFV. However, the document provides measurements methods on radio sites which might be useful for defining energy efficiency according to ETSI NFV's framework.
Recommendation ITU-T L.1332 [i.26]	The document contains metrics definitions used to evaluate the energy efficiency of an entire network consisting of telecommunication equipment and infrastructure equipment.	No reference to ETSI NFV. However, the document but provides metrics which might be useful for defining energy efficiency according to ETSI NFV's framework.
Recommendation ITU-T L.1350 [i.27]	The document contains energy efficiency metrics used to evaluate the energy efficiency of a base station site.	No reference to ETSI NFV. However, the document provides base station metrics which might be useful for defining energy efficiency according to ETSI NFV's framework.
Recommendation ITU-T L.1361 [i.28]	The document contains metrics definition for NFV environments.	Explicit reference to ETSI NFV. The document provides metrics for NFV environment which might be useful for defining energy efficiency according to ETSI NFV's framework.
Recommendation ITU-T L.1310 [i.66]	The document provides the definition of energy efficiency metrics test procedures, methodologies and measurement profiles for the assessment of energy efficiency of telecommunication equipment.	No reference to ETSI NFV. However, the document provides measurements methods and metrics on various kinds of network equipment that can be used in ETSI NFV environments, such as data centres.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

4.3.4 DMTF Redfish

Table 4.3.4-1 summarizes the main technical reports and specifications of DMTF related to energy efficiency.

Table 4.3.4-1: Summary of DMTF technical reports and specifications about energy efficiency

Document	Scope (see note)	Relationship with NFV
Redfish Telemetry API [i.29]	The document describes the Redfish telemetry service used to retrieve telemetry data across a system. The "PowerConsumedWatts" property defines the power consumption of a chassis. The examples in the paper show how these properties can be retrieve and reported.	No reference to ETSI NFV. However, the "PowerConsumedWatts" property might be useful for computing energy efficiency according to ETSI NFV's framework.
DMTF DSP0266 [i.35]	The document describes an interface standard that uses RESTful interface semantics to access a data model to conduct management operations.	No reference to ETSI NFV. However, the Redfish interface is used to access system telemetry - namely the "PowerConsumedWatts" property, which might be useful for computing energy efficiency according to ETSI NFV's framework.
DMTF DSP2046 [i.34]	The document describes the resources and resource properties that make up the Redfish model.	No reference to ETSI NFV. However, the document lists and describes other power metrics which might be useful for computing energy efficiency according to ETSI NFV's framework.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

4.3.5 Others

Table 4.3.5-1 summarizes the main technical reports and specifications of other organizations related to energy efficiency.

Table 4.3.5-1: Summary of other technical reports and specifications about energy efficiency

Document	Scope (see note)	Relationship with NFV
ATIS-0600015.09.2015 [i.30]	The document provides a methodology for determining the energy efficiency and the measurement and reporting of Base Station metrics.	No specific relationship. The referenced document focuses on base station input power and energy efficiency metrics, and NFV does not specify metrics related to specific network functions.
ATIS-0600015.03.2013 [i.67]	The document is related to router and Ethernet switch products based on their position in a network, as well as a methodology to calculate the energy efficiency ratio.	No specific relationship. However, the document provides measurement methods and metrics of various kinds of network equipment that can be used in ETSI NFV environments, such as data centres.
NOTE: The original scope text from the source document has been slightly adapted to fit within the table and the purpose of analysis from an "energy efficiency" perspective.		

5 Use cases

5.1 Overview

Clause 5 describes various use cases concerning the scope of the present document on energy efficiency for NFV. The use cases cover different aspects such as:

- Network-wide and NFV management and orchestration for energy efficiency;
- Energy-related measurements and runtime information, such as the acquisition of relevant power consumption measurements;
- Energy-driven design, such as VNF design, NFVI configurations; and
- Enablers for smart energy usage and decision making, such as the use of energy efficiency policies.

The list of documented use cases is the following:

- Use case #1: Resources power saving management under different load conditions;
- Use case #2: Power saving management function under limited power supply;
- Use case #3: Acquiring infrastructure power consumption data;
- Use case #4: Acquiring NFV power consumption data;
- Use case #5: Processing power state driven VNF/NS deployment;
- Use case #6: Mapping local VNF/NS and physical resources for optimized power consumption;
- Use case #7: Restoration from power saving situations with resource reservation; and
- Use case #8: Dynamic power-state adjustment of compute resources for maintaining desired power budget.

5.2 Use case #1: Resources power saving management under different load conditions

5.2.1 Introduction

This use case is related to changes of the power state of the resources in the NFVI based on different conditions such as traffic load supported by the VNF and NS instances deployed and managed by NFV-MANO.

The goal of this use case is to showcase the scenario in which, under certain load conditions, resources can be put into a power state that results in less power consumption. This can make the overall resource usage from the NFVI more energy efficient.

NOTE: In the present use case, it is assumed that "suspension", "resume", "switch on", and "switch off" are all examples of resources power management actions that are considered when declaring "changing the power state of a resource".

5.2.2 Actors and roles

Table 5.2.2-1 describes the use case actors and roles.

Table 5.2.2-1: Use case #1 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.

5.2.3 Trigger

Table 5.2.3-1 describes the use case trigger.

Table 5.2.3-1: Use case #1 trigger

Trigger	Description
1	A substantial change in traffic load handled by the VNF instances and NS instances has occurred. The monitoring of traffic load is at the VNF/NS level. The change is assumed to either go below or above traffic load thresholds defined by the service provider, e.g. via operational policies.

5.2.4 Pre-conditions

Table 5.2.4-1 describes the use case pre-conditions.

Table 5.2.4-1: Use case #1 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.

5.2.5 Post-conditions

Table 5.2.5-1 describes the use case post-conditions.

Table 5.2.5-1: Use case #1 post-conditions

#	Post-condition	Description
1	A set of the physical resources in the NFVI are set into the requested power state.	In the case the trigger implied crossing down the defined traffic load threshold, this results in physical resources to be "suspended" or "powered off". In the case the trigger implied crossing up the defined traffic load threshold, this results in physical resources being put into "normal mode".
2	(Optionally) NFVI resources for future VNF/NS restoration are reserved.	In the case the trigger implied crossing down the defined traffic load threshold, NFVI resources can be reserved for future VNF/NS restoration when former VNF/NS deployments are expected to be restored (i.e. back to normal operation). Refer to use case #7 in clause 5.8 about restoration management.

5.2.6 Flow description

Table 5.2.6-1 describes the use case flow.

Table 5.2.6-1: Use case #1 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider, either with an explicit request, or via the execution of some operational policy, indicates to the NFV-MANO system to change the power state of NFVI resources, potentially including information about relevant VNF/NS associated to the traffic load changes.
1	NFV-MANO system	In the case the trigger implied crossing down the defined traffic load threshold, the NFV-MANO performs the consolidation of VNF instance deployments (e.g. migration of VNF/VNFCs into a subset of NFVI resources) to free space for other NFVI resources to be suspended.
2	NFV-MANO system	(Optionally) In the case the trigger implied crossing down the defined traffic load threshold, NFV-MANO system reserves the suspended NFVI resources for future VNF/NS restoration, if such resources are expected to be reused when restoring the former VNF/NS deployments.
3	NFV-MANO system -> NFVI	The NFV-MANO system interacts with the NFVI to change the power state of a set of physical resources.
4	NFVI	The NFVI changes the power state of the requested physical resources.
5	NFVI -> NFV-MANO system	The NFVI confirms the NFV-MANO system about the change of the power state of the involved physical resources.
Ends when	NFV-MANO system -> Service provider	The NFV-MANO informs the service provider about the power state changes of the NFVI, which can include details about the actual/average power consumption after the change.

5.3 Use case #2: Power saving management function under limited power supply

5.3.1 Introduction

This use case showcase a scenario in which the goal is to continuously provide VNF and NS service under limited power supply conditions to certain sites composing the NFVI and NFV-MANO system functions. This use case assumes that the NFVI and NFV-MANO system functions are distributed into several sites at different locations, and power supply is demarcated for each site. In this case, the list of services and quality of service provided by the VNF instances, NS instances and the NFV-MANO system are reduced for non-critical services, in exchange of a lower power consumption mode realization.

User stories related to this use case are:

- A service provider can change the power consumption rate consumed by the VNF instance or NS instance to an optimal rate, when power at sites is derived from natural power generation, e.g. solar panels, and power supply is insufficient due to natural conditions, e.g. weather.
- A service provider can modify the power consumption rate consumed by the VNF instance or NS instance, when sites of NFV-MANO system or NFVI supporting the VNF instance or NS instance are in battery-life mode, caused by planned power outage, or emergent power outage situations.
- A service provider receives power saving requests from electric power companies or governments and is requested to reduce power consumption at one or more locations.

The following clauses describe the use case of "power saving management function under limited power supply" which is related to the above user stories.

5.3.2 Actors and roles

Table 5.3.2-1 describes the use case actors and roles.

Table 5.3.2-1: Use case #2 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.
4	Power Supply Manager	The function or entity that monitors power supply for the NFV system.
5	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes e.g. CPU frequency, power-on/off, etc., and the power consumption on each NFVI-Node and other hardware of the NFVI and the NFV-MANO system. See note.
NOTE: The role and capabilities of physical infrastructure management in relation to NFV-MANO are specified in ETSI GS NFV-IFA 053 [i.65].		

5.3.3 Trigger

Table 5.3.3-1 describes the use case trigger.

Table 5.3.3-1: Use case #2 trigger

Trigger	Description
1	The service provider plans a power supply limitation, or the Power Supply Manager notices a power supply change caused by, e.g. weather conditions, switch to battery power supply or a request for power reduction.

5.3.4 Pre-conditions

Table 5.3.4-1 describes the use case pre-conditions.

Table 5.3.4-1: Use case #2 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.
3	VNF and NS LCM policies aimed at reducing power consumption in power-limited situations are available to the NFV-MANO system.	Either via VNFD/NSD or via runtime operational policies, it is assumed that these can define how VNFs and NSs are expected to consume power in power-limited situations, i.e. power supply lower than defined thresholds. To activate the power saving management function under limited power supply, the policies and thresholds specifying which VNF and NS service(s) to stop/deactivate under certain power supply limited cases are defined in advance.
4	NFV-MANO system function priority is defined in NFV-MANO policies.	Policies define the service level (e.g. which VNF/NS lifecycle management, virtualised resource and containerized workload management tasks to prioritize or rate limiting operation request) provided by NFV-MANO system under power-limited situations, i.e. power supply lower than defined thresholds. To activate the power saving management function under limited power supply, the policies and thresholds specifying which NFV-MANO system service(s) to stop/deactivate under certain power supply limited cases are defined in advance.

5.3.5 Post-conditions

Table 5.3.5-1 describes the use case post-conditions.

Table 5.3.5-1: Use case #2 post-conditions

#	Post-condition	Description
1	VNF instances and NS instances are continuously providing critical services under limited power supply conditions.	In the case the remaining power of station(s) and location(s) comprising NFVI-Nodes is below a certain threshold, VNF instances and NS instances service level is limited until full power supply can be restored, e.g. CPU frequencies consumed by the resources of the VNF instances and NS instances are lowered, or redundant components of the VNF instances and NS instances are terminated.
2	NFV-MANO system is continuously providing critical services under limited power supply conditions.	In the case the remaining power of station(s) and location(s) comprising the NFV-MANO system is below a certain threshold, NFV-MANO system functions are limited to critical services until full power supply can be restored, e.g. VNF lifecycle tasks are shut-down or suspended in the NFV-MANO system functions, or resource management is partially limited to shut down or suspend partial VIM services.
3	(Optionally) NFVI resources for use according to the former state of VNF/NS instances are reserved.	In the case NFVI resources are suspended, the NFVI resources can be reserved for future VNF/NS restoration when former VNF/NS instance deployments are expected to be restored (i.e. back to normal operation). Refer to use case #7 about restoration management.
NOTE: The approaches described as possible post-conditions to limit the VNF and NS service level and NFV-MANO system functions to save power are not meant to be exhaustive. Other approaches or actions might be possible.		

5.3.6 Flow description

Table 5.3.6-1 describes the use case flow.

Table 5.3.6-1: Use case #2 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider / Power Supply Manger	The service provider determines a power supply limitation, or the Power Supply Manager notices a power supply reduction.
1	Service provider / Power Supply Manager -> NFV-MANO system	The service provider or the Power Supply Manager notifies the NFV-MANO system about a power limitation and/or shortage condition. The NFV-MANO system evaluates the condition against the policy and/or thresholds.
2	NFV-MANO system	The NFV-MANO system decides the following actions to save power, depending on the conditions and thresholds crossed: • A) to shut-down or suspend particular NFV-MANO system services, e.g. resource management services or VNF/NS lifecycle management services, • B) to lower the power consumption by the NFVI-Node, e.g. by reducing the CPU frequency, and/or • C) to terminate specific virtualised resources used by the redundant components of VNF/NS instances and shut-down or suspend relevant unused NFVI resources.
3	NFV-MANO system -> VNF / NS instance	To handle case C) (as listed in step #2), the NFV-MANO system terminates or scales in redundant components of VNF / NS instance if needed.
4	NFV-MANO system	(Optionally) To handle case C), in the case relevant NFVI resources are shut-down or suspended, NFV-MANO system reserves the affected NFVI resources for future VNF/NS restoration, if such resources are expected to be reused when restoring the former VNF/NS deployments.
5	NFV-MANO system	To handle case A), the NFV-MANO system shuts down or suspends particular NFV-MANO system services, if needed.
6	NFV-MANO system -> Physical Infrastructure Manager	NFV-MANO system requests the Physical Infrastructure Manager to lower the power consumption of NFVI resources, e.g. by decreasing the CPU frequency or shutting down whole NFVI-Nodes.
7	Power Supply Manager -> NFV-MANO system	The Power Supply Manager notifies the NFVO the new power supply amount rate. The NFVO checks if this rate meets the defined threshold.
8		Repeat steps #2 to #6 until power supply meets the defined threshold.
Ends When	NFV-MANO system -> Service provider	The NFV-MANO system informs the service provider about the change of VNF instance and NS instance operation level through the OSS. This can include details about the actual/average power consumption after the change.

5.4 Use case #3: Acquiring infrastructure power consumption data

5.4.1 Introduction

This use case is about acquiring infrastructure power consumption data from the NFVI.

The goal of this use case is to showcase the scenario by which the service provider can acquire infrastructure power consumption information. With such an information, the service provider can determine whether the energy efficiency policies are having a desired effect, update existing policies, set new policies, and/or create/update deployment policies to address power consumption goals.

Infrastructure power consumption refers to power consumption of physical infrastructure equipment in the NFVI-PoP, including both, the NFVI resources dedicated to directly support the instantiation of virtualised compute, storage and network resources and containerized workloads, and other supporting resources such as cooling systems and racks. For compute resources, it also includes power consumption of the compute servers further detailed by main subcomponents such as CPUs and memory.

5.4.2 Actors and roles

Table 5.4.2-1 describes the use case actors and roles.

Table 5.4.2-1: Use case #3 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads. The NFV-MANO system can also perform aggregation of power consumption data across various NFVI-PoPs and, if needed, perform correlation of power consumption data with managed objects (e.g. NS).
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the resources for the deployment and execution of the VNF instances.
4	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes, e.g. CPU frequency, power-on/off. It is also capable of acquiring power consumption data from the different resources and subsystems in the NFVI. See note.
NOTE: The role and capabilities of physical infrastructure management in relation to NFV-MANO are specified in ETSI GS NFV-IFA 053 [i.65].		

5.4.3 Trigger

Table 5.4.3-1 describes the use case trigger.

Table 5.4.3-1: Use case #3 trigger

Trigger	Description
1	The service provider plans to acquire infrastructure power consumption data from the infrastructure.

5.4.4 Pre-conditions

Table 5.4.4-1 describes the use case pre-conditions.

Table 5.4.4-1: Use case #3 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	The Physical Infrastructure Manager is operational.	No additional description.
3	NFVI is operational.	No additional description.

5.4.5 Post-conditions

Table 5.4.5-1 describes the use case post-conditions.

Table 5.4.5-1: Use case #3 post-conditions

#	Post-condition	Description
1	The infrastructure power consumption data has been acquired by the service provider.	No additional description.

5.4.6 Flow description

Table 5.4.6-1 describes the use case flow.

Table 5.4.6-1: Use case #3 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider, either with an explicit request, or via some operational policy execution, indicates to the NFV-MANO system to acquire infrastructure power consumption data of NFVI resources. The acquisition of data can be for a whole NFVI-PoP or subset of NFVI resources in an NFVI-PoP, or aggregated power consumption across various NFVI-PoP.
1	NFV-MANO system -> Physical Infrastructure Manager	NFV-MANO system requests the Physical Infrastructure Manager to collect infrastructure power consumption data from the defined set of NFVI resources.
2	Physical Infrastructure Manager -> NFV-MANO	The Physical Infrastructure Manager returns the requested infrastructure power consumption.
Ends when	NFV-MANO system -> Service provider	The NFV-MANO provides the service provider the requested infrastructure power consumption data (as requested in the "begins when").

5.5 Use case #4: Acquiring NFV power consumption data

5.5.1 Introduction

This use case is about acquiring power consumption data associated to deployed virtualised resources and containerized workloads. The power consumption data can be further associated to the deployed VNF/NS instances since the NFV-MANO knows about which virtualised resources and containerized workloads are composing the VNF/NS instance. Such power consumption is referred further in the use case as "NFV power consumption".

The goal of this use case is to showcase the scenario by which the service provider can acquire NFV power consumption information which is already processed and associated to specific virtualised resources and containerized workloads, or even associated to VNF/NS instances. With such an information, the service provider can determine whether the energy efficiency policies are having a desired effect, update existing policies, set new policies, and/or create/update deployment policies to address power consumption goals.

The use case considers the possibility that NFV-MANO computes the NFV power consumption data based on:

- Infrastructure power consumption. See also use case described in clause 5.4;
- Pre-processed power consumption associated to specific virtualised resources, such as power consumption associated to a virtualised compute resource (e.g. VM); and
- The sets of concerned virtualised resources, containerized workloads, VNF and NS instances, and their associated key performance indicators.

NOTE: The computation of NFV power consumption data is further investigated in a key issue and solutions.

5.5.2 Actors and roles

Table 5.5.2-1 describes the use case actors and roles.

Table 5.5.2-1: Use case #4 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads. The NFV-MANO system can also perform aggregation of power consumption data across various NFVI-PoPs and correlation of power consumption data with managed objects (e.g. NS).
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.
4	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes, e.g. CPU frequency, power-on/off. It is also capable of acquiring power consumption data from the different resources and subsystems in the NFVI. See note.
NOTE: The role and capabilities of physical infrastructure management in relation to NFV-MANO are specified in ETSI GS NFV-IFA 053 [i.65].		

5.5.3 Trigger

Table 5.5.3-1 describes the use case trigger.

Table 5.5.3-1: Use case #4 trigger

Trigger	Description
1	The service provider plans to acquire NFV power consumption data associated to virtualised resources, containerized workloads, VNFs, and/or NSs.

5.5.4 Pre-conditions

Table 5.5.4-1 describes the use case pre-conditions.

Table 5.5.4-1: Use case #4 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	The Physical Infrastructure Manager is operational.	No additional description.
3	NFVI is operational.	No additional description.

5.5.5 Post-conditions

Table 5.5.5-1 describes the use case post-conditions.

Table 5.5.5-1: Use case #4 post-conditions

#	Post-condition	Description
1	The NFV power consumption data associated to virtualised resources and containerized workloads, and/or to VNF and NS instances, has been acquired by the service provider.	No additional description.

5.5.6 Flow description

Table 5.5.6-1 describes the use case flow.

Table 5.5.6-1: Use case #4 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider, either with an explicit request, or via some operational policy execution, indicates to the NFV-MANO system to acquire NFV power consumption data associated to virtualised resources and/or containerized workloads. The request can further indicate specific VNF and/or NS instances for which associated power consumption data is expected, i.e. the total power consumption of all virtualised resources deployed for a VNF instance.
1	NFV-MANO system	The NFV-MANO system processes the request and determines the sets of NFVI resources associated to the deployed virtualised resources and container workloads of the VNF/NS.
2 (see note 2)	NFV-MANO system -> Physical Infrastructure Manager	The NFV-MANO system requests the Physical Infrastructure Manager to collect infrastructure power consumption data from NFVI resources. See note 1.
3 (see note 2)	Physical Infrastructure Manager -> NFV-MANO system	The Physical Infrastructure Manager returns to the NFV-MANO system the requested infrastructure power consumption data.
4 (see note 2)	NFV-MANO system -> NFVI	The NFV-MANO system requests processed power consumption data from subsystems in the NFVI that already provide such data associated to specific virtualised resources (e.g. VM associated power consumption provided by the hypervisor, as virtualization layer solution in the NFVI).
5 (see note 2)	NFVI -> NFV-MANO system	The NFVI returns to the NFV-MANO system the requested power consumption data.
6	NFV-MANO system	The NFV-MANO system computes the power consumption of the virtualised resources and containerized workloads based on the infrastructure power consumption data and processed data from the infrastructure, the KPIs and configuration of the virtualised resources/containerized workloads. If requested, the NFV-MANO system further compiles the power consumption data corresponding to the VNF/NS instances of concern.
Ends when	NFV-MANO system -> Service provider	The NFV-MANO provides to the service provider the requested power consumption data.
NOTE 1: See also use case described in clause 5.4. NOTE 2: Pairwise steps #2-#3 and #4-#5 can be done in any order (i.e. first perform steps #2-#3, or first perform steps #4-#5) or in parallel.		

5.6 Use case #5: Processing power state driven VNF/NS deployment

5.6.1 Introduction

This use case is about considering power state capabilities and information of compute infrastructure resources when handling the instantiation of VNFs/NSs.

The goal of this use case is to illustrate the scenario by which the service provider can request the instantiation of VNF/NS instances taking into consideration service provider requirements for the actual power state of compute infrastructure resources and requirements of VNF/NSs to be deployed regarding the compute infrastructure resources capabilities and states. For instance, a service provider might aim to deploy certain VNF instances on a pool of resources in the NFVI in which the comprising compute resources are in a specific and known low power state mode.

The use case assumes that compute infrastructure resources can be configured to be set at different power states, and that this information is known and made available to the NFV-MANO for performing the proper virtualised resource/containerized workload scheduling and allocation on such compute resources. It is also assumed that power state information of compute infrastructure resources can be made available by various means, e.g. polling either on-demand or periodically, or received via events/notifications from NFVI subsystems responsible for the compute infrastructure control.

NOTE: The present use case does not assume a control of the power state of compute infrastructure resources, which is considered by other use cases or key issue/solution analysis.

As an example of power state capabilities of compute infrastructure resources, typically processors (CPU) can be in two main state sets (refer to Unified Extensible Firmware Interface (UEFI) Forum's Advanced Configuration and Power Interface (ACPI) Specification [i.36]):

- P-states: these states are execution power saving states. In these states, the processor has the capability of running at different voltage and/or frequency levels. Sub-states such as P0, P1, P2, etc. are typically available, the P0 being the highest state resulting in maximum performance, while P1, P2, and so forth, save power at the penalty of decreasing the processing performance.

- C-states: these states are idle power saving states. In these states, the processor is not executing anything. C0 is still considered an operational state, while C1, C2, etc. intend to transition the CPU and different related subsystems through deeper levels of idling such as turning off parts of the CPU like caches, memory paths, etc.

5.6.2 Actors and roles

Table 5.6.2-1 describes the use case actors and roles.

Table 5.6.2-1: Use case #5 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.

5.6.3 Trigger

Table 5.6.3-1 describes the use case trigger.

Table 5.6.3-1: Use case #5 trigger

Trigger	Description
1	The service provider plans to deploy a set of VNF/NS instances considering certain operational and VNF/NS specified requirements about expected power consumption states for the compute infrastructure resources.

5.6.4 Pre-conditions

Table 5.6.4-1 describes the use case pre-conditions.

Table 5.6.4-1: Use case #5 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.
3	NFV-MANO system has information about the power states of compute infrastructure resources in the NFVI.	It is assumed that NFV-MANO can get such information by some polling mechanism (either on demand or periodically) or via events/notifications issued by NFVI subsystems responsible for the compute infrastructure control. Updating information about power states can also take place during the execution of the use case flow.
4	NFV-MANO system has policies and/or information regarding the deployment of VNF/NS instances considering power state of compute infrastructure resources.	Policies and information could be provided and/or managed in various forms such as: part of VNF/NS descriptors or in energy efficiency-related policies executed by NFV-MANO.

5.6.5 Post-conditions

Table 5.6.5-1 describes the use case post-conditions.

Table 5.6.5-1: Use case #5 post-conditions

#	Post-condition	Description
1	The set of VNF/NS instances has been instantiated according to the power state requirements.	No additional description.
2	The set of compute infrastructure resources on which the set of VNF/NS instances has been instantiated remains according to the required power states.	Operational policies could enforce that no changes in the power states of the compute infrastructure resource take place.

5.6.6 Flow description

Table 5.6.6-1 describes the use case flow.

Table 5.6.6-1: Use case #5 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider, either with an explicit request, or via some operational policy execution, indicates to the NFV-MANO system to perform the instantiation of a set of VNF/NS instances considering the power state of compute infrastructure resources.
1	NFV-MANO system	NFV-MANO system processes the request and checks the power state requirements for the deployments, e.g. against specific requirements contained in the VNF/NS descriptors and/or operational policies in force by the NFV-MANO system.
2	NFV-MANO system	If needed, e.g. to have up-to-date information, the NFV-MANO system retrieves information about the power states of relevant compute infrastructure resources.
3	NFV-MANO system	Based on the requirements (see step #1), power state information, and other requirements (e.g. placement, affinity, storage, etc.) the NFV-MANO system performs the scheduling of the NFVI resources to use for the deployment.
4	NFV-MANO system	The NFV-MANO proceeds with the VNF/NS instantiation.
Ends when	NFV-MANO system -> Service provider	The NFV-MANO informs the service provider about the completion of the VNF/NS instantiation.

5.7 Use case #6: Mapping logical VNF/NS and physical resources for optimized power consumption

5.7.1 Introduction

This use case is about determining VNF/NS resource placement optimizations in conjunction with power consumption savings targets.

The goal of this use case is to illustrate the scenario by which the service provider, with the assistance of NFV-MANO, can determine an optimal VNF/NS placement together with NFVI resource power consumption savings. A possible outcome of the optimization process could be the generation of a blueprint providing:

- the feasible changes in the current VNF/NS resource placements or feasible future VNF/NS deployment and corresponding resource placements; and
- the target NFVI resources whose power consumption can be decreased or saved.

One of the key aspects of the optimization process is to consider both aspects indicated below:

- logical composition of VNF/NSs, whose descriptors declare placement constraints and resource requirements; and
- physical resources and network topology, which determines the capabilities offered by the NFV environment, how such capabilities are distributed and how they can be used for VNF/NS deployments.

Specific examples influencing which and how resources in the NFVI can be put into sleep mode, or in low power consumption states based on the information available about VNF/NS requirements are:

- The priority of certain VNF/NSs and their services might be low, thus potentially they could be scaled down or terminated to free up underlying NFVI resources and put those resources into sleep or shut them down. However, there could be other VNF/NS instances making use of those same NFVI resources, thus modifying the power state of those NFVI resources might be unfeasible.
- When deploying VNFs and NSs, resource and placement requirements determine the expected usage of NFVI resources. Moreover, VNFs/NSs could make use of specific types of resources such as accelerators or have specific placement requirements like logical node (e.g. NUMA) placements. This creates a set of placement requirement dependencies among the potential set of NFVI resources that can be used. Modifying the power consumption or putting certain NFVI resource into a low power consumption state can therefore also affect the VNF/NS deployment.

The use case assumes the following inputs:

- VNF and NS descriptors specify the placement constraints, e.g. in terms of affinity/anti-affinity, and resource requirements, e.g. compute, storage, acceleration and other NFVI capabilities.
- In the case of a new deployment optimization, service provider might deliver to NFV-MANO information about the target performance of the VNF/NS deployment (e.g. expected traffic loads/capacity), in addition to indicating the target VNF/NS deployment flavours and/or instantiation levels.
- NFV-MANO tracks the resource consumption of VNF instances and NS instances, i.e. has current mapping information about what VNF/NS instances use what set of NFVI resources.
- NFV-MANO has information about the NFVI physical resources, e.g. the distribution and types of NFVI resources, and how they are interconnected.
- NFV-MANO has information about the current power consumption of NFVI resources and information about the power profiles of the NFVI resources (e.g. supported power states, power consumption under different power states, etc.).

5.7.2 Actors and roles

Table 5.7.2-1 describes the use case actors and roles.

Table 5.7.2-1: Use case #6 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads, and that is responsible for determining the placement of VNFs/NSs.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.
4	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes, e.g. CPU frequency, power-on/off, etc. See note.
NOTE: The role and capabilities of physical infrastructure management in relation to NFV-MANO are specified in ETSI GS NFV-IFA 053 [i.65].		

5.7.3 Trigger

Table 5.7.3-1 describes the use case trigger.

Table 5.7.3-1: Use case #6 trigger

Trigger	Description
1	The service provider plans to deploy or change the deployment of a set of VNF/NS instances considering certain operational and VNF/NS specified requirements, as well as physical resources topology and capabilities of the NFVI.

5.7.4 Pre-conditions

Table 5.7.4-1 describes the use case pre-conditions.

Table 5.7.4-1: Use case #6 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.
3	NFV-MANO system has information about the power consumption of the NFVI.	It is assumed that the NFV-MANO system can get such information by some polling mechanism (either on demand or periodically) or via events/notifications issued by NFVI subsystems responsible for the compute infrastructure control.
4	NFV-MANO system has NFVI resources topology and capabilities information.	No additional description.
5	NFV-MANO system has policies and/or information regarding the deployment of VNF/NS instances considering power consumption requirements as well as placement and other resources requirements.	Policies and information could be provided and/or managed in various forms such as: part of VNF/NS descriptors or in energy efficiency-related policies executed by NFV-MANO.

5.7.5 Post-conditions

Table 5.7.5-1 describes the use case post-conditions.

Table 5.7.5-1: Use case #6 post-conditions

#	Post-condition	Description
1	NFV-MANO and the service provider have information about an optimal placement of VNFs/NSs for deployment or changes of existing deployments considering a reduced power consumption and fulfilment of necessary VNF/NS placement and resources requirements.	Based on the already deployed VNF/NS instances, there could be the case that no optimal placement or changes in current placement of VNF/NS instances can be achieved to save power consumption. The optimization deployment output could be delivered in some form of "blueprint" or analytics report (similar to an MDA's output as reported in ETSI GR NFV-IFA 041 [i.38]).

5.7.6 Flow description

Table 5.7.6-1 describes the use case flow.

Table 5.7.6-1: Use case #6 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider, either with an explicit request, or via some operational policy execution, indicates to the NFV-MANO system to determine an optimal placement or modification of a set of VNF/NS instances considering the VNF/NS placement and resource constraints and policies with the objective to lower the power consumption of the NFVI resources. In the case of a new VNF/NS deployment, optionally, information about potential traffic loads/capacity of the VNF/NS can be provided, in addition to indicating the target deployment flavours and/or instantiation levels.
1	NFV-MANO system	The NFV-MANO system processes the request and checks the placement and resource requirements for the deployments, e.g. against specific requirements contained in the VNF/NS descriptors and/or operational policies in force by the NFV-MANO system.
2	NFV-MANO system -> Physical Infrastructure Manager	If needed, e.g. to have up-to-date information, the NFV-MANO system retrieves information about the NFVI resources power consumption as well as the NFVI resources topology and capabilities support. See note 1.
3	NFV-MANO system	Based on the requirements (see step #1), NFVI resources power consumption (see step #2), NFVI inventory and resources physical topology (see step #2), the NFV-MANO system computes the scheduling of which NFVI resources to be used for the VNF/NS deployment or current deployment changes.
Ends when	NFV-MANO system -> Service provider	The NFV-MANO informs the service provider about the potential scheduling of NFVI resources and VNF/NS deployment or current deployment changes. The NFV-MANO can also use this information internally to proceed with the VNF/NS instantiation/modification, if necessary. See note 2.
NOTE 1: Refer to use case #3 flow description in clause 5.4.6. NOTE 2: See also remarks in the post-conditions in clause 5.6.5 about the possibility that no power saving could be achieved.		

5.8 Use case #7: Restoration from power saving situations with resource reservation

5.8.1 Introduction

The goal of this use case is to illustrate the scenario by which the service provider can manage the restoration of VNFs from power saving states by using the reserved NFVI resources that had been set into low power consumption states. This use case illustrates how NFV-MANO is involved in managing the allocation of reserved NFVI resources considering the VNF/NS placement constraints and resource requirements, the NFVI resource capabilities and the NFVI resources power consumption states and information.

The use case assumes that VNF instances can be set into a power saving state. This could be realized in various forms, such as (not an exhaustive list):

- By a specific mode of operation of the VNF, e.g. a VNF sleep mode, which could involve processes of storing configuration and state information of the VNF instance and terminating certain power consuming VNFC instances or even, if feasible, all of the VNFC instances; and/or

- By associating power consumption to scaling information of the VNF, thus resolving that power savings can be realized by scaling down the VNF to a minimum set of operational VNFC instances. In this case, it is assumed that handling of configuration and state information is performed as part of the scaling process. As part of the VNF scaling, certain power consuming VNFC instances could be terminated.

Importantly, when restoring a VNF instance from a power saving state, the original configuration and state of the VNF instance is assumed to be also restored, hence the VNF instance is back to a fully functional and operational state. For instance, identifiers of the VNF instance and its VNFC instances have the same values as before setting the VNF instance into the power saving state.

NOTE: The implications of restoring configuration and state information of the VNF are further developed in relevant key issues and solutions.

By operating the power state of the VNF instances and being able to set them into some power saving state (e.g. sleep mode), the NS instances deployed by the service provider can also be regarded to be in some sort of power saving state.

To ensure that VNF instances can be restored from power saving states, NFVI physical and virtualised resources could be reserved before such resources are set into a low power consumption state. Also, when restoring the VNF instance from a power saving state, the NFVI reserved resources are restored from their lower power consumption states.

Optionally, setting priority resource restoration and allocation of reserved NFVI resources could be assumed. In case high priority operations are triggered, the reserved NFVI resources could be recommended to be allocated to the high priority workloads first, e.g. healing/scaling of high-priority VNFs/NSs caused by resource failure or overload. In such cases, the original reserved NFVI resources for the VNFs/NSs are not available, hence other reserved NFVI resources which meet requirements of capabilities and constraints for VNFs/NSs are allocated instead.

5.8.2 Actors and roles

Table 5.8.2-1 describes the use case actors and roles.

Table 5.8.2-1: Use case #7 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.
4	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes, e.g. CPU frequency, power-on/off. It is also capable of acquiring power consumption data from the different resources and subsystems in the NFVI.
5	VNFinstance	The VNF instances running on the NFVI, which are restored by the NFV-MANO system using the reserved NFVI resources.

5.8.3 Trigger

Table 5.8.3-1 describes the use case trigger.

Table 5.8.3-1: Use case #7 trigger

Trigger	Description
1	The service provider plans to restore the set of VNF/NS instances from power saving states (e.g. sleep mode) by using the reserved NFVI resources that had previously been set into some low power consumption state.

5.8.4 Pre-conditions

Table 5.8.4-1 describes the use case pre-conditions.

Table 5.8.4-1: Use case #7 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.
3	VNF instances are in a power saving state.	The VNF instances are in a power saving state (consequently, NS instances consume also less power), and their allocated physical and virtualised resources have been released and also set into a low power consumption state.
4	NFVI resources for the NS/VNF instances are in a low power state, and have been reserved by NFV-MANO.	No additional description.

5.8.5 Post-conditions

Table 5.8.5-1 describes the use case post-conditions.

Table 5.8.5-1: Use case #7 post-conditions

#	Post-condition	Description
1	VNF instances are in a fully functional and operational state, and NFVI resources have been re-allocated to the NS/VNF instances.	No additional description.

5.8.6 Flow description

Table 5.8.6-1 describes the use case flow.

Table 5.8.6-1: Use case #7 flow description

Step #	Actor/Role	Action/Description
Begins when	Service provider -> NFV-MANO system	The service provider requests the NFV-MANO system to restore the set of NS/VNF instances from related power saving states (e.g. sleep mode).
1	NFV-MANO system	Based on the information of the NS/VNF instances, NFV-MANO system identifies the reserved NFVI resources associated to the NS/VNF instances.
2	NFV-MANO system -> Physical Infrastructure Manager	The NFV-MANO system requests the Physical Infrastructure Manager to restore the NFVI resources from a low power consumption state.
3	Physical Infrastructure Manager -> NFV-MANO system	The Physical Infrastructure Manager notifies the NFV-MANO system the success of power state restoration of the NFVI resources.
Ends when	NFV-MANO system	The NFV-MANO system restores the VNF instances from their power saving state (e.g. sleep mode). This involves re-instantiating relevant VNFCs of the VNF instance and restoring saved configuration and state information of the VNF instance. See note.
NOTE: Refer also to the description in clause 5.8.1.		

5.9 Use case #8: Dynamic power-state adjustment of compute resources for maintaining desired power budget

5.9.1 Introduction

Enabling smart energy usage within an NFV system is one of the targets of "Green NFV". This use case describes a scenario where NFV-MANO runs in a closed-loop automated fashion to achieve its overall power budget. This use case enables smart and policy-driven energy usage in the NFVI by leveraging information and managing the CPU power states of the compute resources in a dynamic fashion.

The goal of this use case is to achieve smart energy usage in the NFVI compute resources by adjusting their power states based on runtime CPU utilization. NFV-MANO, in a closed-loop coordination with NFVI, constantly monitors the CPU usage of physical CPUs and dynamically adjusts their execution power states (P-states) based on instantaneous CPU utilization. For example, if CPU utilization of a physical core is low in a given sampling interval, the core is considered underutilized and can therefore be put into a higher order P-state (P1, P2, P3 and so on) in proportion to the load demand for optimal power consumption, without compromising on the performance. Alternatively, CPU P-states transition in reverse order for a physical core that is facing high-utilization. For idle case scenarios, if the CPU remains idle for long periods, the physical cores are put into deeper C-states (idle power states) to conserve power. Relevant subsystems in the NFVI carry out this transitioning of CPU power states based on the policies governed by NFV-MANO.

This use case assumes the following:

- Running VNF/NS instances allow for the adjustment of CPU power states as part of information provided in their respective descriptors. The descriptors could further contain information on CPU utilization thresholds and list of allowed power states (P-states and C-states) for the VNFCs.
- Power management policies are defined in the NFV-MANO for achieving an overall power budget in the form of priority-based set of actions.
- Compute infrastructure resources can be configured to transition between different power states at runtime.
- NFVI subsystems or PIM, responsible for compute infrastructure control, are capable of monitoring CPU utilization per physical core at runtime.
- Power state information of compute infrastructure resources can be made available by various means, e.g. polling either on-demand or periodically, or received via events/notifications from NFVI subsystems.
- NFV-MANO has information about the current power consumption of NFVI resources and information about the power profiles of the NFVI resources (e.g. supported power states, power consumption under different power states, etc.).

NOTE: The current use case does not cover the configuration aspects of sampling interval, CPU utilization thresholds and NFV-MANO power management policies. Furthermore, the mechanisms through which the relevant NFVI subsystems control the power states of physical CPUs is also beyond the scope of the present use case. These aspects are further investigated in the key issue analysis and potential solutions clauses.

5.9.2 Actors and roles

Table 5.9.2-1 describes the use case actors and roles.

Table 5.9.2-1: Use case #8 actors and roles

#	Actor and role	Description
1	Service provider	The entity or organization that operates the NFV system.
2	NFV-MANO system	The NFV-MANO system that provides the management capabilities of NSs, VNFs, virtualised resources and containerized workloads.
3	NFVI	The compendium of physical and virtualised resources, including software elements such as hypervisors. The NFVI also includes CIS in case of container-based VNF deployments. The NFVI provides the necessary resources for the deployment and execution of the VNF instances.
4	Physical Infrastructure Manager	The function that manages and controls NFVI-Nodes e.g. CPU frequency, power-on/off, etc.

5.9.3 Trigger

Table 5.9.3-1 describes the use case trigger.

Table 5.9.3-1: Use case #8 trigger

Trigger	Description
1	The NFV-MANO system notices the overall power consumption of deployed VNF/NS instances exceeds a pre-defined threshold according to its active power management policy and determines to implement smart energy usage policies for running VNF/NS instances.

5.9.4 Pre-conditions

Table 5.9.4-1 describes the use case pre-conditions.

Table 5.9.4-1: Use case #8 pre-conditions

#	Pre-condition	Description
1	NFV-MANO system is operational.	No additional description.
2	NFVI is operational.	No additional description.
3	Physical Infrastructure Manager is operational.	No additional description.
4	NFV-MANO system has information about the power state capabilities of compute infrastructure resources in the NFVI.	It is assumed that the NFV-MANO system can get such information by some polling mechanism (either on demand or periodically) or via events/notifications issued by NFVI subsystems responsible for the compute infrastructure control. Updating information about power states can also take place during the execution of the use case flow.
5	NFV-MANO system has policies and/or information regarding the allowed power states of underlying compute resources for the deployed VNF/NS instances.	Policies and information could be provided and/or managed in various forms such as: part of VNF/NS descriptors or in energy efficiency-related policies executed by NFV-MANO.
6	NFV-MANO system has the infrastructure power consumption data.	No additional description. See note 1.
7	NFV-MANO system has NFV power consumption data associated to virtualised resources and containerized workloads, and/or to VNF and NS instances.	No additional description. See note 2.
NOTE 1: The workflow for acquiring infrastructure power consumption data is described in use case #3 (clause 5.4). NOTE 2: The workflow for acquiring NFV power consumption data is described in use case #4 (clause 5.5).		

5.9.5 Post-conditions

Table 5.9.5-1 describes the use case post-conditions.

Table 5.9.5-1: Use case #8 post-conditions

#	Post-condition	Description
1	NFV-MANO system has notified the service provider about power consumption exceeding desired thresholds even after taking the corrective measures according to its power management policies, requesting further intervention.	If the service provider has subscribed to notifications related to power management, the NFV-MANO system delivers them accordingly.
2	NFV-MANO system has been successful in decreasing the overall power consumption of NFVI compute resources according to its power budget targets.	No additional description.

5.9.6 Flow description

Table 5.9.6-1 describes the use case flow.

Table 5.9.6-1: Use case #8 flow description

Step #	Actor/Role	Action/Description
Begins when	NFV-MANO system	The NFV-MANO system notices the overall power consumption exceeding a pre-defined threshold according to its active power management policy.
1	NFV-MANO system -> service provider	The NFV-MANO system informs the service provider about power consumption exceeding desired thresholds.
2	NFV-MANO system	If needed, the NFV-MANO system retrieves up-to-date information on infrastructure power consumption of relevant subsystems in NFVI from the Physical Infrastructure Manager and computes information on NFV power consumption.
3	NFV-MANO system	Based on up-to-date NFV power consumption information associated to VNF and NS instances, the NFV-MANO system identifies the set of VNF/NS instances and their associated components for dynamic power state adjustments. The NFV-MANO system also considers the descriptors of relevant VNF and NS instances to check for information on allowed power state levels (P-states & C-states) for the underlying compute resources used for the various VNF components and/or NS constituents.

Step #	Actor/Role	Action/Description
4	NFV-MANO system	The NFV-MANO system determines a set of physical CPUs associated with the VNF and NS instances identified in step #2. It further checks that there are no conflicting policies vis-à-vis power states between the virtualised resources and/or containerized workloads (components of VNF instances and/or constituents of NS instances) sharing the same physical compute resources.
5	NFV-MANO system -> NFVI	NFV-MANO system instructs the NFVI subsystem responsible for controlling compute resources to adjust power states of a set of physical CPUs according to their runtime CPU utilization. NFV-MANO system also conveys to the NFVI the list of CPU power states that are allowed in the VNF/NS descriptors under "normal" conditions. See note 1.
6	NFVI	The NFVI subsystem for managing physical compute resources constantly monitors CPU usage per each core and assigns them appropriate P-states and/or C-states according to their runtime utilization. See note 2.
7	NFVI -> NFV-MANO system	The NFVI subsystem sends the updated NFVI power consumption information to the NFV-MANO system after the dynamic power states adjustments have been made by the corresponding NFVI subsystem.
8	NFV-MANO system -> service provider	The NFV-MANO system evaluates the effects of dynamic adjustments on overall power consumption. It keeps monitoring the overall power consumption periodically through a polling and/or notification mechanism. If necessary, the NFV-MANO system opts for another power management policy such as suspending or shutting down redundant instances in order to meet the power budget. The NFV-MANO system informs the service provider about the power consumption exceeding desired thresholds, even after taking corrective measures. See note 3.
Ends when	NFV-MANO system	The NFV-MANO system notes the power consumption falling below the specified threshold according to its active power management policy and the associated power budget.
NOTE 1: "Normal" conditions correspond to the NFV system working under normal traffic load and power supply conditions as opposed to the triggers described in use case #1 (clause 5.2) and use case #2 (clause 5.3). NOTE 2: A hypervisor can be the NFVI subsystem that monitors CPU usage at runtime and adjusts P-states and C-states of physical CPU cores. NOTE 3: The NFV-MANO system will stay in this stage by periodically monitoring the power consumption information and might trigger other operational flows (described in use case #2, clause 5.3) according to its power management policies until it achieves the desired power consumption.		

6 Key issue analysis

6.1 Introduction

From the use cases documented in clause 5 and other sources of information such as reports and specifications referenced in clause 2, in the following clauses, key issues regarding energy efficiency and power consumption are identified and documented. The key issues are classified based on the following aspects:

- Components subject to energy consumption in a service provider network (refer to clause 6.2);
- Energy efficiency considerations in management and orchestration of a service provider network (refer to clause 6.3);
- Design for energy efficiency (refer to clause 6.4); and
- Metrics and KPIs related to energy consumption and efficiency (refer to clause 6.5).

6.2 Key issues on energy consumption

The following key issues are identified related to the components that are subject to energy consumption in a service provider network.

Key issue #1.1: Power consuming components in NFV-based network deployments

Use case #3 in clause 5.4 illustrates the scenario in which the service provider can acquire infrastructure power consumption data.

This key issue is about the determination of what components in an NFV-based network deployment environment consume power. Questions that can be formulated are (not fully exhaustive):

- What components in an NFV-based network deployment are subject to power consumption?
- How can the power consumption data of the relevant components be acquired by the service provider and/or NFV-MANO?
- How can the NFV-MANO system be made aware of the power state capabilities of compute infrastructure resources in the NFVI?

Key issue #1.2: Changing power state in NFV-based network deployments

Use case #1 in clause 5.2 illustrates the scenario of varying the power state of NFVI physical resources based on different load conditions.

This key issue is about determining which elements in an NFV-based network deployment environment support their power state to be changed and the mechanisms that enable it. Thus, questions that can be formulated are (not fully exhaustive):

- What components in an NFV-based network deployment are subject to, and capable to be changed with respect to their power consumption?
- How can the power consumption of a component be changed and through which systems?

Key issue #1.3: Power distribution in NFV environments

Use case #3 in clause 5.4 illustrates the scenario by which the service provider can acquire infrastructure power consumption data.

As described in the key issue #1.1, there are several components in a NFV environment that are directly subject to consume power. Specific components in the NFVI are responsible for distributing the power to the power outlets. Based on the classification specified in the DMTF DSP2046 [i.34], the list of relevant power and power distribution equipment includes:

- Related to power distribution:
  - Electrical buses;
  - Floor Power Distribution Units (PDU): a PDU providing feeder circuits for further power distribution;
  - Power shelves;
  - Rack PDUs: a PDU providing outlets for a rack or similar devices; and
  - Transfer switches, including automatic/manual power transfer switches and electrical switchgears.
- Related to power supply:
  - Power supply units.

The key issue #1.3 concerns whether the components related to power distribution and supply are expected to be included in the calculation of the NFVI power consumption , and if so, the means to meter such power consumption and make it available to the NFV-MANO for further consideration and quantification. Hence, questions that can be formulated related to this key issue are (not fully exhaustive):

- Are the components related to power distribution also subject to power consumption, and if so, how it can be measured?
- Are the components related to power distribution to be considered in the calculation of the power consumption of the NFVI?

6.3 Key issues on energy efficiency management

The following key issues are identified related to aspects of energy efficiency management.

Key issue #2.1: Management and orchestration processes considering energy efficiency

Use case #1 in clause 5.2 illustrates the scenario in which different load conditions can be a trigger to orchestrate and manage the power state of the NFVI with the target to implement energy efficiency policies. In addition, use case #2 in clause 5.3 illustrates scenarios in which the service provider can modify the power state consumed by VNFs and NSs based on different conditions such as under limited power supply. Finally, use case #6 in clause 5.7 describes the scenario in which the NFV-MANO can determine an optimal VNF/NS placement considering NFVI resource power consumption savings.

The key issue #2.1 is about listing and describing the NFV management and orchestration capabilities that can be used to apply energy efficiency policies. Questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the NFV management and orchestration capabilities that can be leveraged for applying energy efficiency policies?
- How can the service provider determine if the implemented energy efficiency policies are achieving the expected goals?
- How can the energy efficiency policies be provided to and executed by the NFV-MANO?
- What sources of information can be considered to determine the optimal placement of VNFs/NSs delivering power consumption savings according to the energy efficiency policies?
- Which placement constraints need to be considered while deploying VNF/VNFC instances on compute resources with physical CPUs whose power state can be dynamically or statically managed?

Key issue #2.2: Resiliency, availability and restoration during power management processes

Use case #7 in clause 5.8 describes the case in which VNF instances can be restored from power saving states with the assumption that VNF instances can be fully restored from a configuration, state and resource perspective to lower the impact on the service availability. Use case #2 in clause 5.3 and use case #1 in clause 5.2 illustrate also steps in which NFVI resources to be powered down could be reserved to facilitate the restoration of the VNF instance from a power saving state.

The key issue #2.2 is about describing the management processes that can maintain the resiliency and availability levels of service during and after performing power management actions. Questions that can be formulated related to this key issue are (not fully exhaustive):

- How can a VNF instance be restored from power saving states without losing or corrupting configuration and state information?
- How can a VNF instance resources' availability be guaranteed across power management state changes?
- What are the impacted workloads and deployments supported by the NFVI components when the power consumption of these is changed?

Key issue #2.3: Efficient power state management of NFVI compute resources for heterogenous workloads

On the one hand, use case #5 in clause 5.6 assumes that the power states of compute infrastructure resources do not change once the set of VNF/NS instances has been deployed on the NFVI according to their power state requirements. NFV-MANO system can enforce this using a set of operational policies regarding energy efficiency. On the other hand, use case #8 in clause 5.9 describes a scenario where the NFV-MANO system performs dynamic power state management of physical compute resources in the NFVI to maintain a desired power budget under normal operations.

Conflicting requirements and policies might arise for power state management among the virtualised workloads in a heterogeneous environment where multiple workloads share the same underlying physical resources but have different power management policies. In that regard, the present key issue concerns the configuration aspects related to power state management of NFVI compute resources that consider the conflicting power management policies and power consumption requirements of virtualised workloads sharing the same physical resources. Relevant questions that can be formulated related to this key issue are (not fully exhaustive):

- What aspects need to be considered by the NFV-MANO for dynamically adjusting power states of compute resources in the NFVI to avoid power management policy conflicts arising from different virtualised workloads sharing the same physical resources?
- What are the optimal power management policies for heterogeneous virtualised workloads sharing the same physical resources?

6.4 Key issues on energy efficient design

The following key issues are identified related to aspects of designing NFV-based network deployments for energy efficiency.

Key issue #3.1: VNF/NS driven energy efficiency design

Use case #2 in clause 5.3 illustrates the case in which the power consumption of VNF and NS instances can be changed to an optimal rate depending on certain conditions, such as under limited power supply. This implies that VNFs and NSs can be designed in a way that take into consideration the optimization of their power consumption, and for instance, relevant descriptors are assumed to define the power consumption levels and/or requirements of VNFs an NSs. In addition, use case #5 in clause 5.6 illustrates the scenario in which compute resources can be set into certain power state. Such power state information can be made available to NFV-MANO, and together with information provided by VNF/NS descriptors related to power state requirements of compute resources, an optimized energy efficient deployment can be performed. Furthermore, use case #6 in clause 5.7 describes the case in which the VNF/NS placement can be optimized from a power consumption perspective based on the knowledge of specific placement requirements and dependencies. And finally, use case #7 in clause 5.8 also assumes that VNF instances can be set into a power saving state.

The key issue #3.1 concerns the design aspects of VNFs and NSs that influence the power consumption of the deployed respective instances and how these can be expressed towards the NFV-MANO. Thus, related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the design aspects of VNF and NS that influence their power consumption (e.g. redundancy configurations of VNFs/NSs tailored to low power consumption)?
- How can VNFs and NSs expose/describe energy efficiency and/or power consumption requirements?
- How can power saving states be defined for a VNF and be managed?

Key issue #3.2: NFV-MANO driven energy efficiency

Use case #2 in clause 5.3 exemplifies the scenario in which NFV-MANO system functions can be shut down or partially suspended in response to known limited power supply towards the NFVI hosting VNF/NS instances managed by NFV-MANO.

The key issue #3.2 concerns the design aspects of the NFV-MANO and to the capabilities offered by NFV-MANO that can influence the power consumption, as well as to the aspects that can be considered for implementing energy efficiency policies for NFV-MANO functions. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the design aspects of NFV-MANO systems that influence their power consumption (e.g. distribution of functional entities implementing NFV-MANO functional blocks and functions across the NFV-based network deployment)?
- Which NFV-MANO capabilities can be tailored to implement power consumption and/or energy efficiency policies?
- How can NFV-MANO capabilities be configured during the runtime of the NFV-MANO system to accommodate certain power consumption targets?

Key issue #3.3: NFVI driven energy efficiency

Use case #5 in clause 5.6 describes the scenario in which compute resources can be set into certain power state, which in turn can be a tool for driving energy efficiency design of the NFVI.

The key issue #3.3 concerns the design aspects of the NFVI and the capabilities offered by the NFVI resources that can influence the power consumption and can be considered for implementing energy efficiency policies. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the design aspects of the NFVI that influence its power consumption (e.g. choice of compute resources to assemble the NFVI)?
- Which NFVI capabilities and resources can be tailored to implement power consumption and/or energy efficiency policies?
- What are the design considerations for pooling physical and virtual compute resources in the NFVI to enable power-state driven VNF/NS deployments?

6.5 Key issues on energy related metrics and KPIs

Key issue #4.1: Energy and power consumption related metrics and KPIs of NFVI resources

Use case #1 in clause 5.2 includes steps in which the NFVI and NFV-MANO can provide details about the actual/average power consumption after applying the request to change the power state of a set of physical resources. A similar scenario is showcased in the use case #2 in clause 5.3. Finally, use case #3 in clause 5.4 focuses on illustrating the capability through which the service provider and NFV-MANO can acquire power consumption data from the infrastructure resources.

The key issue #4.1 is about determining what physical infrastructure resources in the NFVI can be monitored in terms of power consumption and energy, and how such monitoring data can be performed and made available to NFV-MANO and the service provider. The key issue is also about defining related energy efficiency and energy consumption KPIs for physical resources in the NFVI. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the energy efficiency and power consumption KPIs of physical resources in the NFVI?
- How can energy efficiency be defined for physical resources in the NFVI?
- How can NFV-MANO and the service provider acquire power consumption and energy efficiency related measurements of physical resources in the NFVI? What interfaces can be involved in the acquisition of power consumption?

Key issue #4.2: Energy and power consumption related metrics and KPIs of virtualised resources and containerized workloads

Use case #3 in clause 5.4 focuses on illustrating the capability through which the service provider and NFV-MANO can acquire and compute power consumption data associated to NFV-MANO managed objects such as virtualised resources and containerized workloads.

The key issue #4.2 is about determining what virtualised resources and containerized workloads can be monitored in terms of power consumption and energy, and how such monitoring data can be performed and made available to NFV-MANO and the service provider. The key issue is also about defining related energy efficiency and energy consumption KPIs for virtualised resources and containerized workloads. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the energy efficiency and power consumption KPIs of virtualised resources and containerized workloads?
- How can energy efficiency be defined for virtualised resources and containerized workloads?
- How can NFV-MANO and the service provider acquire power consumption and energy efficiency related measurements of virtualised resources and containerized workloads? What interfaces can be involved in the acquisition of power consumption?

Key issue #4.3: Energy and power consumption related metrics and KPIs of VNFs

Use case #3 in clause 5.4 focuses on illustrating the capability through which the service provider and NFV-MANO can acquire and compute power consumption data associated to VNF instances managed by NFV-MANO.

According to ETSI TS 128 554 [i.37], the energy consumption of a VNF is not measured, instead it is estimated. As being an estimation, its value might not be accurate enough. Therefore, other forms of measuring the energy consumption of VNFs are still an open issue.

The key issue #4.3 is about determining how to monitor and/or compute energy and power consumption associated to VNF instances, and to make such data available to NFV-MANO and the service provider. The key issue is also about defining related energy efficiency and energy consumption KPIs of VNFs. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the energy efficiency and power consumption KPIs of VNFs?
- How can energy efficiency be defined for a VNF?
- How can NFV-MANO and the service provider acquire power consumption and energy efficiency related measurements of VNFs? What interfaces can be involved in the acquisition of power consumption?

Key issue #4.4: Energy and power consumption related metrics and KPIs of NSs

Use case #3 in clause 5.4 focuses on illustrating the capability through which the service provider and NFV-MANO can acquire and compute power consumption associated to NS instances managed by NFV-MANO.

The key issue #4.4 is about determining how to monitor and/or compute energy and power consumption associated to NS instances, and to make such data available to NFV-MANO and the service provider. The key issue is also about defining related energy efficiency and energy consumption KPIs of NSs. Related questions that can be formulated related to this key issue are (not fully exhaustive):

- What are the energy efficiency and power consumption KPIs of NSs?
- How can energy efficiency be defined for an NS?
- How can NFV-MANO and the service provider acquire power consumption and energy efficiency related measurements of NSs? What interfaces can be involved in the acquisition of power consumption?

7 Framework and potential solutions

7.1 Introduction

Clause 7 introduces and describes potential solutions related to energy efficiency and power management, which aim to address one or more of the key issues documented in clause 6 of the present document.

Solutions are documented in the following manner:

- an introduction providing background and conceptual information on which the solution is based;
- a description of the technical solution;
- a reference to the key issues that the solution is addressing; and
- a description of technical gaps in the ETSI NFV architectural framework and/or referenced ETSI NFV specifications, in case there is any.

Table 7.1-1 summarizes the list of solutions documented and the key issues to which they relate to.

NOTE: The present document version does not describe any solution related to key issue #1.3 on "power distribution in NFV environments". The key issue concerns mainly the physical construction of the NFVI-PoP/sites and how the power is distributed among the resources comprising the NFVI-PoP. For information about relevant power distribution elements, refer to the description in the key issue #1.3 in clause 6.2.

Table 7.1-1: Summary of solutions and related key issues

Solution	Title	Related key issues
Solution #1	Listing and description of power consuming elements in NFV deployments	#1.1: Power consuming components in NFV-based network deployments #1.2: Changing power state in NFV-based network deployments #3.3: NFVI driven energy efficiency #4.1: Energy and power consumption related metrics and KPIs of NFVI resources
Solution #2	Changing the power state of NFVI components	#1.2 #3.3
Solution #3	Resources reservation for VNFs under power management	#2.2: Resiliency, availability and restoration during power management processes
Solution #4	State and configuration management during VNF power management	#2.2
Solution #5	Power profiles of a VNF	#3.1: VNF/NS driven energy efficiency design
Solution #6	VNF power manager/controller	#3.1 #4.3: Energy and power consumption related metrics and KPIs of VNFs
Solution #7	Power state based zones in NFVI for static and dynamic power state management of compute resources	#3.3 #2.1: Management and orchestration processes considering energy efficiency #2.3: Efficient power state management of NFVI compute resources for heterogenous workloads
Solution #8	Power-state aware allocation of compute resources for VNF/NS orchestration and lifecycle management	#2.1 #3.1
Solution #9	Energy efficiency policies	#2.1
Solution #10	NFV-MANO state management for energy efficiency	#3.2: NFV-MANO driven energy efficiency
Solution #11	NFV-MANO capabilities for energy efficiency	#2.1
Solution #12	VIM for acquiring virtualised resource power consumption	#4.2: Energy and power consumption related metrics and KPIs of virtualised resources and containerized workloads
Solution #13	CISM for acquiring containerized workload power consumption	#4.2
Solution #14	NFVO for acquiring NS power consumption	#4.4: Energy and power consumption related metrics and KPIs of NSs
Solution #15	NFV power metrics function for acquiring multiple kinds of power consumption data	#4.2 #4.3 #4.4
Solution #16	Energy efficiency of NFVI physical resources	#4.1
Solution #17	Energy efficiency of virtualised resources and containerized workloads	#4.2
Solution #18	Energy metrics related to virtualised resources and containerized workloads	#4.2
Solution #19	Energy efficiency of VNFs	#4.3
Solution #20	Energy metrics related to VNFs	#4.3

Solution	Title	Related key issues
Solution #21	Energy efficiency of NSs	#4.4
Solution #22	Energy metrics related to NSs	#4.4
Solution #23	PIM for acquiring power consumption of physical resources in the NFVI	#4.1

7.2 Potential solutions

7.2.1 Solution #1: Listing and description of power consuming elements in NFV deployments

7.2.1.1 Introduction

Power consuming components concern physical components that consume power in order to operate and perform their functionality. Figure 7.2.1.1-1 illustrates an overview of the main power consuming components in an NFVI-PoP. In figure 7.2.1.1-1, these are surrounded by the "dotted box" and are listed as follows:

- Compute hardware,
- Storage hardware,
- Network hardware,
- NFVI Node racks infrastructure, and
- NFVI facilities infrastructure.

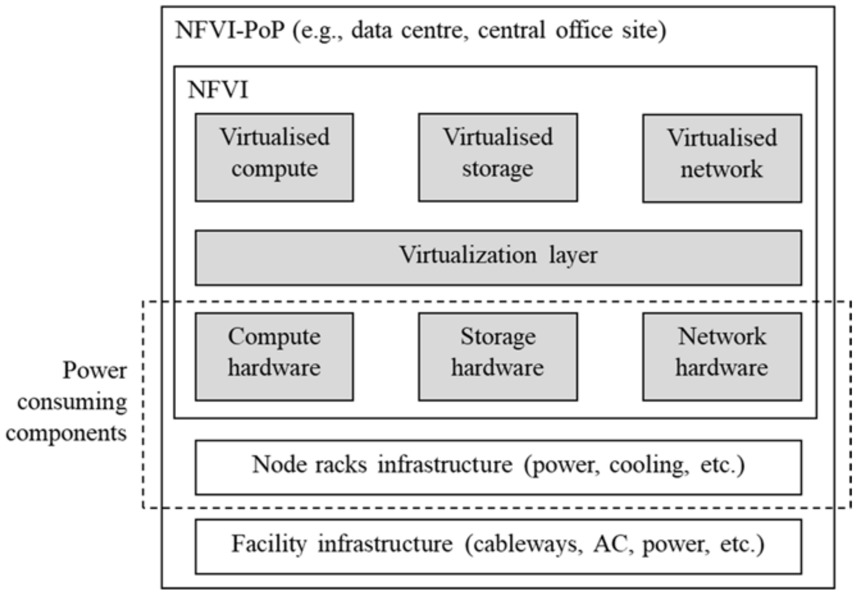

Figure 7.2.1.1-1: Power consuming components in an NFVI-PoP

Physical components can in turn be composed of more granular elements which also add into the power consumption. For instance, Redfish DMTF DSP2046 [i.34] defines "EnvironmentMetrics" applicable to many types of components such as:

- Chassis: sheet-metal confined spaces and zones such as racks, enclosures, chassis and other containers. These can in turn contain compute servers, network devices, storage resources and other subsystems such as cooling fan units;
- Compute server: memory, processors, PCIe devices, network adapters and network ports;
- Storage resource: storage drives, storage controller, storage controller ports, hard drives in compute resources;

- Network device: switch, router, network device ports; and
- Others: media controllers.

According to DMTF DSP2046 [i.34], information about power consumption can be made available with the "PowerWatts" attribute associated to the various types of components listed above. The attribute indicates the power consumption of the component, in watts. The presence of this attribute indicates to a consumer the metrics related to the component's power consumption.

In addition, for chassis components, additional power consumption metrics can be obtained using the DMTF DSP2046 [i.34], such as:

- PowerCapacityWatts: indicates the amount of power that can be allocated to the chassis, in watts.
- PowerConsumedWatts: indicates the actual power that the chassis consumes, in watts,
- PowerMetrics: provides various power readings about the chassis, such as "AverageConsumedWatts", "MaxConsumedWatts" and "MinConsumedWatts" derived over a certain period of time.

7.2.1.2 Solution description

The present solution considers the capability of NFV-MANO to hold an inventory of physical resources present in the NFVI. Furthermore, for each one of the inventoried resources, the information in the inventory indicates whether the component is subject to power consumption, and if so, the levels of power consumption under different conditions, in the case that such information is available. The inventory of physical resources can be realized by a specialized functional block or function responsible for the management of physical resources, e.g. a Physical Infrastructure Management (PIM) function, as being specified in ETSI GS NFV-IFA 053 [i.65].

NOTE: The present solution does not cover how, potentially, other NFV-MANO entities can acquire power consumption and management related information.

Figure 7.2.1.2-1 illustrates an example of the solution.

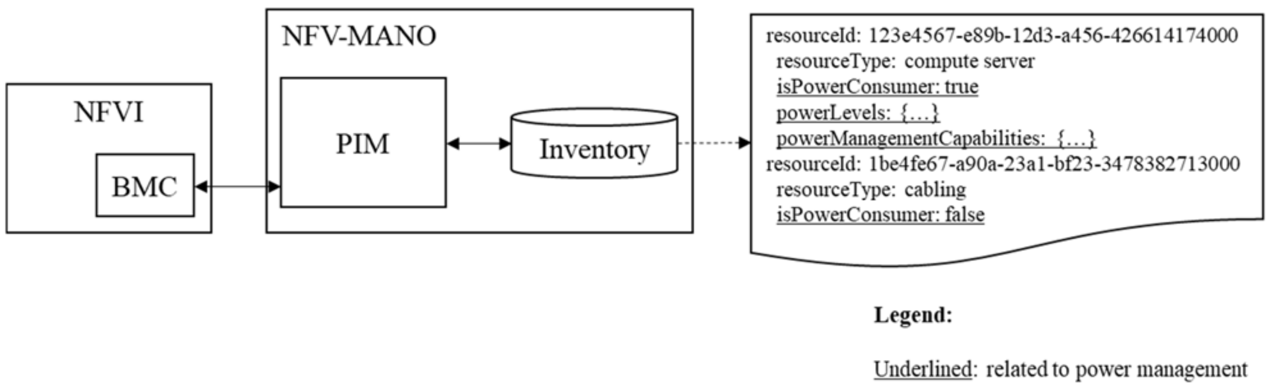

Figure 7.2.1.2-1: Illustrative example of Solution #1

The resources inventory provides information for each of the inventoried resources:

- whether it consumes power;
- the power level(s), including information about supported power states, current power state and current power consumption (or relevant metrics); and
- power management capabilities, e.g. whether the power states of the resource can be managed and how these can be managed such as statically or dynamically via some policies or configuration.

The solution envisions that the PIM interacts with relevant Baseboard Management Controllers (BMC) that implement an interface enabling remote management capabilities of the various systems and components in the NFVI. An example of such type of service and interface is DMTF's Redfish DSP0266 [i.35].

7.2.1.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #1.1,
- Key issue #1.2,
- Key issue #3.3; and
- Key issue #4.1.

7.2.1.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #1.1: The NFV-MANO to support an inventory of physical resources in the NFVI.

Gap #1.2: The NFV-MANO to support discovery, storage and use of information about power consumption characteristics of physical resources in the NFVI.

7.2.2 Solution #2: Changing the power state of NFVI components

7.2.2.1 Introduction

As described in Solution #1, a majority of components (resources) in the NFVI consume power. Likewise, some of the components offer power management capabilities.

Three main groupings of NFVI components are considered for performing power management:

- Physical resources and associated components;
- Processor level subsystems; and
- Virtualization layer components.

Regarding physical resources and associated components, DMTF DSP2046 Redfish [i.34] defines diverse types related to device-level power management. Power management actions (also defined as "reset types") can be performed on the following components: Chassis, ComputerSystem, Processor, Drive (for storage), MediaController, Memory, Switch, Port, and Power Supply.

Standard values of reset types (referred as "ResetType" in DMTF DSP2046 Redfish [i.34]) are listed in table 7.2.2.1-1.

Table 7.2.2.1-1: Types of reset actions according to DMTF DSP2046 Redfish [i.34]

Value	Description
ForceOff	Turn off the unit immediately (non-graceful shutdown).
ForceOn	Turn on the unit immediately.
ForceRestart	Shut down immediately and non-gracefully, and restart the system.
GracefulRestart	Shut down gracefully and restart the system.
GracefulShutdown	Shut down gracefully and power off.
Nmi	Generate a diagnostic interrupt, Non-Maskable Interrupt (NMI) on x86 systems, to stop normal operations, complete diagnostic actions, and typically, halt the system.
On	Turn on the unit.
Pause	Pause execution on the unit, but do not remove power. This is typically a feature of virtual machine hypervisors.
PowerCycle	Power cycle the unit. Behaves like a full power removal, followed by a power restore to the resource.
PushPowerButton	Simulate the pressing of the physical power button on the unit.
Resume	Resume execution on the paused unit. This is typically a feature of virtual machine hypervisors.
Suspend	Write the state of the unit to disk before powering off. This allows for the state to be restored when powered back on.

Not all components necessarily support all possible reset actions. According to DMTF DSP2046 Redfish [i.34], this is indicated for each of the components with an "AllowableValues" property.

From the list in table 7.2.2.1-1, it is of interest to remark the following actions that can modify the power consumption incurred by the component:

- To turn off the component: power consumption will be decreased by turning off/shutting down the unit.
- To suspend the component: power consumption will be decreased by powering off the component but allowing to also write the state for later restoration.
- To pause the component: the unit pauses execution without removing power, which can decrease power consumption by not incurring in actual usage of the resource.

Power management performed through the virtualization layer components enables to perform power actions on virtualised resources. An example solution in the case of VM-based virtualization is the libvirt [i.39], an API, daemon and management tools for managing platform virtualization. The API is used for the interaction between the virtualization layer software (i.e. hypervisor in the case of VM-based virtualization) and the virtualization manager and orchestrator. More specifically, the libvirt-domain API [i.40] offers capabilities to perform management actions on "domains" (i.e. a VM), including:

- Suspending an active domain, wherein the domain is frozen without further access to CPU resources and I/O, but memory is still used by the domain at the hypervisor level.
- Resuming a suspended domain.
- Shutting down a domain, in which the domain's operating system (OS) is stopped, but it can be then recovered.
- Creating a domain, in which the domain's OS is launched.

In the case of OS container-based virtualization, commonly used solutions for the orchestration and management of containerized workloads do not consider the capabilities to suspend, pause, or shutdown sets of one or more OS container (e.g. a Pod in the case of Kubernetes®); the sets of one or more OS container are either simply created or terminated.

NOTE 1: The Kubernetes® Word Mark and Kubernetes Logo are either registered trademarks/service marks or trademarks/service marks of The Linux Foundation, in the United States and other countries and are used with the Linux Foundation's permission.

Regarding processor power management, UEFI Forum's ACPI Specification [i.36] specifies the main states in which a processor (or CPU) can be:

- C-states: these are power states that a processor can use to reduce power consumption. This can be done on a per-core level, or on a CPU package level (by powering down portions of the core package) or both. C0 is an active/operational state where instructions are still executed. The C1 through Cn power states are processing sleeping states in which the processor consumes less power compared to C0.
- P-states: these are execution performance states in which the processor has the capability of running at different voltage and/or frequency levels, thus incurring in various levels of power consumption. These states are available while the processor is in the C0 state. Sub-states such as P0, P1, P2, etc. are typically available, P0 being the highest state resulting in maximum performance, while P1, P2, and so forth, save power at the penalty of decreasing the processing performance.

Core C-states can be controlled by the OS as defined by the UEFI Forum's ACPI Specification [i.36]. P-states can typically be managed from user space applications running on the OS. For instance, in Linux®, the management of P-states can be done via the kernel system file system "sysfs" using a management routing, known as "governor", that decides how to control the frequency of the processor in response to the workload.

NOTE 2: Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

In VM-based virtualization technology, hypervisors support setting custom power management policies for each physical CPU core. Additionally, there is support for dynamic transitioning between P-states based on runtime CPU utilization for each physical core. However, many typical physical chipsets do not support setting a custom P-state for individual cores and instead only support the same state for all the physical cores of the CPU.

NOTE 3: The relevant NFVI subsystem, a hypervisor for example, controls CPU power states by adjusting the voltage and/or frequency levels (V/F levels) for physical CPUs, optionally on a per-core basis.

In OS-based virtualization technology, it is possible to manage the power of physical CPU cores using CIS cluster enhancement capabilities (CCEC). An example is the Kubernetes® Power Manager [i.41].

There can be two ways by which the relevant NFVI subsystem can manage power states of the underlying CPU resources, either setting a static or custom power state per physical core or dynamically managing the power states based on per-core CPU utilization.

Figure 7.2.2.1-1 illustrates an exemplary relationship among processor power states, as defined by UEFI Forum's ACPI Specification [i.36]. The example shows that P-states can be transitioned while the processor is in the state C0 through a "throttling" process.

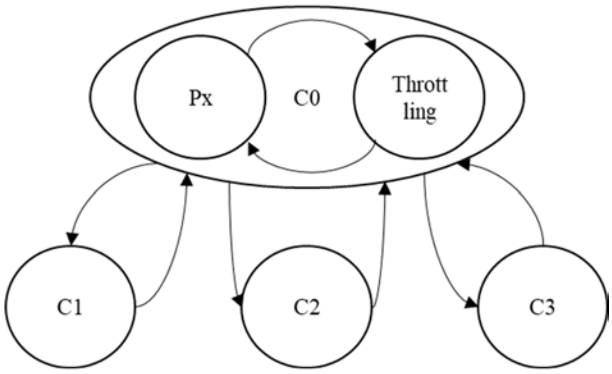

Figure 7.2.2.1-1: Example of processor power state transitions

7.2.2.2 Solution description

The present solution considers a framework in which NFV-MANO components (functions and functional blocks) interact with specific entities in the NFVI and the VNF/VNF generic OAM functions to perform power management of the three NFVI components considered in clause 7.2.2.1.

Figure 7.2.2.2-1 illustrates an overview of the present solution. References to the tagged interactions overlayed on the framework's diagram are provided hereafter in the description of the solution.

NOTE: In figure 7.2.2.2-1, the greyed boxes highlight the components whose power state is subject to be changed.

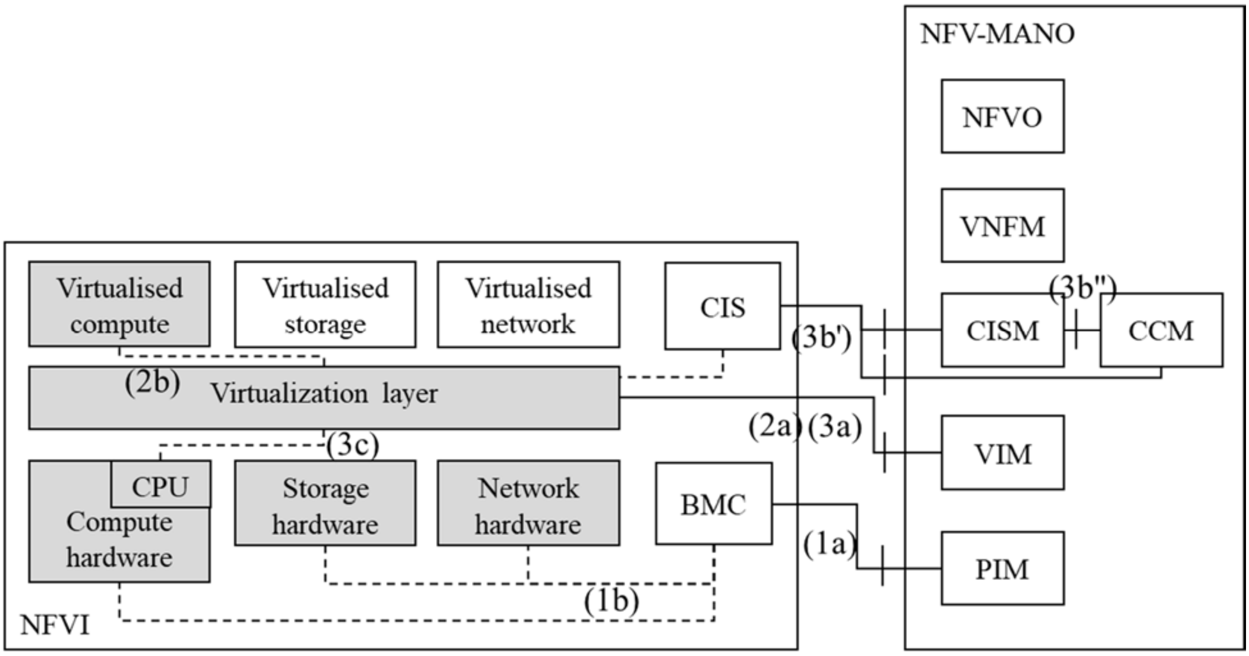

Figure 7.2.2.2-1: Overview of Solution #2

For managing the power state of the physical resource and associated components, a specialized functional block or function in NFV-MANO responsible for the management of physical resources, e.g. a PIM function, interacts with a BMC, which is represented in figure 7.2.2.2-1 within the NFVI, as shown on the interaction tagged as (1a). The BMC enables remote management capabilities of the various system and components in the NFVI, including the management of power states when these are supported by the resources. Such interaction can be realized by using standard interfaces, such as those specified in the DMTF's Redfish DSP0266 [i.35]. Based on the requested actions by the PIM to the BMC, the latter further interacts with the corresponding resources as shown on the interactions tagged as (1b).

For managing the power state of the VM-based virtualisation containers, the VIM interacts with the virtualization layer, i.e. hypervisor, as indicated on the interaction (2a). This interaction can be supported by interfaces such as the libvirt-domain API specified in [i.40]. If needed, the virtualization layer further interacts with the VM as shown in the interaction tagged as (2b).

For managing the power state of a CPU, two variants are considered depending on the form of virtualization:

- For VM-based virtualization, the VIM instructs the virtualization layer's OS, as shown on the interaction tagged as (3a) to change the power state of the CPU (interaction (3c)).
- For OS-based virtualization supporting OS container-based deployments, the CCM, directly (interaction (3b')) or via the CISM (interaction (3b")), instructs the CIS node's OS to change the power state of the CPU (interaction (3c)). If the CIS cluster node is realized as a VM, the relevant entity (e.g. hypervisor) instructs the virtualization layer's OS to change the power state of the underlying CPU.

7.2.2.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #1.2; and
- Key issue #3.3.

7.2.2.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #2.1: The NFV-MANO (namely PIM, VIM and CCM as depicted in figure 7.2.2.2-1) to support invoking operations to manage the power state of corresponding managed components in the NFVI.

7.2.3 Solution #3: Resources reservation for VNFs under power management

7.2.3.1 Introduction

As introduced in key issue #2.2 and relevant use cases, it is assumed that VNF instances can be restored from power saving states. The virtualised resources/containerized workloads used for the deployment of the VNF are aspects to be considered to ensure the availability and restoration of the VNF.

A VNF, once deployed, is comprised of one or more VNFC instances supported (or realized) by resources in the NFVI. In the case of VM-based virtualization, VNFC instances are realized by VMs; in the case of OS container-based virtualization, VNFC instances are realized by containerized workloads in the form of sets of one or multiple OS containers (e.g. a Pod). A VNF instance could undergo changes in their resources structure due to power management actions. For instance, to decrease the power consumption of a VNF instance, a set of VNFC instances can be terminated. Once it is determined that the power state of the VNF instance can get back to normal consumption, the set of terminated VNFC instances are re-instantiated.

The present solution addresses aspects to ensure the availability and restoration of the VNF instance resources under power state management actions.

7.2.3.2 Solution description

The present solution considers leveraging virtualised resources reservation capabilities provided by the NFV-MANO framework to guarantee the availability of resources for a VNF instance when such VNF instance undergoes changes in its resources structure (e.g. number and composition of the virtualised resources) due to power management actions. ETSI GS NFV-IFA 010 [i.42] specifies functional requirements of the NFV-MANO functional blocks and functions, including specific requirements regarding virtualised resources reservation. ETSI GS NFV-IFA 005 [i.43] specifies the requirements and interfaces produced by the VIM related to virtualised resources reservation management.

For the reservation of virtualised resources, two cases are devised:

- Case #1: prior to perform the modification of the virtualised resource, i.e. to terminate such a resource, the NFVO requests to the VIM to reserve a virtualised resource with equivalent characteristics as the one to be terminated.
- Case #2: during the step performing the termination of the virtualised resource, the reservation of a virtualised resource with equivalent characteristics is done simultaneously.

In the case #1, additional factors are to be considered. Depending on the capacity of the NFVI and due to the current consumption of virtualised resources, including the ones from the VNF instance being modified, it can be that reserving the requested virtualised resources is not successful. For instance, assuming a simple scenario in which the NFVI's capacity is comprised of 4 virtualised resources, and all these resources are currently being used; a request to reserve a new virtualised resource will not be possible if there is no information available for the infrastructure manager that one of the virtualised resources is expected to be terminated, and thus NFVI capacity be made available at a future point in time. This issue could be addressed by providing the capability to indicate in the virtualised resource reservation request information indicating that an equivalent virtualised resource with the same characteristics will be terminated at a later stage.

In the case #2, currently the NFV-MANO framework does not support the capability of the VNFM to perform and be authorized to request virtualised resources reservation management (refer to ETSI GS NFV-IFA 006 [i.44]); only the NFVO is capable and authorized to (refer to ETSI GS NFV-IFA 005 [i.43] and ETSI GS NFV-IFA 010 [i.42]). Changing the responsibility of functional blocks/functions in the NFV-MANO framework regarding virtualised resources reservation is not expected by the present solution. The present solution also assumes, as supported by the NFV-MANO framework, that the NFVO can discover or learn about which virtualised resources are affected by the VNF power management processes through VNF LCM operation granting mechanisms.

In the scenario of containerized workload resources, the CISM offers the capability to set quotas on namespaces (see ETSI GS NFV-IFA 010 [i.42] and ETSI GS NFV-IFA 040 [i.45]). In this case, if the quota is not modified during the power management processes, the same quota value applies before and after the modification of the containerized workloads. In such a case, it is expected that re-instantiating the containerized workloads at a future stage will be possible as long as the request and limit values are kept within the boundaries supported by the established quota and such quota has not been modified.

Figure 7.2.3.2-1 illustrates a high-level overview of interactions concerning VNFM, NFVO and VIM for the case of VM-based VNFs. Two major processes are depicted:

- (A) the steps during the power management action to decrease power consumption, and
- (B) the steps when returning the power consumption levels of the VNF to "normal".

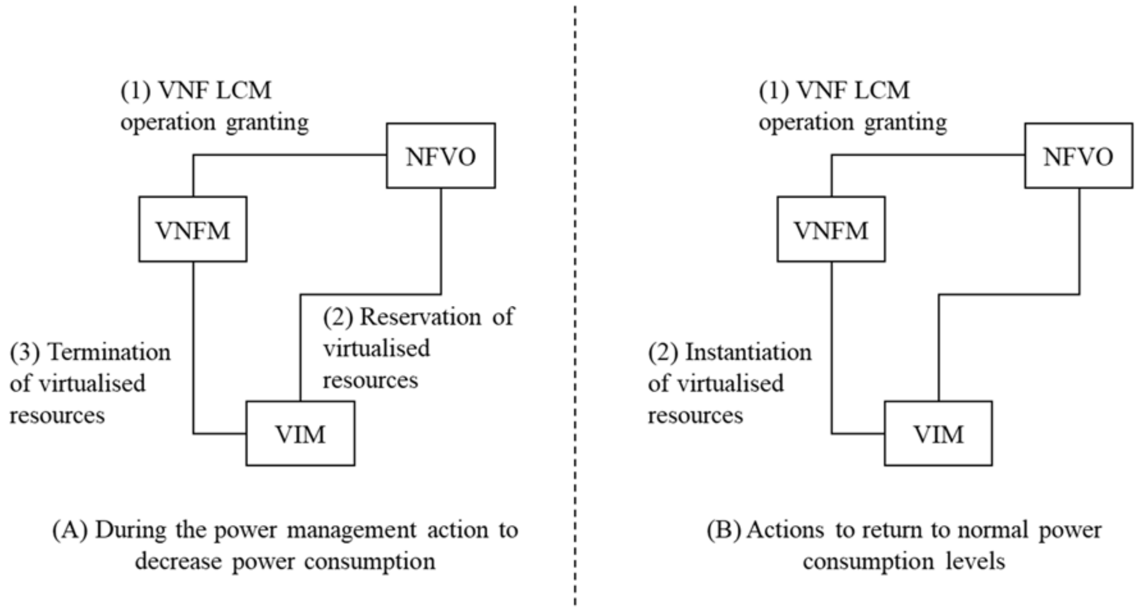

Figure 7.2.3.2-1: High-level interactions in Solution #3

In (A), the following key steps are realized:

(1) A VNF LCM operation has been processed by the VNFM which concerns VNF power state management, and following standard VNF LCM procedures, the VNFM requests a granting to the NFVO. In the granting request, the VNFM provides information about which virtualised resources are expected to be terminated as part of the VNF LCM.

(2) The NFVO requests the VIM the reservation of virtualised resources of equivalent characteristics as the ones that are expected to be terminated. After this step, the NFVO can approve the granting request.

(3) The VNFM proceeds with requesting the VIM the termination of the relevant virtualised resources.

In (B), the key steps are:

(1) A VNF LCM operation is processed by the VNFM, e.g. to return the VNF power state to "normal" levels. Following standard VNF LCM procedures, the VNFM requests a granting to the NFVO. The granting request indicates the virtualised resources that are expected to be (re-)instantiated. The NFVO checks the granting request and approves it, returning along the identifiers of the virtualised resources that have been reserved in the process (A).

(2) The VNFM proceeds with requesting the VIM the instantiation of the relevant virtualised resources by providing the information about the virtualised resources that have been reserved and that are to be used for allocation during the instantiation.

7.2.3.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #2.2.

7.2.3.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #3.1: The VIM and produced virtualised resource reservation interfaces to provide the capability to indicate during virtualised resource reservation requests that certain resources will be terminated. The information about resources that are expected to be terminated can provide additional information to the VIM about capacity availability to be considered for virtualised resources reservation.

Gap #3.2: The VIM and its produced virtualised resource reservation interface to provide the capability to reserve a virtualised resource with equivalent characteristics to the one that is being requested to be terminated.

7.2.4 Solution #4: State and configuration management during VNF power management

7.2.4.1 Introduction

As introduced in key issue #2.2 and relevant use cases, it is assumed that VNF instances can be restored from power saving states. State and configuration of the VNF and virtualised resources/containerized workloads used for the deployment of the VNF are aspects to be considered to ensure the availability and restoration of the VNF.

A VNF, once deployed, is comprised of one or more VNFC instances supported (or realized) by resources in the NFVI. For its operation, the VNF instance is also configured; the whole configuration can include virtualization-dependent and virtualization-independent configuration data, as described by ETSI GR NFV-EVE 022 [i.47]. For the VNF to realize its service and network functionality, the VNF typically makes use of state information. A certain amount of state information can be expected to be available to the VNF instance locally, while some state information can be kept and stored separately from the VNF instance. In case that VNF resources are affected during the processes of VNF power management, e.g. certain components are terminated, then the state and configuration data hold by these resources can also be impacted.

The present solution addresses aspects to ensure the availability and restoration of the VNF state and configuration data under power state management actions.

7.2.4.2 Solution description

The present solution set considers mechanisms to store the state and configuration data of a VNF instance when it is known that such data can be impacted by the VNF power state management actions.

To enable the recovery of the state and configuration data, it is expected that the processes to backup such data will be performed before the virtualised resources and/or containerized workloads are affected. After the VNF instance is restored to its "normal" power state, the state and configuration data can be restored.

Different sub-solutions are devised, such as:

- Solution #4.1: In case the VNF instance uses and/or is managed by VNF generic OAM functions, the VNF configuration manager collects and backups all the configuration and state data of the VNF. In turn, the VNF configuration manager can either realize or interact with a Configuration Server, as reported in ETSI GR NFV-EVE 022 [i.47] for the purpose of storing/backing up the collected data.

- Solution #4.2: Use the VNF snapshotting feature supported by the VNFM, as specified in ETSI GS NFV-IFA 010 [i.42] and ETSI GS NFV-IFA 007 [i.46].

- Solution #4.3: Bind the configuration and state data of the VNF to a specific subset of VNFCs that are not part of the subset of resources that are affected during the VNF power state management action. In this case, the VNFM is expected to know which VNFC cannot be terminated. It is noted that actions other than terminating the VNFC instance, for instance, changing the CPU states of the compute resources used by the VNFC instance, could still be possible.

7.2.4.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #2.2.

7.2.4.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #4.1: The VNFD to support indicating which VNFC cannot be terminated during processes of resource optimization, such as those triggered during VNF state management operations.

NOTE: The support of "Configuration Server" functionality as reported in ETSI GR NFV-EVE 022 [i.47] is not reported as a gap considering that such capability specification can be proceeded by the corresponding feature work that triggered ETSI GR NFV-EVE 022 [i.47].

7.2.5 Solution #5: Power profiles of a VNF

7.2.5.1 Introduction

As introduced in key issue #3.1 and relevant use cases, it is assumed that VNF instances can be managed in terms of their power states. Regarding this feature, aspects to be considered include how to determine which aspects of a VNF are subject to change the power consumption, and how the expected power consumption can be described, so that NFV-MANO can execute corresponding energy efficiency policies.

7.2.5.2 Solution description

The present solution considers the capability provided by the VNFD for a VNF designer to define "power profiles", which indicate, for each power state a VNF can be set into, the following information:

- A description of the power profile including its name, e.g. sleep-mode, low-power, high-performance, etc., and its properties;
- An (estimated or lab-test computed) value of power consumption, e.g. as computed under certain tested and controlled scenarios;
- The amount and structure of virtualised resources (e.g. VM and/or OS containers), which is defined as part of the scale levels;
- Power state characteristics of virtual compute resources required by the VNF components for each power profile. Examples of these power state characteristics are vCPU power state management policies (e.g. static and/or dynamic), operating P/C states for vCPUs, etc.; and
- The service capacity of the VNF at the defined power state. This information is valuable for the network operator to perform a proper planning of the network in relation to the targeted energy efficiency policies.

The power profiles of the VNF are made available to the NFV-MANO and/or other management functions managing the VNF, e.g. a VNF generic OAM function. This can be done by enhancing the VNFD, specified in ETSI GS NFV-IFA 011 [i.48], with the description of the power profiles supported by the VNF. This can also enable the case that such profiles can be mapped/bound to specific scale levels of a VNF and to the existing description of requirements regarding the virtualised resources requirements of the VNF. Regarding the information about service capacity of the VNF, this might be provided as part of the VNFD or in the form of additional artifacts in the VNF Package.

With the VNF power profiles information in the VNFD, the NFV-MANO can process the power consumption characteristics of the VNF as part of the VNF LCM and perform the execution of energy efficiency policies.

Figure 7.2.5.2-1 illustrates a high-level representation of VNFD enhancement with the description of VNF power profiles.

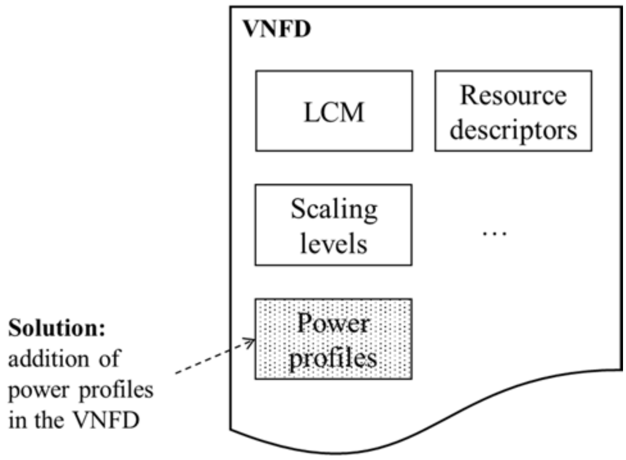

Figure 7.2.5.2-1: Enhancing the VNFD with VNF power profiles

7.2.5.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #3.1.

7.2.5.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #5.1: The VNFD to provide the capability to describe power profiles of a VNF.

7.2.6 Solution #6: VNF power manager/controller

7.2.6.1 Introduction

As introduced in key issue #3.1 and relevant use cases, it is assumed that VNF instances can be managed in terms of their power states. An aspect to consider in this use case is defining which entities will be able to control and manage the power states of the VNF. Furthermore, key issue #4.3 and relevant use cases highlight the need of monitoring the VNF instances in terms of power consumption and energy. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements of VNF instances, and which interfaces are involved in such an acquisition.

7.2.6.2 Solution description

The present solution considers the existence of a management function or controller responsible for managing the power states of a VNF. It can do so based on the information of power profiles, as described in Solution #5. The management function is also capable to process and provide power consumption information and relevant measurements about the VNF instances that are managed.

Different sub-options can be considered regarding the placement of the VNF power management functionality:

- Solution #6.1: as a VNF generic OAM function, e.g. a VNF power management function. VNF generic OAM functions framework is studied in ETSI GR NFV-EVE 019 [i.49], and interfaces and information model are specified in ETSI GS NFV-IFA 049 [i.50].

- Solution #6.2: management functionality added to the VNFM, e.g. as part of the overall VNF LCM functionality specified in ETSI GS NFV-IFA 010 [i.42] and ETSI GS NFV-IFA 007 [i.46].

In both sub-solutions, the solution considers that the entity responsible for the VNF power management exposes standard interfaces or operations to enable the power management of VNF instances. In the case of Solution #6.1, the VNF power management function produces and exposes a VNF power management interface. In the case of Solution #6.2, the VNF LCM interface can be leveraged and enhanced to support relevant operations.

The standard interfaces (depicted as Itf-a and Itf-b in figure 7.2.6.2-1) and operations enable a consumer to:

- Change the power state/mode of the VNF instance, e.g. set the VNF instance into sleep mode, low-power mode, etc. (through Itf-a);
- Query about the current power state/mode of the VNF instance (through Itf-a); and
- Retrieve power consumption and energy related KPIs of the VNF (through Itf-b).

Figure 7.2.6.2-1 illustrates a high-level representation of the sub-options described in the present solution.

NOTE 1: Figure 7.2.6.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the VNF power management function as main actor in the present solution.

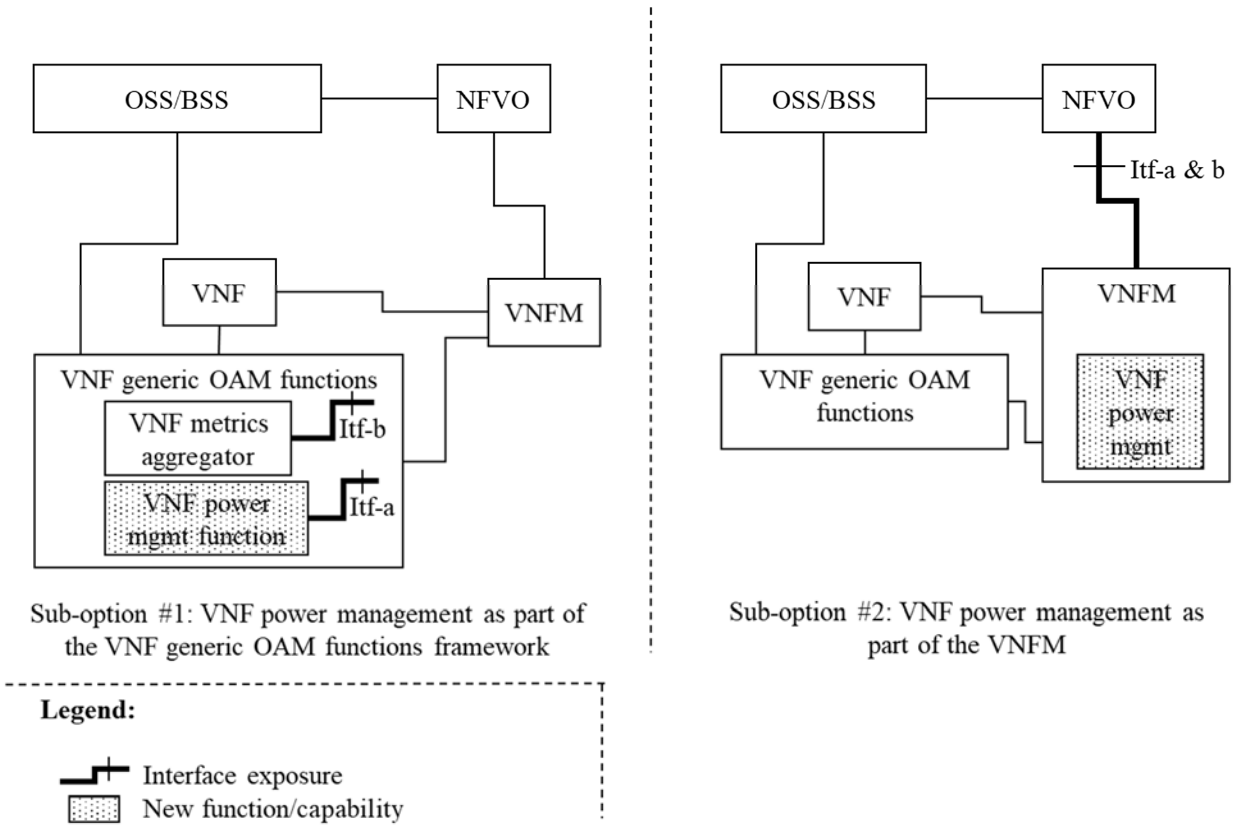

Figure 7.2.6.2-1: Sub-options in the Solution #6

With regards to the exposure of the power consumption and energy metrics, the solution envisions that performance management interfaces are exposed. This is depicted in figure 7.2.6.2-1 as Itf-b. In the case of integrating the VNF power management function as a VNF generic OAM function, the Metrics aggregator VNF generic OAM function produces the Metrics Exposure Interface, as specified in ETSI GS NFV-IFA 049 [i.50], which can be extended to provide the power consumption and energy metrics. In the case of integrating the VNF power management function into the VNFM, the solution envisions the reuse of the VNF Performance Management interface exposed by the VNFM (refer to clause 7.4 of ETSI GS NFV-IFA 007 [i.46] or clause 7.4 of ETSI GS NFV-IFA 008 [i.58]). Furthermore, the solution envisions the definition of new performance metrics related to power consumption and energy information associated to VNF instances. The definition of performance measurements related to power consumption and energy information leverages the performance measurements specification of ETSI GS NFV-IFA 027 [i.57].

NOTE 2: The present solution does not focus on the aspects about how to compute the power consumption of the corresponding managed objects, which is defined by other solutions described in the present document.

7.2.6.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #3.1; and
- Key issue #4.3.

7.2.6.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #6.1: In case the VNF power management function is part of the VNFM, the VNFM to support managing the power state or modes of a VNF instance.

Gap #6.2: In case the VNF power management function is part of the VNF generic OAM framework, the VNF generic OAM functions framework to define and support a specific VNF generic OAM function responsible for managing the VNF power state or modes.

Gap #6.3: In case the VNF power management function is part of the VNFM, the VNFM to support providing power consumption and energy metrics associated to its managed objects via the VNF Performance Management interface.

Gap #6.4: In case the VNF power management function is part of the VNF generic OAM framework, the Metrics aggregator function to support providing power consumption and energy metrics associated to its managed objects (i.e. the VNF) via the Metrics Exposure interface.

7.2.7 Solution #7: Power state based zones in NFVI for static and dynamic power state management of compute resources

7.2.7.1 Introduction

Use cases related to power-state driven VNF/NS deployment and dynamic power state management of compute resources discussed in clause 5 assume some power state management to be performed in the NFVI. As discussed in key issue #3.3 related to NFVI driven energy efficiency, a mechanism for pooling physical and virtual compute resources can be considered in the NFVI to enable power-state driven VNF/NS deployments.

Use case #5, described in clause 5, assumes that once the VNF/NS instances have been deployed using the pre-set CPU power states (P-states and C-states), the set of underlying compute resources remains in the required power states. On the other hand, use case #8 suggests that the runtime power states of underlying compute resources belonging to a VNF/NS instance can be dynamically altered based on CPU utilization, to save power even further. These scenarios can be realized by maintaining logical zones in the NFVI, where the physical resources can be pooled based on power state management policies.

The present solution addresses aspects related to resource pooling in the NFVI, enabling NFV-MANO to perform power state driven deployments and avoid conflict between static and dynamic power state management policies for underlying compute resources.

NOTE: The present solution focuses on power state management of compute resources. The solution might also, either fully or partially, be applicable to other kinds of NFVI resources, such as storage and network, but this is not analyzed by the solution.

7.2.7.2 Solution description

The present solution proposes to categorize the NFVI compute resources into static and dynamic zones based on the CPU power management policies, i.e. static or dynamic modes of CPU power state management. This kind of zoning can be done on multiple levels considering the scale of the NFVI. The static and dynamic zones can be maintained on either or both of the following levels:

- Cluster/resource zone level; and
- Node level.

Cluster/resource zone level would mean categorization of nodes or servers that either belong to a static or a dynamic zone. An entity such as the VIM, the CCM or the PIM can manage these zones on a need basis.

Node level zoning implies that some of the CPUs (or CPU cores) are assigned to the static zone, while others are associated to the dynamic zone. An NFVI subsystem operating on the node level (such as a hypervisor or an OS scheduler) can manage these zones based on actual demand and runtime utilization.

The two zones are described below:

1) Static zone: The power states of CPU cores belonging to the static zone cannot be adjusted dynamically at runtime during the lifecycle of the VNF/VNFC instance (see use case #5 in clause 5.6). The CPU operating states (P-states and C-states) are determined by the vCPU profiles described in the VNFD of the associated VNF/VNFC.

2) Dynamic zone: The power states of CPU cores in the dynamic zone can be adjusted dynamically and autonomously, based on runtime CPU utilization, between a range of P-states to conserve power even further (see use case #8 in clause 5.9). The range of operating P/C-states is expected to be provided in the VNFD for each VNFC.

The NFV-MANO system is expected to be aware of the zones that are being maintained in the NFVI, and thereby, the NFV-MANO system is able to instantiate VNF/NS components on the compute resources belonging to the appropriate zones. In case of statically managed CPUs, the NFV-MANO system provides the operating P/C-states for the vCPU(s) based on the information provided in the VNFD (see Solution #5 in clause 7.2.5). While for the dynamically managed CPUs, a range of CPU power states is advertised to the NFVI subsystem responsible for adjusting the CPU power states based on runtime utilization.

For the case of OS container-based deployments, the CIS cluster node tagging scopes can be used to categorize CIS cluster nodes into static and dynamic zones, and further expressing CPU profiles for each node, such as supported P and C-states.

7.2.7.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #3.3;
- Key issue #2.1; and
- Key issue #2.3.

7.2.7.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #7.1: The NFV-MANO to support power state management-based pooling of physical resources in the NFVI.

Gap #7.2: The NFV-MANO to support discovery, storage and use of information about power state aware zoning and pooling of physical resources in the NFVI.

Gap #7.3: The NFV-MANO entities (i.e. PIM, VIM and CCM) to support invoking operations to manage power states of corresponding resource pools in the NFVI.

Gap #7.4: The VNFD to provide the capability to describe CPU power management related information and constraints on a per VNFC basis.

7.2.8 Solution #8: Power-state aware allocation of compute resources for VNF/NS orchestration and lifecycle management

7.2.8.1 Introduction

Key issue #2.1, described in clause 6, raises questions about leveraging NFV management and orchestration capabilities for applying energy efficiency policies and determining optimal placement of VNF/NS instances according to those policies, thereby delivering power consumption savings. To support the concepts of power state driven VNF/NS deployment and mapping logical VNF/NS and physical resources for optimized power consumption, discussed in use case #5 and use case #6 respectively, NFV-MANO can consider power state requirements while allocating compute resources for VNF/NS components.

A VNF instance can be comprised of multiple VNFCs, each with its own set of compute requirement parameters such as number of CPUs, affinity/anti-affinity constraints, CPU pinning, etc. These characteristics are described for both VM-based and OS container-based VNFs in their corresponding fields in the VNFD, as specified in ETSI GS NFV-IFA 011 [i.48]. Incorporating CPU power states while allocating compute resources for VNFCs can enable the NFV-MANO system to control per-VNF power consumption on a more granular level.

The present solution addresses aspects related to power state aware allocation of virtualised compute resources while instantiating or scaling (out) VNF/NS instances.

NOTE: The present solution focuses on power state management of compute resources. The solution might also, either fully or partially, be applicable to other kinds of NFVI resources, such as storage and network, but this is not analyzed by the solution.

7.2.8.2 Solution description

This solution proposes to enhance orchestration and management capabilities of NFV-MANO by introducing power state aware allocation of compute resources for VNF/NS instances. As supported by the NFV framework, a VNF can be realized using VMs or OS containers. For the case of VM-based VNFs, VNFCs are realized using VMs. For OS container based VNFs, VNFCs are realized using containerized workloads from groups of one or more OS containers (e.g. a Pod).

Either realized by a VM or a group of one or more OS containers, compute characteristics of a VNFC can be described using a VDU, as specified in ETSI GS NFV-IFA 011 [i.48]. In the present solution, the VDU for each VNFC contains information related to power states of its CPUs, such as power state management policy (static or dynamic), desired P/C-states for CPUs, etc. This information is used by NFV-MANO while allocating compute resources in the NFVI.

Following sub-solutions can be devised to implement power state aware allocation of compute resources:

- Solution #8.1: VDUs in the VNFDs carry additional information related to CPU power state requirements in the relevant fields for both VMs and OS containers realizing the VDUs. These CPU power profiles of VDUs can be linked with the relevant VNF power profiles in the VNFD as described in Solution #5.
- Solution #8.2: For the case of VM-based VNFs, the NFVO and the VNFM provide power state related parameters, such as CPU states management policy, desired P/C-states per vCPU, etc. while requesting the VIM to allocate compute resources for VNF instances.
- Solution #8.3: For the case of containerized workloads, NFV-MANO selects the appropriate CIS cluster that meets the desired CPU power profile indicated in the VNFD, similar to the CIS cluster selection and VNF placement workflow specified in ETSI GS NFV-IFA 036 [i.51]. Supported CPU profiles can be advertised using CIS cluster node tags, as specified in ETSI GS NFV-IFA 036 [i.51]. Furthermore, the VNFM provides power state related parameters while requesting CPU resources for the OS container, e.g. CPU states management policy, desired P/C-states for the CPU, etc. to the CISM. The CISM is expected to deploy the containerized VNFCs on CIS cluster nodes that support the desired CPU power states.

Combination of above-mentioned sub-solutions can enable power state aware allocation of compute resources during lifecycle management operations for VM-based and OS container-based VNFs.

7.2.8.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #2.1; and
- Key issue #3.1.

7.2.8.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #8.1: The information elements over VIM produced interfaces, as specified in ETSI GS NFV-IFA 005 [i.43] and ETSI GS NFV-IFA 006 [i.44], for allocation of virtualised compute resources to provide the means to convey per-CPU power state requirements in the relevant information elements.

Gap #8.2: Operations and relevant information elements related to orchestration and lifecycle management of VNF/NS instances to consider the exchange of information related to CPU power profiles for VNFCs in several reference points and interfaces within NFV-MANO, for example the Or-Vnfm reference point specified in ETSI GS NFV-IFA 007 [i.46].

Gap #8.3: The VDU descriptions in the VNFD, as specified in ETSI GS NFV-IFA 011 [i.48], to describe CPU power profiles for VMs or OS containers realizing the VDU, such as power state management policy, desired P/C-states for CPUs, etc.

Gap #8.4: The CIS cluster node tagging scopes, specified in ETSI GS NFV-IFA 036 [i.51], to consider how to express CIS cluster node's supported CPU power profiles.

7.2.9 Solution #9: Energy efficiency policies

7.2.9.1 Introduction

As introduced in key issue #2.1 and relevant use cases, NFV management and orchestration capabilities can be used to apply energy efficiency policies. NFV-MANO can execute relevant policies to enforce the objectives of the network operator in achieving proper power consumption of VNF and NS instances.

7.2.9.2 Solution description

The present solution considers the modelling, design and use of NFV-MANO policies specific to energy efficiency. As reported by ETSI GR NFV-IFA 023 [i.52] and ETSI GR NFV-IFA 042 [i.56] and later by the NFV-MANO reference point interface specifications (e.g. ETSI GS NFV-IFA 007 [i.46]) and the NFV-MANO policy information modelling in ETSI GS NFV-IFA 048 [i.53], NFV-MANO supports both, policy management interfaces enabling the transfer, activation, termination, etc. of policies between a Policy Administration Point (PAP) and the NFV-MANO functional blocks and functions, acting as Policy Functions (PF), as well as the modelling of policies. The modelling in ETSI GS NFV-IFA 048 [i.53] supports Event-Condition-Action (ECA) type of policies.

Table 7.2.9.2-1 describes the current and potential NFV-MANO functional blocks, as PFs for energy efficiency related policies, that are considered by the present solution, as well as the scope of their policies.

Table 7.2.9.2-1: Energy efficiency policy, managed objects and examples

PF	Managed objects	Policy examples (see note)
NFVO	NS, VNF, NFVI resources	• If power consumption in NFVI-PoP #1 is above a defined threshold, then deploy new NS/VNF instances to an alternate NFVI-PoP. • If there is potential power outage or decrease in capacity is expected in NFVI-PoP #1 where VNF instance #1 is instantiated, then change the power state of VNF instance #1 to a lower power consumption state (see solutions #5 and #6 described in clause 7.2.5 and clause 7.2.6, respectively).

PF	Managed objects	Policy examples (see note)
VNFM	VNF, VNFC	• If power consumption of VNF instance #1 is above a specified threshold and time of day is within a certain time period, then scale down VNF instance #1 to a smaller scale/instantiation level. • If VNF instance #2 CPU metrics indicate low CPU usage and VNF indicator #A shows low capacity usage of the VNF instance, then request the CISM/VIM to change the power state of the CPU supporting the containerized workload/virtualised compute resources of the VNF to a low-power consuming state.
VIM	Virtualised resources	• If a request to instantiate a VM is received requiring X amount of vCPU at frequency Y and time of day is within certain time period, then allocate the VM on a resource pool in the NFVI based on CPU resources of generation Z, where generation Z is assumed to be more efficient in terms of power consumption.
CISM	Containerized workloads	• If a request to instantiate a containerized workload is received requiring X amount of vCPU and time of day is within a certain time period, then allocate the containerized workload on a CIS cluster node labelled with label Z, where label Z is assumed to be associated to CIS cluster nodes which have CPUs that are more efficient in terms of power consumption.
WIM	MSCS, MSNC	• If a request to instantiate an MSCS is received to interconnect service connection points of NFVI-PoP (sites) #1 and #2, then setup the necessary MSNC over the WAN avoiding network devices in domain area X, in which is known that domain area X is powered by a more expensive power source.
PIM	Physical resources	• If compute resource #1 has no workloads or virtualised resources instantiated on it and the time of day is within time period X, then power down the compute resource.
NOTE: The list of examples is not meant to be exhaustive, and these are only for illustrative purposes.		

To further support the solution, it is assumed that the NFV-MANO functional blocks and functions can collect KPIs related to power consumption associated to the diverse managed objects. The solution also assumes that VIM, CISM and PIM provide interfaces with the capability to request changing power states of NFVI resources associated to the virtualised resources and containerized workloads realizing VNF and NS instances.

7.2.9.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #2.1.

7.2.9.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #9.1: The NFV-MANO to support capabilities related to collecting metrics related to power consumption and computing power consumption associated to its managed objects (e.g. a VNF instance).

NOTE: ETSI NFV specifications about policy models do not constrain the scope of policies that can be provided as long as necessary events, conditions and actions of the policies can be mapped to corresponding capabilities supported by NFV-MANO's functional requirements, interfaces and models.

7.2.10 Solution #10: NFV-MANO state management for energy efficiency

7.2.10.1 Introduction

As introduced in key issue #3.2 and relevant use cases, the NFV-MANO system is expected to be managed to address energy efficiency and power consumption objectives. Under some scenarios, it might be necessary to restrict the capabilities of the NFV-MANO system in order to decrease the power consumption of the NFV-MANO system or of its managed objects, such as VNF and NS instances.

7.2.10.2 Solution description

The present solution considers the capability of NFV-MANO functional blocks and functions to be managed, as specified by the functional requirements in clause 10 of ETSI GS NFV-IFA 010 [i.42]. As specified in ETSI GS NFV-IFA 031 [i.54], the "NFV-MANO configuration and information management interface" supports the "Change State operation", which enables an interface consumer to change the operational and administrative state of the NFV-MANO functional entity application and/or its provided NFV-MANO service interfaces. In terms of operational state, the following three actions are supported:

- START: to start the managed entity (i.e. NFV-MANO functional entity application and NFV-MANO service interface).
- STOP: to stop the managed entity.
- RESTART: to stop and start again the managed entity.

By stopping the whole application or specific NFV-MANO services, power consumption incurred by the NFV-MANO functional entity can be decreased.

In terms of administrative states, the following actions are supported by the referenced ETSI NFV specifications:

- LOCK: to lock the managed entity (i.e. NFV-MANO functional entity application and NFV-MANO service interface).
- UNLOCK: to unlock the managed entity.

The meaning of administrative states is specified in detail in clause 5.7 of ETSI GS NFV-SOL 009 [i.55]. In summary, when the managed entity is in locking state, the entity is not allowed to handle requests for creating new managed object instances (e.g. a VNFM to create a new VNF instance), while service is maintained for existing managed objects. In the locked state, the managed entity is completely discharged from service and not allowed to handle requests for creating new managed object instances.

To make a more granular adjustment of the administrative state with the purpose to fine-tune the power consumption by the NFV-MANO functional entity, the present solution considers the capability to enable suspending the handling of requests on NFV-MANO service interfaces produced by the NFV-MANO functional entity for existing managed objects:

- SUSPEND: to suspend the managed entity, i.e. enables requesting the NFV-MANO functional entity to suspend processing any new requests on existing managed objects being handled by a certain NFV-MANO service interface.
- RESUME: to resume the operations and processing on the managed entity.

NOTE: It is assumed that proper evaluation has been done by the consumer of the NFV-MANO management capabilities when requesting operational and administrative state changes to NFV-MANO functional entities to avoid impact on critical services.

Furthermore, the peer NFV-MANO functional entity consumers can also be aware of the administrative and operational state of the producer NFV-MANO functional entity. The solution also envisions the capability to perform state management from an NFV-MANO service interface consumer perspective. In this regard, the suspension/resume of the consumption of NFV-MANO service interfaces is defined:

- SUSPEND: to suspend in an NFV-MANO functional entity consumer the consumption of an NFV-MANO service interface produced by a peer NFV-MANO functional entity, i.e. enables requesting the NFV-MANO functional entity to suspend consuming, therefore, generating any new requests, to a certain NFV-MANO service interface produced by the peer entity.
- RESUME: to resume the consumption of the NFV-MANO service interface.

By adjusting the administrative states, it is expected that less processing cycles are needed by the NFV-MANO functional entity due to the locking and/or suspension of the tasks and therefore of processing requirements by the managed entity. A decrease in processing requirements can potentially decrease the power consumption incurred by the NFV-MANO functionality application.

7.2.10.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #3.2.

7.2.10.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #10.1: The NFV-MANO management to support suspending a managed entity such as NFV-MANO functional entity application and NFV-MANO service interfaces.

NOTE: Suspension stops processing any new requests on existing managed objects being handled by a certain NFV-MANO service interface.

Gap #10.2: The NFV-MANO management to support suspending the consumption of NFV-MANO service interfaces produced by a peer NFV-MANO functional entity at a targeted NFV-MANO functional entity consumer.

7.2.11 Solution #11: NFV-MANO capabilities for energy efficiency

7.2.11.1 Introduction

As introduced in key issue #2.1 and relevant use cases, NFV-MANO offers a set of capabilities that can be leveraged for applying energy efficiency policies.

7.2.11.2 Solution description

The present solution considers listing the set of capabilities of NFV-MANO that can be leveraged for applying energy efficiency policies. NFV-MANO capabilities related to the managed objects, such as NS, VNF, virtualised resources, etc. can be grouped into two categories:

- Resource-related capabilities, whose actions imply a change in the allocation and use of NFVI resources by the managed objects.
- Non-resource related capabilities, whose actions do not change the use of NFVI resources by the managed objects.

At large, actions implying an increase in usage of NFVI resources, such as instantiation and scaling out, will likely have an effect of increasing the power consumption in the NFVI. Whereas actions implying a decrease in usage of NFVI resources, such as termination and scaling in, will likely have the contrary effect, i.e. decrease the power consumption in the NFVI.

Table 7.2.11.2-1 lists the resource-related capabilities provided by NFV-MANO and its functional blocks and functions by differentiating those that potentially increase/decrease power consumption.

Table 7.2.11.2-1: Resource-related capabilities by NFV-MANO

NFV-MANO entity	Capability	Increasing power consumption	Decreasing power consumption
NFVO	NS LCM	Instantiation, healing and scaling up/out of NS instances.	Termination and scaling down/in of NS instances.
VNFM	VNF LCM	Instantiation, healing, scaling out, start operation and snapshotting of VNF instances.	Termination and scaling in of VNF instances.
VIM	Virtualised resources management	Instantiation and snapshotting of virtualised resources.	Termination of virtualised resources.

NFV-MANO entity	Capability	Increasing power consumption	Decreasing power consumption
CISM	Containerized workloads management (including OS container compute, network and storage workloads)	Instantiation of containerized workloads.	Termination of containerized workloads.
CCM	CIS cluster LCM	Instantiation and scaling out of CIS clusters.	Termination and scaling in of CIS clusters.
WIM	MSCS management	Instantiation and update (adding MSCS endpoints) of MSCS instances.	Termination and update (removing MSCS endpoints) of MSCS instances.

7.2.11.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #2.1.

7.2.11.4 Gap analysis

No gaps are identified.

7.2.12 Solution #12: VIM for acquiring virtualised resource power consumption

7.2.12.1 Introduction

Key issue #4.2 and relevant use cases highlight the need of monitoring the virtualised resources in terms of power consumption and energy. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements about virtualised resources, and which interfaces are involved in such an acquisition.

7.2.12.2 Solution description

The present solution considers enhancing the VIM to make it capable to process and provide power consumption information and relevant energy-related measurements about the managed virtualised resources.

As the VIM is responsible for the management of virtualised resources, it can associate the power consumption and energy metrics to the corresponding managed objects.

NOTE 1: The present solution does not focus on the aspects about how to compute the power consumption of the corresponding managed objects, which is defined by other solutions described in the present document.

With regards to the exposure of the power consumption and energy metrics, the solution envisions the reuse of the Virtualised Resource Performance Management interface exposed by the VIM (refer to clause 7.7 of ETSI GS NFV-IFA 005 [i.43] and clause 7.7 of ETSI GS NFV-IFA 006 [i.44]) and the definition of new performance metrics related to power consumption and energy information associated to the various virtualised resource managed objects such as the virtual compute, virtual network and virtual storage. By defining new metrics and reusing the existing PM interface, the capabilities of monitoring via PM Jobs and Thresholds can be leveraged. The definition of performance measurements related to power consumption and energy information leverages the performance measurements specification of ETSI GS NFV-IFA 027 [i.57].

Figure 7.2.12.2-1 illustrates a high-level representation of the solution depicting the VIM providing relevant metrics to its consumers such as the NFVO, the VNFM, and the CCM over standard interfaces.

NOTE 2: Figure 7.2.12.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the VIM as main actor in the present solution.

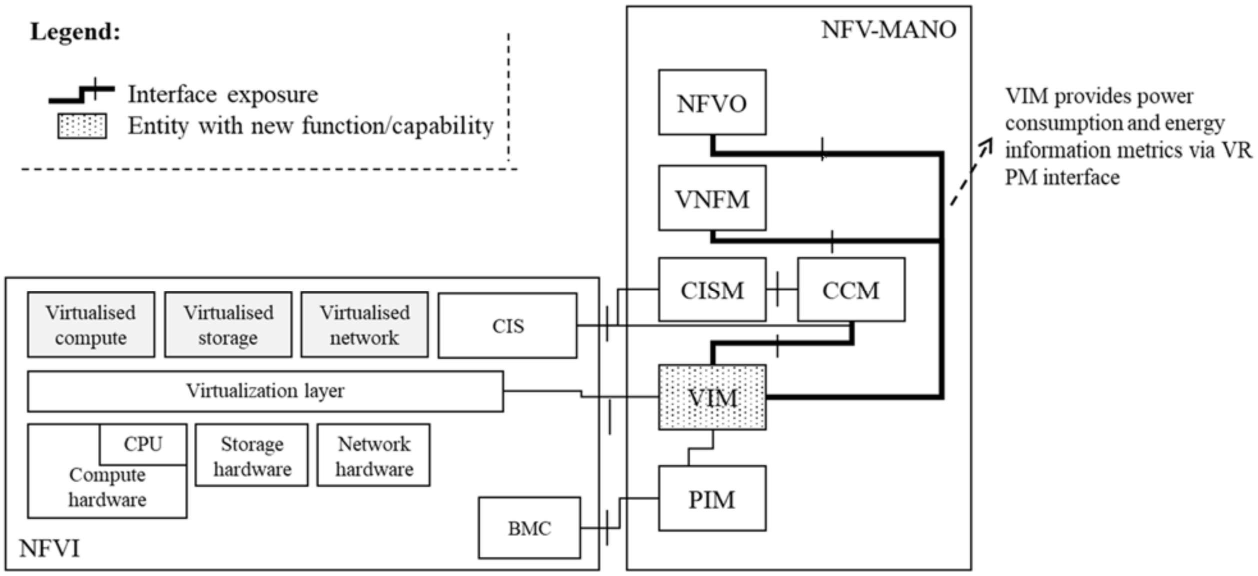

Figure 7.2.12.2-1: High-level representation of Solution #12

7.2.12.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.2.

7.2.12.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #12.1: The VIM to support providing power consumption and energy metrics associated to its managed objects via the Virtualised Resources Performance Management interface.

7.2.13 Solution #13: CISM for acquiring containerized workload power consumption

7.2.13.1 Introduction

Key issue #4.2 and relevant use cases highlight the need of monitoring the containerized workloads in terms of power consumption and energy. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements about containerized workloads, and which interfaces are involved in such an acquisition.

7.2.13.2 Solution description

The present solution considers enhancing the CISM to make it capable to process and provide power consumption information and relevant measurements about the managed containerized workloads.

As the CISM is responsible for the management of containerized workloads, it can associate the power consumption and energy metrics to the corresponding managed objects.

NOTE 1: The present solution does not focus on the aspects about how to compute the power consumption of the corresponding managed objects, which is defined by other solutions described in the present document.

With regards to the exposure of the power consumption and energy metrics, the solution envisions that performance management interface(s) (one or multiple, depending on the granularity of containerized workloads management and their type such as MCIO-C, MCIO-S and MCIO-N types) is exposed by the CISM. Furthermore, the solution envisions the definition of new performance metrics related to power consumption and energy information associated to the various containerized workload managed objects. The definition of performance measurements related to power consumption and energy information leverages the performance measurements specification of ETSI GS NFV-IFA 027 [i.57].

Figure 7.2.13.2-1 illustrates a high-level representation of the solution depicting the CISM providing relevant metrics to its consumers such as the NFVO, the VNFM, and the CCM over standard interfaces.

NOTE 2: Figure 7.2.13.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the CISM as main actor in the present solution.

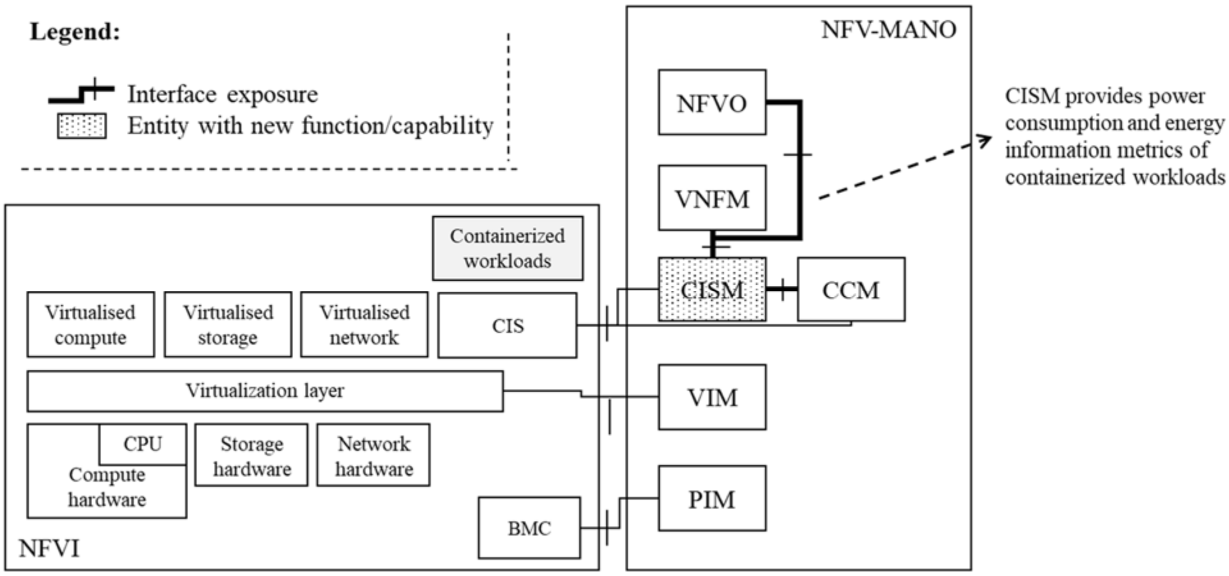

Figure 7.2.13.2-1: High-level representation of Solution #13

7.2.13.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.2.

7.2.13.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #13.1: The CISM to support producing a performance management interface(s) to provide performance data about its managed objects.

Gap #13.2: The CISM to support providing power consumption and energy metrics associated to its managed objects via corresponding interface(s) (see gap #13.1).

7.2.14 Solution #14: NFVO for acquiring NS power consumption

7.2.14.1 Introduction

Key issue #4.4 and relevant use cases highlight the need of monitoring the NS instances in terms of power consumption. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements about NS instances, and which interfaces are involved in such an acquisition.

7.2.14.2 Solution description

The present solution considers enhancing the NFVO to make it capable to process and provide power consumption information and relevant measurements about the managed NS instances.

As the NFVO is responsible for the management of NS instances, it can associate the power consumption and energy metrics to the corresponding managed objects.

NOTE 1: The present solution does not focus on the aspects about how to compute the power consumption of the corresponding managed objects, which is defined by other solutions described in the present document.

With regards to the exposure of the power consumption and energy metrics, the solution envisions the reuse of the NS Performance Management interface exposed by the NFVO (refer to clause 7.5 of ETSI GS NFV-IFA 013 [i.59]) and the definition of new performance metrics related to power consumption and energy information associated to the NS instance. By defining new metrics and reusing the existing PM interface, the capabilities of monitoring via PM Jobs and Thresholds can be leveraged. The definition of performance measurements related to power consumption and energy information leverages the performance measurements specification of ETSI GS NFV-IFA 027 [i.57].

Figure 7.2.14.2-1 illustrates a high-level representation of the solution depicting the NFVO providing relevant metrics to its consumers, such as the OSS/BSS, over standard interfaces.

NOTE 2: Figure 7.2.14.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the NFVO as main actor in the present solution.

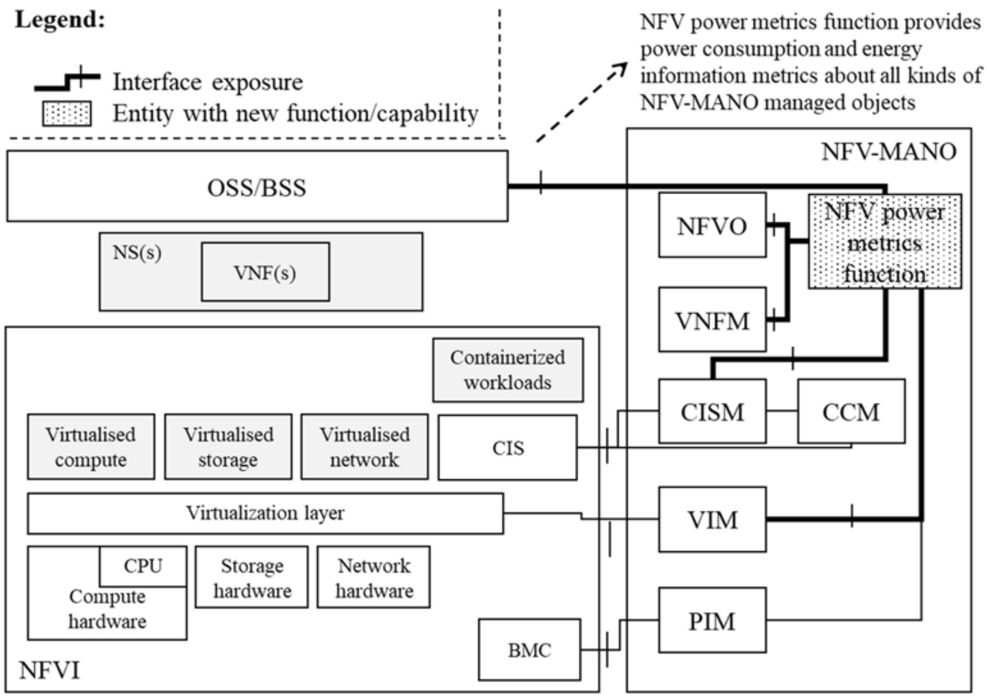

Figure 7.2.14.2-1: High-level representation of Solution #14

7.2.14.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.4.

7.2.14.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #14.1: The NFVO to support providing power consumption and energy metrics associated to its managed objects via the NS Performance Management interface.

7.2.15 Solution #15: NFV power metrics function for acquiring multiple kinds of power consumption data

7.2.15.1 Introduction

Key issues #4.2, #4.3 and #4.4 and relevant use cases highlight the need of monitoring the virtualised resources, containerized workloads, VNF instances and NS instances in terms of power consumption and energy. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements about this set of managed objects, and which interfaces are involved in such an acquisition.

7.2.15.2 Solution description

The present solution considers enhancing the NFV-MANO architectural framework by adding a new dedicated NFV power metrics function. The management function is capable to process and provide power consumption information and relevant measurements about all the objects that are managed by NFV-MANO. The NFV power metrics function might be integrated with other NFV-MANO functional blocks or functions.

NOTE 1: The present solution does not make an assessment of the feasibility of such integration possibility.

The NFV power metrics function is responsible for collecting and computing, when necessary, the power consumption of the corresponding managed objects. The solution assumes that the NFV power metrics function is aware of the managed objects, even though it is not responsible for their management. Which managed objects and which specific instances of them are being handled by NFV-MANO can be known to the NFV power metrics function by interacting with the rest of NFV-MANO functional blocks and functions. For instance, to know which VNF instances are deployed, the NFV power metrics function interacts with one or more VNFMs or NFVOs. To collect the underlying NFVI resource power consumption, the NFV power metrics function interacts with the controller, e.g. BMC, in the NFVI, either directly or via the PIM. The power consumption and energy metrics collected from the NFVI can then be used to associate and compute the power consumption and other energy metrics of the respective NFV-MANO managed objects.

NOTE 2: The present solution does not focus on the aspects about how to compute the power consumption of the corresponding managed objects, which is defined by other solutions described in the present document.

With regards to the exposure of the power consumption and energy metrics, the solution envisions that a power metrics interface is exposed by the NFV power metrics function. Furthermore, the solution envisions the definition of new performance metrics related to power consumption and energy information associated to the various kinds of NFV-MANO managed objects. The definition of performance measurements related to power consumption and energy information leverages the performance measurements specification of ETSI GS NFV-IFA 027 [i.57].

Figure 7.2.15.2-1 illustrates a high-level representation of the solution depicting the NFV power metrics function providing relevant metrics to its consumers such as the NFVO, the VNFM, and the OSS/BSS over standard interfaces.

NOTE 3: Figure 7.2.15.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the NFV power metrics function as main actor in the present solution.

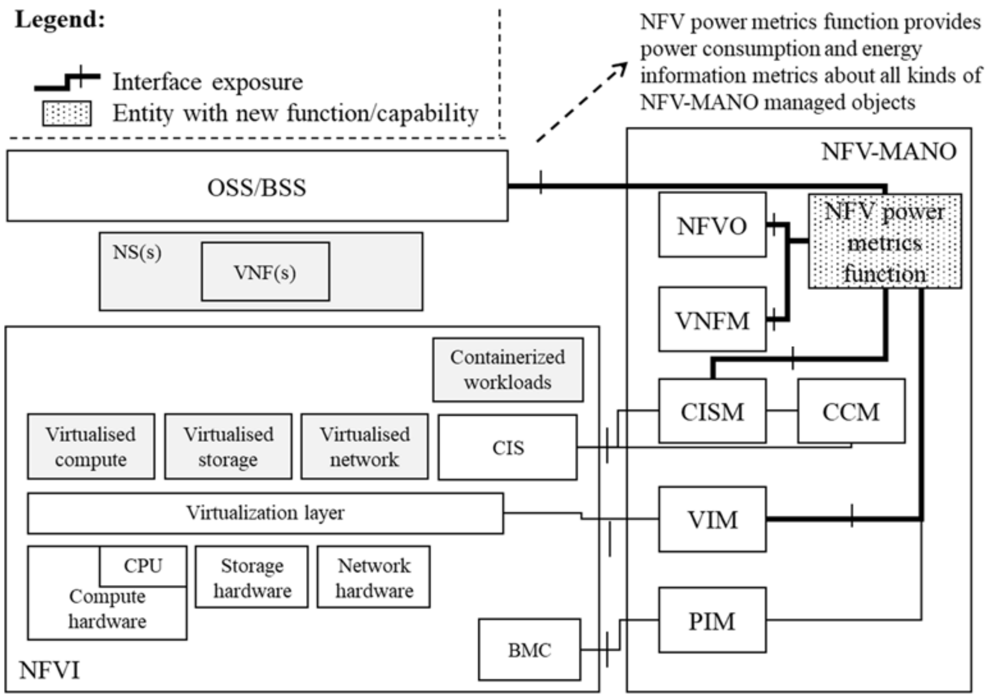

Figure 7.2.15.2-1: High-level representation of Solution #15

7.2.15.3 Key issues addressed

The present solution aims at addressing aspects of the following key issues described in clause 6:

- Key issue #4.2;
- Key issue #4.3; and
- Key issue #4.4.

7.2.15.4 Gap analysis

The referenced ETSI NFV specifications in the present document do not specify:

Gap #15.1: The NFV-MANO architectural framework to support an NFV power metrics function capable to provide power consumption and energy metrics of all kinds of NFV-MANO managed objects.

Gap #15.2: The NFV power metrics function (see gap #15.1) to support producing a power metrics interface.

7.2.16 Solution #16: Energy efficiency of NFVI physical resources

7.2.16.1 Introduction

Key issue #4.1 and relevant use cases identify the need to define energy efficiency for physical resources in the NFVI. For instance, determining the energy efficiency of compute resources in a reliable, accurate and reproducible manner provides to NFV-MANO and the service provider valuable information to compare the efficiency between the resources, which can be a trigger for further developing energy efficiency and power saving policies for NFV-MANO.

7.2.16.2 Solution description

The present solution considers reusing and referencing standard defined energy efficiency metrics. Energy efficiency metric is commonly defined as the relation between the service delivered by a system and the energy consumed by the system when delivering such a service. Depending on the relation calculation factors, either efficiency per power consumption unit, or power consumption per service performance unit can be derived.

For instance, ETSI EN 303 470 [i.20] specifies metrics regarding energy efficiency of compute servers. For the definition of the metrics, the Server Efficiency Rating Tool (SERT®) [i.60] is used. The tool collects measurements and uses an accompanying efficiency metric to assess computer server energy efficiency for various certification and regulatory programs, such as ENERGY STAR® certification and Lot9 Ecodesign Regulation (EU) 2019/424 [i.68] by the European Union.

NOTE 1: SERT® is a registered trademark of Standard Performance Evaluation Corporation (SPEC).

NOTE 2: ENERGY STAR® name and mark are registered trademarks owned by the U.S. Environmental Protection Agency (EPA).

The "Active State Efficiency Metric", which measures the efficiency of a compute server, is specified in clause 5.1.2.2 of ETSI EN 303 470 [i.20]. Clause 5.2 of ETSI EN 303 470 [i.20] specifies the "Idle State Metric" as well. ETSI EN 303 470 [i.20] defines the active state and idle states as follows:

- "active state: operational state of a computer server (as opposed to the idle state) in which the computer server is carrying out work in response to prior or concurrent external requests (e.g. instruction over the network)."
- "idle state: operational state of a computer server in which the operating system and other software have completed loading but is not performing any useful work."

For data storage equipment, ETSI EE also specifies relevant energy efficiency metrics in ETSI EN 303 804 [i.61]. Clauses 7.5 and 7.6 of ETSI EN 303 804 [i.61] specify the "active state periodic energy efficiency" and the "idle state energy efficiency" of the data storage equipment, respectively.

NOTE 3: At the time of developing the present document, ETSI EN 303 804 [i.61] was a draft version, and not a published one, and the referenced content might not be present in the final published version.

For network equipment, clause 11 of Recommendation ITU-T L.1310 [i.66] specifies the energy efficiency ratio for routers and Ethernet switches, in turn based on the metrics adopted from ATIS-0600015.03.2013 [i.67].

7.2.16.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.1.

7.2.16.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #16.1: A metric defining the energy efficiency of NFVI resources, including compute, storage and network.

7.2.17 Solution #17: Energy efficiency of virtualised resources and containerized workloads

7.2.17.1 Introduction

Key issue #4.2 and relevant use cases identify the need to define energy efficiency for virtualised resources and containerized workloads. For instance, determining the energy efficiency of a virtualised compute resource (i.e. a virtual machine) in a reliable, accurate and reproducible manner provides to NFV-MANO and the service provider valuable information to compare the efficiency between the resources, which can be a trigger for further developing energy efficiency and power saving policies for NFV-MANO.

7.2.17.2 Solution description

The present solution considers the definition of an energy efficiency metric of a virtualised compute resource (Eff_{vm}) using the following formula:

$$Eff_{vm} = \frac{Perf_{vm}}{Pow_{vm}}$$

wherein $Perf_{vm}$ corresponds to the service performance of the virtual compute resource and Pow_{vm} is the power consumption associated to the virtual compute resource over a defined time interval.

Similarly, the energy efficiency metric for a compute-related containerized workload (Eff_{cw}) (i.e. a group of one or more OS containers, e.g. a Pod in the case of Kubernetes®) can be defined as:

$$Eff_{cw} = \frac{Perf_{cw}}{Pow_{cw}}$$

wherein $Perf_{cw}$ corresponds to the service performance of the group of one or more OS containers and Pow_{cw} is the power consumption associated to the group of one or more OS containers.

NOTE 1: For the definition of power consumption metrics of virtualised compute resource and group of one or more OS containers refer to solution #18 in clause 7.2.18.

NOTE 2: The present solution does not define how the performance value associated to a virtualised compute resource or a group of one or more OS containers is computed, and it is highlighted as a specific technical gap.

7.2.17.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.2.

7.2.17.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #17.1: A metric defining the energy efficiency of a virtualised compute resource.

Gap #17.2: A metric defining the energy efficiency of a group of one or more OS containers.

Gap #17.3: A metric defining the performance of a virtualised compute resource.

Gap #17.4: A metric defining the performance of a group of one or more OS containers.

7.2.18 Solution #18: Energy metrics related to virtualised resources and containerized workloads

7.2.18.1 Introduction

Key issue #4.2 and relevant use cases identify the need to define energy metrics for virtualised resources and containerized workloads. By defining relevant energy metrics, NFV-MANO and the service provider can clearly identify the monitoring parameters of the relevant NFV-MANO managed objects and use that information either to execute energy related policies or perform additional computations to determine the power consumption and other energy related metrics associated to other managed objects.

7.2.18.2 Solution description

The present solution considers the definition of the following energy metrics for virtualised compute resources and groups of one or more OS containers (hereafter referred as the monitored object):

- minimum power consumption, over a defined time interval;
- mean power consumption, over a defined time interval;
- maximum power consumption, over a defined time interval; and
- total power consumption, over a defined time interval.

For the computation of the power consumption, the present solution considers two sub-solutions:

- Solution #18.1: Estimated value based on the relative resource usage of the monitored object with respect to other monitored objects placed on the same NFVI compute resource.
- Solution #18.2: Measured value provided by an entity (e.g. an OS tool) which is capable of measuring the power consumption of the monitored object.

Solution #18.1:

An example of estimated value is provided in ETSI TS 128 554 [i.37]. Clause 6.7.3.1.4 of ETSI TS 128 554 [i.37] specifies the "estimated virtual compute resource instance energy consumption based on mean vCPU usage". In this case, the energy consumption of a virtual compute resource instance is estimated as a proportion of the energy consumption of the NFVI node on which the resource instance runs. The proportion is obtained by dividing the vCPU mean usage of the virtual compute resource instance by the sum of the vCPU mean usage of all virtual compute resource instances running on the same NFVI node.

This solution only produces the mean power consumption.

Solution #18.2:

In the case of a "real" measured value of the monitored object, this can be obtained by determining the power being used by the processes running on the compute resource. This is made possible by quantifying the total process time (i.e. the amount of CPU time used by a process, which in Linux® is measured in Jiffies [i.62]) and collecting information about the power usage by the compute resource. Current and typical compute servers provide sensors that measure how much power is being used by CPUs and GPUs, in watts. A power reading can then be derived by counting the process time of a process when it is running and comparing it to the power that is being consumed at the time.

An example of a solution using this concept is the Scaphandre metrology agent [i.63], which supports exposing power reading metrics to virtual machines and OS containers. A similar concept is the solution named Kepler [i.64] which focuses on measuring, in the case of Kubernetes®, the energy consumption by Pods, i.e. groups of one or more OS containers.

Based on these measurements, the following energy metrics can be defined as:

- Minimum power consumption: it equals to the smallest value from the time series of individual power consumption measurements of the monitored object over the defined time interval.
- Mean power consumption: it equals to the arithmetic mean of all individual power consumption measurements of the monitored object over the defined time interval.
- Maximum power consumption: it equals to the greatest value from the time series of individual power consumption measurements of the monitored object over the defined time interval.
- Total power consumption: it equals to the sum of all individual power consumption measurements of the monitored object over the defined time interval.

7.2.18.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.2.

7.2.18.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #18.1: A set of metrics related to energy/power consumption of a virtualised compute resource.

Gap #18.2: A set of metrics related to energy/power consumption of a group of one or more OS containers.

NOTE: As the solution focuses on compute type of virtualised resources or containerized workloads, no gaps are derived with respect to other kinds of resources, such as storage and network.

7.2.19 Solution #19: Energy efficiency of VNFs

7.2.19.1 Introduction

Key issue #4.3 and relevant use cases identify the need to define energy efficiency of a VNF. For instance, determining the energy efficiency of a VNF in a reliable, accurate and reproducible manner provides to NFV-MANO and the service provider valuable information to compare the efficiency between the VNFs, which can be a trigger for further developing energy efficiency and power saving policies for NFV-MANO.

7.2.19.2 Solution description

The present solution considers the definition of an energy efficiency metric of a VNF instance (Eff_{vnf}) using the following formula:

$$Eff_{vnf} = \frac{Perf_{vnf}}{Pow_{vnf}}$$

wherein $Perf_{vnf}$ corresponds to the service performance of the VNF instance and Pow_{vnf} is the power consumption associated to the VNF instance over a defined time interval.

NOTE: For the definition of power consumption metrics of a VNF instance refer to solution #20 in clause 7.2.20.

The determination of the service performance of a VNF will greatly depend on the functionality and purpose of the VNF, which from an NFV-MANO point of view is not specified. However, some other generic performance metrics can be considered, such as the number of incoming/outgoing bytes of all the VNF external CPs. In this case,

$$Perf_{vnf} = \sum_{i} ByteIncomingVnfExtCp_i + ByteOutgoingVnfExtCp_i$$

where i denotes the instance of a VNF external CP. A similar approach is also followed in the definition of NG-RAN data energy efficiency specified in clause 6.7.1 of ETSI TS 128 554 [i.37].

7.2.19.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.3.

7.2.19.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #19.1: A metric defining the energy efficiency of a VNF.

7.2.20 Solution #20: Energy metrics related to VNFs

7.2.20.1 Introduction

Key issue #4.3 and relevant use cases identify the need to define energy metrics associated to VNF instances. By defining relevant energy metrics, NFV-MANO and the service provider can clearly identify the monitoring parameters of the relevant NFV-MANO managed objects and use that information either to execute energy related policies or perform additional computations to determine the power consumption and other energy related metrics associated to other managed objects.

7.2.20.2 Solution description

The present solution considers the definition of the following energy metrics for VNF instances:

- Actual power consumption, at a defined time;

- Minimum power consumption, over a defined time interval;
- Mean power consumption, over a defined time interval;
- Maximum power consumption, over a defined time interval; and
- Total power consumption, over a defined time interval.

For the computation of the power consumption, the present solution considers two sub-solutions:

- Solution #20.1: Estimated value based on estimated power consumption of each individual VNFC instance comprising the VNF instance, whose power consumption is in turn derived from the estimated power consumption of the virtual compute resource or group of one or more OS containers realizing the VNFC.
- Solution #20.2: Computed value based on the measured value of power consumption of the virtual compute resources or group of one or more OS containers realizing the VNFCs that comprise the VNF instance.

Solution #20.1:

An example of estimated value is provided in ETSI TS 128 554 [i.37]. Clause 6.7.3.1.3 of ETSI TS 128 554 [i.37] specifies the "estimated VNFC energy consumption". In this case, the energy consumption of a VNFC instance is the same as the "estimated virtualised compute resource instance energy consumption based on mean vCPU" (see also solution #18.1 in clause 7.2.18).

This solution only produces the mean power consumption.

Solution #20.2:

In the case of a computed value based on the measured values of power consumption, solution #18.2 in clause 7.2.18 describes the approach for retrieving the measured values of power consumption of the virtualised compute resources and the groups of one or more OS containers (hereafter referred as monitored resource object). These resources realize the VNFC, and therefore, the power consumption of these resources represent the power consumption of the VNFC. Therefore, the various power consumption metrics can be defined as:

- Actual power consumption: it equals to the sum of measurements at a specific time of individual power consumption measurements of each monitored resource object realizing each individual VNFC instance:

$$ActualPowVnf_i = \sum_j PowMonResObj_{j,i}$$

  Where i represents the index in the time series and j represents the instance of monitored resource object realizing a VNFC instance, and $PowMonResObj_{j,i}$ represents the power consumption of the monitored resource object j at time i.

- Minimum power consumption: it equals to the smallest value of the time series resulting from the sum of measurements of individual power consumption measurements of each monitored resource object realizing each individual VNFC instance at a specific time:

$$MinPowVnf = min_i\left(\sum_j PowMonResObj_{j,i}\right)$$

- Mean power consumption: it equals to the arithmetic mean value of the time series resulting from the sum of measurements of individual power consumption measurements of each monitored resource object realizing each individual VNFC instance at a specific time:

$$MeanPowVnf = \frac{1}{n}\sum_i^n\left(\sum_j PowMonResObj_{j,i}\right)$$

- Maximum power consumption: it equals to the greatest value of the time series resulting from the sum of measurements of individual power consumption measurements of each monitored resource object realizing each individual VNFC instance at a specific time:

$$MaxPowVnf = max_i\left(\sum_j PowMonResObj_{j,i}\right)$$

- Total power consumption: it equals to the sum of all values of the time series resulting from the sum of measurements of individual power consumption measurements of each monitored resource object realizing each individual VNFC instance at a specific time:

$$TotalPowVnf = \sum_i^n\left(\sum_j PowMonResObj_{j,i}\right)$$

NOTE: As the solution leverages the metrics defined in Solution #18.2 of virtualised compute resources or containerized workloads, metrics of constituents of the VNF other than VNFC instances are not analyzed by the present solution.

7.2.20.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.3.

7.2.20.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #20.1: A set of metrics related to energy/power consumption of a VNF instance.

7.2.21 Solution #21: Energy efficiency of NSs

7.2.21.1 Introduction

Key issue #4.4 and relevant use cases identify the need to define energy efficiency of an NS. For instance, determining the energy efficiency of an NS in a reliable, accurate and reproducible manner provides to NFV-MANO and the service provider valuable information to compare the efficiency between the NS, which can be a trigger for further developing energy efficiency and power saving policies for NFV-MANO.

7.2.21.2 Solution description

The present solution considers the definition of an energy efficiency metric of an NS instance (Eff_{ns}) using the following formula:

$$Eff_{ns} = \frac{Perf_{ns}}{Pow_{ns}}$$

wherein $Perf_{ns}$ corresponds to the service performance of the NS instance and Pow_{ns} is the power consumption associated to the NS instance over a defined time interval.

NOTE: For the definition of power consumption metrics of an NS instance refer to solution #22 in clause 7.2.22.

The determination of the service performance of an NS will greatly depend on the functionality and purpose of the NS instance and the individual VNF instances that comprise it, which from an NFV-MANO point of view is not specified. However, some other generic performance metrics can be considered, such as the number of incoming/outgoing bytes of all the VNF external CPs of all VNF instances which are part of the NS instance. In this case,

$$Perf_{ns} = \sum_{j} Perf_{vnf,j} = \sum_{j} \left(\sum_{i} ByteIncomingVnfExtCp_{i,j} + ByteOutgoingVnfExtCp_{i,j} \right)$$

where i denotes the instance of VNF external CP and j denotes the instance of a VNF. A similar approach is also followed in the definition of energy efficiency of network slices in clause 6.7.2 of ETSI TS 128 554 [i.37].

7.2.21.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.4.

7.2.21.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #21.1: A metric defining the energy efficiency of an NS instance.

7.2.22 Solution #22: Energy metrics related to NSs

7.2.22.1 Introduction

Key issue #4.4 and relevant use cases identify the need to define energy metrics associated to NS instances. By defining relevant energy metrics, NFV-MANO and the service provider can clearly identify the monitoring parameters of the relevant NFV-MANO managed objects and use that information either to execute energy related policies or perform additional computations to determine the power consumption and other energy related metrics associated to other managed objects.

7.2.22.2 Solution description

The present solution considers the definition of the following energy metrics for NS instances:

- Actual power consumption, at a defined time;
- Minimum power consumption, over a defined time interval;
- Mean power consumption, over a defined time interval;
- Maximum power consumption, over a defined time interval; and
- Total power consumption, over a defined time interval.

For the computation of the power consumption, the present solution considers two sub-solutions:

- Solution #22.1: Estimated value based on estimated power consumption of each individual VNF instance comprising the NS instance, whose power consumption is in turn derived from the estimated power consumption of the virtual compute resources or groups of one or more OS containers realizing the VNF.
- Solution #22.2: Computed value based on the measured value of power consumption of the virtual compute resources or groups of one or more OS containers realizing the VNFC instances that comprise the VNF instance, that in turn comprise the NS instance.

NOTE: As the solutions leverage the metrics defined in Solutions #20.1 and #20.2, which are based on virtualised compute resources or containerized workloads, metrics of constituents of the VNF other than VNFC instances are not considered by the present solution.

Solution #22.1:

An example of estimated value is provided in ETSI TS 128 554 [i.37], but applied to a 3GPP network slice, to the radio access network, or to the 5G core network. In such a case, the value is estimated, because for the computation of the power consumption of the VNF instances that comprise the various parts of the network, the estimated value is reused, as described in solution #20.1 in clause 7.2.20.2. Clause 6.7.3.3 of ETSI TS 128 554 [i.37] defines the network slice energy consumption as the summation of the individual energy consumptions of NFs comprising the network slice, while clause 6.7.3.2 of ETSI TS 128 554 [i.37] does the same for the 5G network core.

Solution #22.2:

In the case of a computed value based on the measured values of power consumption, solution #20.2 in clause 7.2.20 describes the approach for computing the measured values of power consumption of a VNF instance. The energy metrics are calculated the same way as defined in solution #20.2, but substituting the $PowMonResObj_{j,i}$ by $ActualPowVnf_{j,i}$.

7.2.22.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.4.

7.2.22.4 Gap analysis

The ETSI NFV specifications referenced in the present document do not specify:

Gap #22.1: A set of metrics related to energy/power consumption of an NS instance.

7.2.23 Solution #23: PIM for acquiring power consumption of physical resources in the NFVI

7.2.23.1 Introduction

Key issue #4.1 and relevant use cases highlight the need of monitoring NFVI resources in terms of power consumption and energy efficiency. In particular, it becomes relevant to determine how NFV-MANO and the service provider can acquire power consumption and energy efficiency related measurements about the physical resources in the NFVI, and which interfaces are involved in such acquisition.

PIM, as a specialized function responsible for the management of physical resources, can be utilized to expose power consumption and energy related measurements to other NFV-MANO components or the service provider. Additionally, NFV-MANO functional blocks and functions can rely on power consumption information and relevant energy-related measurements of NFVI physical resources to process and provide similar information and measurements about their respective managed objects, such as virtualised resources, containerized workloads, VNF instances and NS instances.

7.2.23.2 Solution description

The present solution considers the use of PIM for processing and providing power consumption information and relevant energy-related measurements about the NFVI physical resources. This information can be utilized by other NFV-MANO entities, e.g. VIM, VNFM, etc., as well as new functional elements, e.g. NFV power metrics function, described in clause 7.2.15, to associate and compute the power consumption and energy metrics of their respective NFV-MANO managed objects.

Solution #1, described in clause 7.2.1, illustrates how PIM can be used to acquire power consumption metrics by interacting with BMCs of relevant NFVI resource(s). Additionally, Solution #1 proposes an inventory of all the physical resources managed by PIM, where PIM can keep the real-time power consumption information and relevant energy-related measurements associated with different physical resources in the NFVI.

The present solution assumes that PIM is responsible for managing different types of physical resources in the NFVI. Therefore, it can associate the power consumption and energy metrics (acquired from respective BMCs) to the corresponding managed objects.

NOTE 1: In case of PIM-managed objects, power consumption and energy related metrics are obtained from their respective BMCs. At the time of developing the present document, the specification of the exact types of managed objects and their hierarchy/level of granularity is still under development in ETSI GS NFV-IFA 053 [i.65]. Some computation or aggregation of energy related measurements can take place inside PIM for associating accurate power consumption against each managed object. This computation is outside the scope of the present solution.

NOTE 2: The present solution does not define how the NFV-MANO entities compute the power consumption of their corresponding managed objects, which is defined by other solutions in the present document.

With regards to the exposure of the power consumption and energy metrics by PIM, the present solution can be divided into the following two aspects:

- Aspect #23.1: PIM exposing a Physical Resource Performance Management interface for providing power consumption and energy metrics.
- Aspect #23.2: Definition of performance metrics related to power consumption and energy information associated to various PIM-managed objects such as physical compute, network and storage resources in the NFVI.

Figure 7.2.23.2-1 illustrates a high-level representation of the solution depicting the PIM exposing relevant metrics to its consumers such as NFVO, VNFM, VIM, CCM, and NFV power metrics function (introduced in Solution #15 described in clause 7.2.15).

NOTE 3: Figure 7.2.23.2-1 does not represent an accurate NFV architectural framework and focuses only on the points of interaction of the PIM as main actor in the present solution.

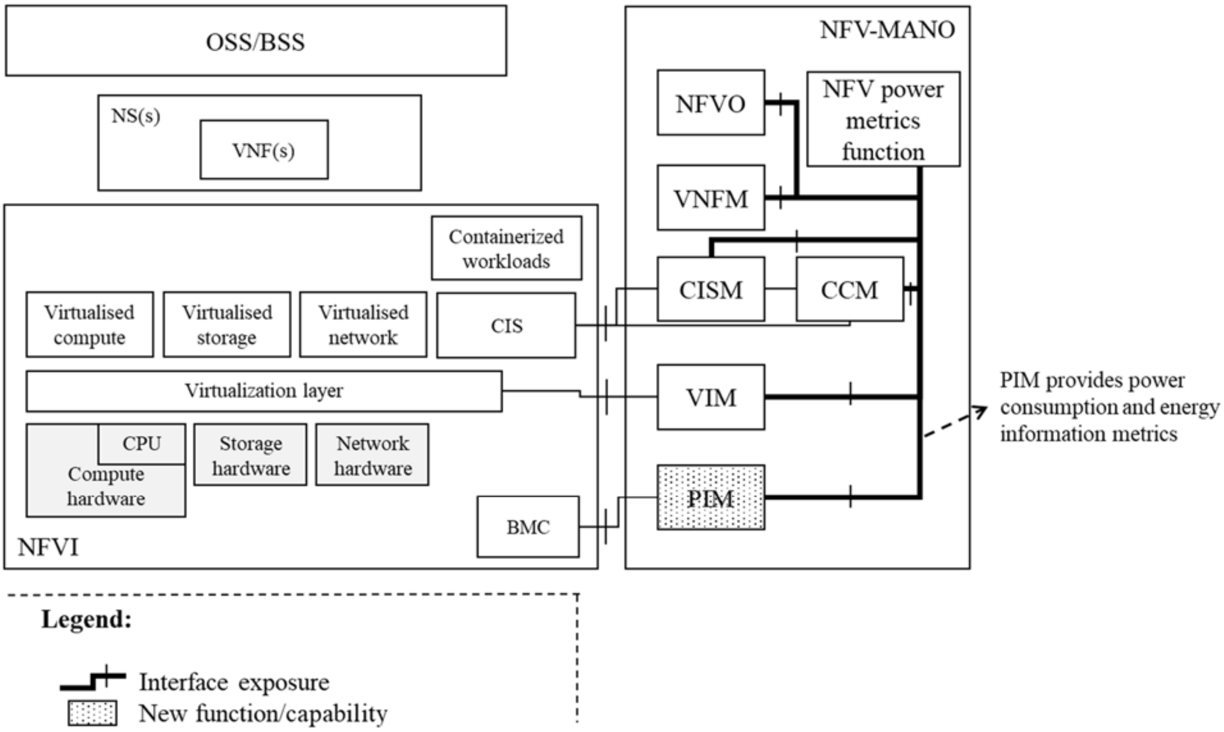

Figure 7.2.23.2-1: High-level representation of Solution #23

7.2.23.3 Key issues addressed

The present solution aims at addressing aspects of the following key issue described in clause 6:

- Key issue #4.1.

7.2.23.4 Gap analysis

The referred ETSI NFV specifications in the present document do not specify:

Gap #23.1: PIM to support providing power consumption and energy metrics associated to its managed objects via a standard interface.

Gap #23.2: Definition of performance measurements related to power consumption and energy information of physical resources.

7.3 Evaluation of solutions

7.3.1 Overview

Clause 7.3 provides an evaluation of the solutions documented in clause 7.2.

Some solutions cover aspects associated to more than one key issue. The evaluation of the solution is performed within the context of each associated key issue.

Where available, if more than one solution target the same key issue and represent alternatives to resolve the key issue, the evaluation reflects such relationships. Furthermore, if a solution contains sub-solutions, the analysis considers a comparison among the sub-solutions.

7.3.2 Evaluation of solutions for key issues on energy consumption

Table 7.3.2-1 provides an analysis of pros and cons of the different solutions documented on key issues related to energy consumption, as introduced in clause 6.2.

Table 7.3.2-1: Pros/cons analysis of solutions for key issues on energy consumption

Solution	Pros	Cons	Alternatives
Solution #1	• Information and inventory of power consuming components are made available by an identified component in the framework. • Existing NFV-MANO functional blocks and functions defined by the referenced ETSI NFV specifications are not impacted for managing information of power consuming components in the NFVI.	• Correlation of power consumption between power consuming components and NFV-MANO managed objects (e.g. a virtualised resource, a VNF, etc.) is not made available by the solution. See note.	Alternative to: • None
Solution #2	• Existing NFV-MANO functional blocks and functions (VIM, CISM and CCM) are leveraged for managing the power states for resources that are associated to their managed objects. • Power state management of "physical" power consuming components is performed by a function (PIM) without impacting existing NFV-MANO functional blocks and functions.	• Correlation of power consumption with power state changes is not made available to the NFV-MANO functional blocks and functions responsible for virtualised and containerized resources management (VIM, CCM and CISM). • Power state management performed via existing NFV-MANO functional blocks and functions (VIM, CISM and CCM) has an impact only on parts of the power consuming components, and not the full set. See note.	Alternative to: • None
NOTE: This con can be minimized by combining this solution with Solution #23. Evaluation of Solution #23 is provided in clause 7.3.5 of the present document.			

In addition, the following specific compatibility aspects are derived from the analysis:

- Solutions #1 and #2: Solution #1 provides a mechanism for inventorying which elements are power consuming components and which ones can be power managed, while Solution #2 provides a mechanism to perform the power state management on the power consuming elements that are inventoried.
- Solutions #1 and #23: Solution #1 provides the mechanism for PIM to collect measurements about energy associated to NFVI resources, while mechanism provided in Solution #23 can be used to process and provide power consumption information of physical components to NFV-MANO functional blocks/functions, which can then correlate and associate power consumption of physical components to their respective NFV-MANO managed objects.
- Solutions #2 and #23: Solution #2 provides a mechanism to perform the power state management on the power consuming elements in the NFVI, while Solution #23 provides information about changes in power consumption of NFVI resources based on the power state changes to the relevant NFV-MANO functional blocks and functions.

7.3.3 Evaluation of solutions for key issues on energy efficiency management

Table 7.3.3-1 provides an analysis of pros and cons of the different solutions documented on key issues related to energy efficiency management, as introduced in clause 6.3.

Table 7.3.3-1: Pros/cons analysis of solutions for key issues on energy efficiency management

Solution	Pros	Cons	Alternatives
Solution #3	• Full reuse of existing NFV-MANO capabilities regarding virtualised resource reservation, with only minor additional enhancements needed. • Additional information provided regarding the usage of virtualised resources (which ones will be terminated) can make the resource reservation process more efficient.	• More coupling between "virtualised resources management" and "virtualised resources reservation management".	Alternative to: • None
Solution #4.1	• Specific configuration and state data can be targeted for backup. • Greater power consumption savings can be achieved, since potentially all VNFCs or needed components of the VNF can be safely terminated.	• Storage of the configuration and state data incurs also additional power consumption. • Undocumented procedures of restoring configuration and state data using VNF generic OAM functions.	Alternative to: • Solution #4.2; and • Solution #4.3.
Solution #4.2	• Full reuse of existing NFV-MANO capabilities regarding VNF snapshotting, without additional enhancements identified to be performed. • Greater power consumption savings can be achieved, since potentially all VNFCs or needed components of the VNF can be safely terminated.	• VNF snapshotting makes use of additional resources (storage resources) which also incur additional power consumption. • VNF snapshotting is not able to differentiate which data to backup, only the whole VNF or individual VNFC instances are snapshotted. • VNF snapshot restoration can be a lengthy procedure.	Alternative to: • Solution #4.1; and • Solution #4.3.

Solution	Pros	Cons	Alternatives
Solution #4.3	• No impact on the configuration and state data (i.e. no backup procedure is needed). • No restoration of configuration and state data is needed.	• Power consumption savings can only be achieved partially, only on the subset of VNFC instances that can be terminated. • VNF LCM procedures can be impacted if there is no well-defined correlation between VNF deployment flavours and the VNFC instances that need to be preserved, i.e. inconsistent runtime information between what is deployed via VNF deployment flavours and what VNFCs are really instantiated and preserved during the power management procedure.	Alternative to: • Solution #4.1; and • Solution #4.2.
Solution #7	• Avoids conflicts among deployed components on the NFVI due to dynamic power state management. • Power state ranges enable a wider spectrum of possibilities to deploy the VNF or its components in dynamic zones.	• Fully optimal power consumption might not be achieved due to the use of static zones. • Resource sharing is impacted, as NFVI resources qualify either as static or dynamic in terms of power management, and not all available resources in the NFVI can be considered for deployment. For optimal resource utilization, an infrastructure management entity such as VIM or PIM is expected to manage these zones, i.e. to assign or de-assign NFVI resources to a particular zone, considering the runtime demand.	Alternative to: • None
Solution #9	• Enables defining at runtime the operation and behaviour of the NFV-MANO regarding its managed objects considering power consumption/management events, conditions and actions.	• Additional variables (actions derived due to events related to power consumption/management) need to be considered, which can impact the policy conflict resolution.	Alternative to: • None
Solution #11	• None identified.	• None identified.	Alternative to: • None

In addition, the following specific compatibility aspects are derived from the analysis:

- Solutions #3 and #4: Solution #3 provides the mechanism to ensure the availability of resources during the restoration process, while Solution #4 provides a mechanism to backup configuration and state data for resources that will be terminated but need to be restored later using reserved resources.
- Solutions #7, #9, #11: all solutions are compatible with any other solution relevant to key issues on energy efficiency management.
- Solution #9 and #11: Solution #11 describes the capabilities of NFV-MANO that can be used as actions in energy related policies foreseen by Solution #9.

7.3.4 Evaluation of solutions for key issues on energy efficient design

Table 7.3.4-1 provides an analysis of pros and cons of the different solutions documented on key issues related to energy efficient design, as introduced in clause 6.4.

Table 7.3.4-1: Pros/cons analysis of solutions for key issues on energy efficient design

Solution	Pros	Cons	Alternatives
Solution #1	• Information and inventory of power consuming components are made available to NFV-MANO and the network operator to contrast power and energy related information about the composition of the NFVI. • Information and inventory of information regarding the power state management capability of the NFVI provide additional capabilities for tailoring more specific NFVI configuration.	• None identified.	Alternative to: • None
Solution #2	• Enables implementing actions for energy efficient policies regarding power state management of physical and virtual resources in the NFVI.	• None identified.	Alternative to: • None
Solution #5	• VNF power profiles provide necessary information for a network operator and NFV-MANO to correlate the VNF composition and structure with power consumption and power state management.	• VNF power profiles might only be able to provide an estimation about the power consumption of the VNF, in particular if the profiles are defined based on deployment test cases and NFVI which is not exactly the same as the one where the VNF is to be deployed.	Alternative to: • None
Solution #6.1	• Support for VNF power management, which is not provided in the existing NFV-MANO architectural framework. • As part of the VNF generic OAM framework, the VNF power management function is "closer" to the VNF and other dedicated VNF generic OAM functions to correlate necessary information regarding the VNF's application. • No impact, neither on existing VNFM capabilities, nor on other VNF generic OAM functions. • Possibility to provide generic VNF power management capabilities from a PaaS platform perspective, while VNF generic OAM functions are considered PaaS services.	• As the VNF generic OAM framework does not directly interact with VIM/CISM for virtualised resources/containerized workloads management, there is dependency with the VNFM to further perform power state management actions on the relevant resources. This con might be minimized by enabling VNF generic OAM functions, as PaaS services, to directly trigger actions towards the PaaS platform. • As the VNF generic OAM framework is not a direct consumer of the VNFD, transferring to the VNF power management function the relevant VNF power profiles information is necessary.	Alternative to: • Solution #6.2.

Solution	Pros	Cons	Alternatives
Solution #6.2	• Support for VNF power management, which is not provided in the existing NFV-MANO architectural framework. • As part of the VNFM, the VNF power management function is "closer" to the functionalities to manage the virtualised resources and containerized workloads, which are the ones that eventually incur in actual power consumption associated to the VNF and on which lower-level power state management actions can be performed.	• The VNFM will have to interact with the VNF generic OAM functions and the VNF instance to gather additional telemetry to correlate necessary information and determine the power management actions.	Alternative to: • Solution #6.1.
Solution #7	• Clear delineation of characteristics that determines the pooling based on power management, i.e. either part of static or dynamic zones.	• Static/dynamic zoning adds a new variable of NFVI disaggregation and partitioning that can impact resource sharing.	Alternative to: • None
Solution #8	• Enables allocating resources for VNFC/VNF considering power state constraints.	• More complex virtualised resource/containerized workload placement to be performed by VIM and CISM, as additional constraints are to be considered for scheduling the allocation of resources to virtualised resources and containerized workloads of VNFC/VNF.	Alternative to: • None
Solution #10	• More granular state management of the NFV-MANO services, from a producer perspective, enables to selectively decrease the power consumption of the NFV-MANO functional entity.	• None identified.	Alternative to: • None

In addition, the following specific compatibility aspects are derived from the analysis:

- Solutions #5 and #6: Solution #5 provides the power profiles to be used for VNF power management, while Solution #6 provides the mechanism to use the power profiles.
- Solutions #1, #2, #7 and #8: Solution #1 provides mechanisms for inventorying the NFVI and determining its capabilities regarding power consumption; Solution #2 provides the mechanisms to realize power state management; Solution #7 provides the design principle and mechanism to configure the NFVI with power management related zoning principles (static and dynamic zones); and Solution #8 provides the mechanism to set the power state for resources that can be deployed either in static or dynamic zones.
- Solutions #7 and #8: Solution #8 proposes power state aware allocation of compute resources for VNF/VNFCs. The placement of virtualised resources or containerized workloads on the appropriate compute nodes according to the desired CPU power profiles can be simplified by leveraging power state-based zones as proposed by Solution #7.

7.3.5 Evaluation of solutions for key issues on energy related metrics and KPIs

Table 7.3.5-1 provides an analysis of pros and cons of the different solutions documented on key issues related to energy metrics and KPIs, as introduced in clause 6.5.

Table 7.3.5-1: Pros/cons analysis of solutions for key issues on energy related metrics and KPIs

Solution	Pros	Cons	Alternatives
Solution #1	• Energy metrics provide means to determine the power consumption of NFVI resources.	• None identified.	Alternative to: • None
Solution #6	• Correlation between the VNF, as a managed object, and its energy metrics done by a single functional block/function, i.e. the VNF power management function. • Existing PM or metrics interfaces exposed by the functional block/function can be reused without any additional impact.	• In the case of VNF power management function realized by VNFM, the VNFM capabilities to be extended to be able to compute respective energy metrics.	Alternative to: • Solution #15.
Solution #12	• Correlation between the virtualised resource, as a managed object, and its energy metrics done by a single functional block/function, i.e. the VIM. • Existing PM interfaces exposed by the functional block/function can be reused without any additional impact.	• VIM capabilities to be extended to be able to compute respective energy metrics.	Alternative to: • Solution #15.
Solution #13	• Correlation between the containerized workload, as a managed object, and its energy metrics done by a single functional block/function, i.e. the CISM. • Reuse of PM interfaces exposed by the functional block/function. See note 1.	• CISM capabilities to be extended to be able to compute respective energy metrics.	Alternative to: • Solution #15.
Solution #14	• Correlation between the NS, as a managed object, and its energy metrics done by a single functional block/function, i.e. the NFVO. • Existing PM interfaces exposed by the functional block/function can be reused without any additional impact.	• NFVO capabilities to be extended to be able to compute respective energy metrics.	Alternative to: • Solution #15.
Solution #15	• Existing NFV-MANO functional blocks/functions are not impacted, no additional capabilities needed regarding the computation of respective energy metrics. • Many possibilities of correlating power consumption and energy related metrics due to the access to all relevant state information and power consumption information, e.g. the function can instantly use compute VNF power consumption to derive NS related power consumption.	• It is expected that huge amounts of data would need to be correlated and processed by the NFV power metrics function. • Lots of state information about NFV-MANO managed objects would need to be collected and duplicated by the NFV power metrics functions for correlation with the power consumption information collected from the infrastructure.	Alternative to: • Solution #6, #12, #13, and #14.
Solution #16	• Energy efficiency metrics provide a means to compare the efficiency of NFVI resources.	• None identified.	Alternative to: • None

Solution	Pros	Cons	Alternatives
Solution #17	• Energy efficiency metrics provides a means to compare the efficiency of virtualised resources and containerized workloads.	• None identified.	Alternative to: • None
Solution #18.1	• Energy metrics provide means to determine the power consumption of virtualised resources and containerized workloads.	• It only provides mean power consumption. • The value that is provided is an estimated value and not an actual measured one.	Alternative to: • Solution #18.2.
Solution #18.2	• Energy metrics provide means to determine the power consumption of virtualised resources and containerized workloads. • A more accurate and actual measured value for the energy metrics is provided. • Various energy metrics can be defined.	• The tools used to collect actual energy measurements from the system can incur in additional power consumption.	Alternative to: • Solution #18.1.
Solution #19	• Energy efficiency metrics provide a means to compare the efficiency of VNFs.	• None identified.	Alternative to: • None
Solution #20.1	• Energy metrics provide means to determine the power consumption of VNFs.	• It only provides mean power consumption. • The value that is provided is an estimated value and not an actual measured one.	Alternative to: • Solution #20.2.
Solution #20.2	• Energy metrics provide means to determine the power consumption of VNFs. • A more accurate and actual measured value for the energy metrics is provided. • Various energy metrics can be defined.	• None identified.	Alternative to: • Solution #20.1.
Solution #21	• Energy efficiency metrics provide a means to compare the efficiency of NSs.	• None identified.	Alternative to: • None
Solution #22.1	• Energy metrics provide means to determine the power consumption of NSs.	• The value that is provided is an estimated value and not an actual measured one.	Alternative to: • Solution #22.2.
Solution #22.2	• Energy metrics provide means to determine the power consumption of NSs. • A more accurate and actual measured value for the energy metrics is provided. • Various energy metrics can be defined.	• None identified.	Alternative to: • Solution #22.1.
Solution #23	• Correlation between the physical resource, as a managed object, and its energy metrics done by a single functional block/function, i.e. the PIM. • Exposure of energy related measurements of physical resources by PIM to other NFV-MANO components, which can then correlate and associate power consumption of physical components to their respective NFV-MANO managed objects.	• PIM capabilities to be defined for processing and computing relevant energy metrics. See note 2.	Alternative to: • None

Solution	Pros	Cons	Alternatives
NOTE 1: The specification of PM service interfaces produced by the CISM cannot be referenced by the present document version due to the relevant specification work being in progress. NOTE 2: At the time of developing the present document, the specification of PIM, its requirements and produced service interfaces is currently in progress in ETSI GS NFV-IFA 053 [i.65].			

In addition, the following specific compatibility aspects are derived from the analysis:

- Solutions #6, #19 and #20: Solution #6 provides the mechanism to collect measurements about energy associated to VNFs, while Solutions #19 and #20 define the relevant energy related metrics.
- Solutions #1 and #16: Solution #1 provides the mechanism to collect measurements about energy associated to NFVI resources as well as specified metrics related to power consumption (among others), while Solution #16 provides the mechanism to define the energy efficiency of NFVI resources.
- Solutions #12, #17 and #18: Solution #12 provides the mechanism to collect measurements about energy associated to virtualised resources, while Solutions #17 and #18 define the relevant energy related metrics.
- Solutions #13, #17 and #18: Solution #13 provides the mechanism to collect measurements about energy associated to containerized workloads, while Solutions #17 and #18 define the relevant energy related metrics.
- Solutions #14, #21 and #22: Solution #14 provides the mechanism to collect measurements about energy associated to NSs, while Solutions #21 and #22 define the relevant energy related metrics.
- Solutions #1 and #23: Solution #1 provides the mechanism for PIM to collect measurements about energy associated to NFVI resources, while Solution #23 enables PIM to process and provide power consumption information and relevant energy-related measurements of NFVI resources to other NFV-MANO components.

8 Recommendations

8.1 Overview

The present clause 8 documents recommendations about potential enhancements, changes, or clarifications to existing ETSI NFV specifications. The recommendations are derived based on the gap analysis performed in the documentation of the potential solutions and the evaluation of solutions documented in clause 7.

The recommendations are categorized and elaborated as follows:

- architecture and framework aspects (refer to clause 8.2);
- functional aspects (refer to clause 8.3);
- descriptors and other information/data model artefacts (refer to clause 8.4);
- interfaces and associated information/data model (refer to clause 8.5); and
- other recommendations, if any (refer to clause 8.6).

Finally, clause 8.7 describes gaps identified in the present document for which no recommendations are further derived.

8.2 Recommendations related to the NFV architectural framework

The present clause provides recommendations focusing on enhancements to the NFV architectural framework, identifying potential new functions or functional blocks, and interactions among functional blocks.

Table 8.2-1 provides the recommendations related to the NFV architectural framework.

Table 8.2-1: Recommendations related to the NFV architectural framework

Identifier	Recommendation description	Comments and/or traceability
greennfv.arch.001	It is recommended that a requirement be specified for the NFV architectural framework to support an inventory of physical resources in the NFVI.	Refer to gap #1.1. This recommendation is related to the specification of physical infrastructure management function in the NFV-MANO.
greennfv.arch.002	It is recommended that a requirement be specified for the NFV architectural framework to support the discovery, collection and use of information about power consumption characteristics of physical resources in the NFVI.	Refer to gaps #1.2 and #23.1. This recommendation is related to the specification of physical infrastructure management function in the NFV-MANO.
greennfv.arch.003	It is recommended that a requirement be specified for the NFV architectural framework to support the discovery, collection and use of information about power state aware zoning and pooling of physical resources in the NFVI.	Refer to gaps #7.1 and #7.2.
greennfv.arch.004	It is recommended that a requirement be specified for the NFV architectural framework to support the collection of measurements related to energy metrics associated to all the NFV-MANO managed objects.	Refer to gap #9.1.
greennfv.arch.005	It is recommended that a requirement be specified for the NFV architectural framework to support the power management of VNF instances.	Refer to gaps #6.1 and #6.2. This recommendation can be proceeded with the specification of a "VNF power management function" as also described in recommendations greennfv.func.004 and greennfv.func.011.

8.3 Recommendations related to functional aspects

The present clause provides recommendations focusing on functional aspects of the functional blocks of the NFV architectural framework, identifying specific new or extended functionality of the NFV architectural framework functional blocks and functions.

Table 8.3-1 provides the recommendations related to functional aspects.

Table 8.3-1: Recommendations related to functional aspects

Identifier	Recommendation description	Comments and/or traceability
greennfv.func.001	It is recommended that a requirement be specified for the VIM to support managing the power states of the managed virtualised compute resources.	Refer to gaps #2.1 and #7.3
greennfv.func.002	It is recommended that a requirement be specified for the CCM to support managing the power states of CIS cluster nodes.	Refer to gaps #2.1 and #7.3.
greennfv.func.003	It is recommended that a requirement be specified for the PIM to support managing the power states of the physical resources in the NFVI that are subject to power consumption.	Refer to gaps #2.1 and #7.3.

Identifier	Recommendation description	Comments and/or traceability
greennfv.func.004	It is recommended that a requirement be specified for VNF power management function to support managing the power states of a VNF.	Refer to gaps #6.1 and #6.2. The recommendation does not take any preferred sub-solution, as documented in Solution #6 in clause 7.2.6. Further evaluation can be considered in derived normative work.
greennfv.func.005	It is recommended that a requirement be specified for the VIM to support scheduling and instantiating virtualised resources considering power state requirements.	Derived from gap #8.1.
greennfv.func.006	It is recommended that a requirement be specified for the CISM to support scheduling and instantiating containerized workloads considering power state requirements.	Derived from gap #8.1.
greennfv.func.007	It is recommended that a requirement be specified for the CCM to support additional tagging scopes for CIS cluster nodes related to CPU power profiles.	Refer to gap #8.4.
greennfv.func.008	It is recommended that a requirement be specified for the PIM to support providing power consumption and energy metrics associated to its managed objects.	Refer to gap #23.1 A standard interface (e.g. Physical Resources Performance Management interface) can be specified for providing the relevant metrics.
greennfv.func.009	It is recommended that a requirement be specified for the VIM to support providing power consumption and energy metrics associated to its managed objects.	Refer to gap #12.1. The Virtualised Resources Performance Management interface produced by the VIM can be leveraged for providing the relevant metrics.
greennfv.func.010	It is recommended that a requirement be specified for the CISM to support providing power consumption and energy metrics associated to its managed objects.	Refer to gaps #13.1 and #13.2.
greennfv.func.011	It is recommended that a requirement be specified for the VNF power management function to support providing power consumption and energy metrics associated to VNFs.	Refer to gaps #6.3 and #6.4. The recommendation does not take any preferred sub-solution, as documented in Solution #6 in clause 7.2.6. Further evaluation can be considered in derived normative work.
greennfv.func.012	It is recommended that a requirement be specified for the NFVO to support providing power consumption and energy metrics associated to its managed objects.	Refer to gap #14.1. The NS Performance Management interface produced by the NFVO can be leveraged for providing the relevant metrics.

8.4 Recommendations related to NFV descriptors and other artifacts

The present clause provides recommendations focusing on NFV descriptors, packaging and other artefacts.

Table 8.4-1 provides the recommendations related to NFV descriptors, packaging and other artefacts.

Table 8.4-1: Recommendations related to NFV descriptors and other artefacts

Identifier	Recommendation description	Comments and/or traceability
greennfv.desc.001	It is recommended that a requirement be specified for the VNFD to support the description of power profiles of a VNF.	Refer to gap #5.1.
greennfv.desc.002	It is recommended that a requirement be specified for the VNFD to support the description of CPU power state related information and constraints on a per VDU basis.	Refer to gaps #7.4 and #8.3.

8.5 Recommendations related to interfaces and information model

The present clause provides recommendations focusing on interfaces and associated information.

Table 8.5-1 provides the recommendations related to interfaces and associated information.

Table 8.5-1: Recommendations related to interfaces and information model

Identifier	Recommendation description	Comments and/or traceability
greennfv.if.001	It is recommended that a requirement be specified for the virtualised resource reservation interface(s) produced by the VIM to support providing information about which virtualised resources are expected to be terminated.	Refer to gap #3.1. Such information can be used by the VIM to perform a more efficient virtualised resource reservation by knowing what capacity will be made available by virtualised resources that are expected to be terminated.
greennfv.if.002	It is recommended that a requirement be specified for the virtualised resource reservation interface(s) produced by the VIM to support providing references to existing allocated virtualised resources during virtualised resource reservation creation procedures.	Refer to gap #3.2. Such information can be used by the VIM to identify the characteristics of which virtualised resources to reserve without having to provide the characteristics again, as these can be determined by reference.
greennfv.if.003	It is recommended that a requirement be specified for the virtualised compute resource management interface produced by the VIM to support providing power state requirements for the virtualised resources to be allocated.	Refer to gap #8.1.
greennfv.if.004	It is recommended that a requirement be specified for the VNF and NS LCM interfaces produced by the VNFM and NFVO, respectively, to enable the parameterization of power state related requirements associated to VDUs to be deployed.	Refer to gap #8.2. This recommendation is related to the "VNF flexible" feature specification, which enables means to parameterize VDU values during LCM operations.
greennfv.if.005	It is recommended that a requirement be specified for the "NFV-MANO configuration and information management interface" produced by an NFV-MANO functional entity to support suspending a managed entity such as the NFV-MANO functional entity application and NFV-MANO service interfaces.	Refer to gap #10.1.
greennfv.if.006	It is recommended that a requirement be specified for the "NFV-MANO configuration and information management interface" produced by an NFV-MANO functional entity to support suspending the consumption of NFV-MANO service interfaces produced by a peer NFV-MANO functional entity.	Refer to gap #10.2.

Identifier	Recommendation description	Comments and/or traceability
greennfv.if.007	It is recommended that a requirement be specified to define a metric about energy efficiency of NFVI physical resources, including compute, storage and network.	Refer to gap #16.1. The metric can be defined by referencing relevant standards.
greennfv.if.008	It is recommended that a requirement be specified to define a metric about energy efficiency of a virtualised compute resource.	Refer to gap #17.1
greennfv.if.009	It is recommended that a requirement be specified to define a metric about energy efficiency of a group of one or more OS containers.	Refer to gap #17.2
greennfv.if.010	It is recommended that a requirement be specified to define a metric about performance of a virtualised compute resource.	Refer to gap #17.3
greennfv.if.011	It is recommended that a requirement be specified to define a metric about performance of a group of one or more OS containers.	Refer to gap #17.4
greennfv.if.012	It is recommended that a requirement be specified to define a set of metrics related to energy/power consumption of NFVI physical resources.	Refer to gap #23.2.
greennfv.if.013	It is recommended that a requirement be specified to define a set of metrics related to energy/power consumption of a virtualised compute resource.	Refer to gap #18.1
greennfv.if.014	It is recommended that a requirement be specified to define a set of metrics related to energy/power consumption of a group of one or more OS containers.	Refer to gap #18.2
greennfv.if.015	It is recommended that a requirement be specified to define a metric about energy efficiency of a VNF.	Refer to gap #19.1
greennfv.if.016	It is recommended that a requirement be specified to define a set of metrics related to energy/power consumption of a VNF.	Refer to gap #20.1
greennfv.if.017	It is recommended that a requirement be specified to define a metric about energy efficiency of an NS.	Refer to gap #21.1
greennfv.if.018	It is recommended that a requirement be specified to define a set of metrics related to energy/power consumption of an NS.	Refer to gap #22.1

8.6 Other recommendations

The present clause provides recommendations for categories other than those documented in the previous clauses.

Table 8.6-1 provides other recommendations.

Table 8.6-1: Other recommendations

Identifier	Recommendation description	Comments and/or traceability
N/A	N/A	N/A

8.7 Gaps without recommendations

The present clause lists the gaps identified in potential solutions for which recommendations are not derived.

NOTE 1: The purpose of the present clause is to track gaps from which normative work is not foreseen.

- Gap #4.1 in clause 7.2.4.4: Solution #4.3 introduced in clause 7.2.4.2 is not recommended. See note 2.
- Gaps #15.1 and #15.2 in clause 7.2.15.4: NFV power metrics function is not recommended.

NOTE 2: The referenced ETSI NFV specifications do not specify the behaviour and composition of the VNF applications, to which this solution has dependency.

9 Conclusion

The present document studies the aspects of NFV that have an impact on power consumption and energy efficiency. Use cases are developed regarding different areas of concern from which key issues are subsequently identified. Based on the list of key issues, potential solutions on how to leverage NFV capabilities or extend them are described and potential technical gaps of NFV are identified. Recommendations for additional normative work relevant to the scope of the present document are finally derived.

The set of recommendations indicate the need to perform additional normative specification work to enhance the capabilities of the NFV architectural framework to better support energy efficiency aspects. The areas for which additional normative work are identified are:

- recommendations related to NFV architectural aspects (refer to clause 8.2);
- recommendations related to functional aspects of the NFV architectural framework and its functional blocks (refer to clause 8.3);
- recommendations related to NFV descriptors and other artefacts (refer to clause 8.4);
- recommendations related to interfaces and associated information models (refer to clause 8.5); and
- other recommendations (refer to clause 8.6).

Annex : Change history

Date	Version	Information about changes
October 2021	0.0.1	Implementation of skeleton and scope contributions approved at EVE#168: - NFVEVE(21)000080: FEAT29 EVE021 Skeleton - NFVEVE(21)000081: FEAT29 EVE021 Scope
December 2021	0.1.0	Implementation of contributions approved at EVE#174 and agreed at EVE#179: - NFVEVE(21)000107r1: FEAT29 EVE021 Clause 4.1 Introduction to energy efficiency - NFVEVE(21)000133r1: FEAT29 EVE021 Clause 4.2 NFV and energy efficiency - NFVEVE(21)000134r1: FEAT29 EVE021 Clause 4.3.1 Energy efficiency in 3GPP Rapporteur actions: - From 107r1: -- The reference [i.ref-mas2020] was deleted, as it is nowhere referenced in the content. -- New reference [i.8] for ETSI EN 303 471 was added, as the document was referenced in the text, but no reference in clause 2.2 has been provided. - Clause 3.1 and 3.3: reference to ETSI GR NFV 003 was corrected, i.e. use GR instead of GS.
February 2022	0.2.0	Implementation of contributions approved at EVE#185: - NFVEVE(21)000017r1: EVE021 Energy Efficiency - clause 4.3 Rapporteur actions: - From 017r1: -- Extended the information and title in the references included from the contribution. -- Added a Rapporteur's Note: the reference to Redfish Telemetry API is not clear. -- Clause 4.3.2, 4.3.3, 4.3.4 and 4.3.5: for all the new tables, a table caption and a text reference has been added following EDR guidelines. The "see note" from the 2nd column header has been deleted, because there are not any notes in the tables to reference.
March 2022	0.3.0	Implementation of contributions agreed at EVE#186: - NFVEVE(22)000027r1: FEAT29 - EVE021 - Proposed update to clause 4.3 - NFVEVE(22)000037r1: FEAT29 EVE021 Clause 5 UC resources sleep resume
April 2022	0.4.0	Implementation of contribution approved at EVE#190: - NFVEVE(22)000043: FEAT29 EVE021 Clause 2.2 correct reference of DMTF Redfish - NFVEVE(22)000051: FEAT29 EVE021 Clause 2.2 additional references for DMTF Redfish - NFVEVE(22)000052r2: FEAT29 EVE021 Clause 5 UC Power saving management function under limited power supply
June 2022	0.5.0	Implementation of contributions approved at EVE#198: - NFVEVE(22)000075r2: FEAT29 EVE021 Clause 5 UC acquiring infrastructure power consumption data - NFVEVE(22)000076r2: FEAT29 EVE021 Clause 5 UC acquiring NFV power consumption data - NFVEVE(22)000077r1: FEAT29 EVE021 Clause 5 UC processing power state driven deployment
August 2022	0.6.0	Implementation of contributions approved at EVE#200 and EVE#202: - NFVEVE(22)000096r1: FEAT29 EVE021 Clause 6 Key issues - NFVEVE(22)000116r1: FEAT29 EVE021 Clause 5 UC power consumption and resources placement optimization
October 2022	0.7.0	Implementation of contributions approved at EVE#208 and EVE#209: - NFVEVE(22)000115r2: FEAT29 EVE021 Clause 5 UC Restoration management from resource reservation in power saving situations - NFVEVE(22)000162: FEAT29 EVE021 Clause 5 UC on dynamic power state management - NFVEVE(22)000161r1: FEAT29 EVE021 Clause 6 Update key issues Rapporteur's actions: - Change the colour of some text in clause 6 to black (it was red for not specific reason).

Date	Version	Information about changes
December 2022	0.8.0	Implementation of contributions approved at EVE#222, EVE#228: - NFVEVE(22)000175: FEAT29 EVE021 Clause 5 UC on dynamic power state management - actors, conditions and flow description - NFVEVE(22)000204: FEAT29 EVE021 Clause 6 Re-organize questions between key issues - NFVEVE(22)000205: FEAT29 EVE021 Clause 7.1 Solutions introduction - NFVEVE(22)000199r2: FEAT29 EVE021 Clause 7 Solutions key issues #1.1 and #1.2 - NFVEVE(22)000206: FEAT29 EVE021 Clause 7 Solutions key issues #1.2 - NFVEVE(22)000209: FEAT29 EVE021 Clause 6 Key issues related to CPU power state management
December 2022	0.9.0	Implementation of contributions approved at EVE#229: - NFVEVE(22)000219: FEAT29 EVE021 Clause 7 Solutions key issues #2.2 resources - NFVEVE(22)000220r1: FEAT29 EVE021 Clause 7 Solutions key issues #2.2 state and config - NFVEVE(22)000221r1: FEAT29 EVE021 Clause 7 Solutions key issues #3.1 VNF power profiles - NFVEVE(22)000222: FEAT29 EVE021 Clause 7 Solutions key issues #3.1 VNF power management
March 2023	0.10.0	Implementation of contributions approved at EVE#239 - NFVEVE(23)000041: FEAT29 EVE021 Clause 7 Solutions key issues #3.3 Power state based zones in NFVI - NFVEVE(23)000042: FEAT29 EVE021 Clause 7 Solutions key issue #2.1 Power state aware allocation of compute resources - NFVEVE(23)000043r1: FEAT29 EVE021 Clause 7 Solutions key issues #2.1 energy policies - NFVEVE(23)000044r1: FEAT29 EVE021 Clause 7 Solutions key issue #3.2 NFV-MANO power consumption - NFVEVE(23)000045: FEAT29 EVE021 Clause 7 Updating key issues by existing solutions - NFVEVE(23)000046r1: FEAT29 EVE021 Clause 7 Solutions key issue #2.1 NFV-MANO capabilities - NFVEVE(23)000047: FEAT29 EVE021 Clause 7.1 Update summary of solutions and related key issues
March 2023	0.10.1	Implementation of contributions approved at EVE#240: - NFVEVE(23)000051: FEAT29 EVE021 Clause 7.1 Update summary of solutions and related key issues (2nd round) - NFVEVE(23)000052: FEAT29 EVE021 Clause 2.2 and 7.2.9.2 Fix missing reference
May 2023	0.11.0	Implementation of contributions agreed at EVE#243: - NFVEVE(23)000065: FEAT29 EVE021 Clause 7.1 Further update summary of solutions and related key issues (3rd round) - NFVEVE(23)000066: FEAT29 EVE021 Clause 7 Solutions key issue #4.2 VIM for acquiring VR power consumption - NFVEVE(23)000067: FEAT29 EVE021 Clause 7 Solutions key issue #4.2 CISM for acquiring containerized power consumption - NFVEVE(23)000068: FEAT29 EVE021 Clause 7.2.6 Update solution to address key issue #4.3 of acquiring VNF power consumption - NFVEVE(23)000069: FEAT29 EVE021 Clause 7 Solutions key issue #4.4 NFVO for acquiring NS power consumption - NFVEVE(23)000070: FEAT29 EVE021 Clause 7 Solutions key issue #4.2, #4.3 and #4.4 NFV power manager for acquiring power consumption

Date	Version	Information about changes
May 2023	0.12.0	Implementation of contributions agreed at EVE#244 and EVE#245: - NFVEVE(23)000074: FEAT29 EVE021 Clause 7.1 Further update summary of solutions (4th round) - NFVEVE(23)000078: FEAT29 EVE021 Clause 7 Solutions key issue #4.1 energy efficiency of NFVI physical resources - NFVEVE(23)000081: FEAT29 EVE021 Clause 7 Solutions key issue #4.2 energy efficiency and metrics of virtualised resources and containerized workloads - NFVEVE(23)000082: FEAT29 EVE021 Clause 7 Solutions key issue #4.3 energy efficiency and metrics of VNF - NFVEVE(23)000083: FEAT29 EVE021 Clause 7 Solutions key issue #4.4 energy efficiency and metrics of NS Rapporteur actions: - Clause 7.2.18.2, added missing reference [i.61] in the first paragraph under Solution #18.b. - Addressed and deleted editor's note in clause 7.2.20.2 about rapporteur action to put correct references to other solutions. - Addressed and deleted editor's note in clause 7.2.22.2 about rapporteur action to put correct references to other solutions.
June 2023	0.13.0	Implementation of contributions agreed at EVE#248: - NFVEVE(23)000098r1: FEAT29 EVE021 Clause 7 Solutions key issue #4.1 PIM for acquiring PR power consumption - NFVEVE(23)000101: FEAT29 EVE021 Clause 5.1 Overview of use cases - NFVEVE(23)000102r1: FEAT29 EVE021 Clean-up of various editor's notes and empty template clauses - NFVEVE(23)000105: FEAT29 EVE021 Clause 9 Conclusions Rapporteur actions: - Deleted an empty row between references [i.61] and [i.62]. - Table 5.3.2-1: in the new sentence added to the action #5 of PIM, added an "on" between "consumption" and "each". - Table 5.5.2-1: added "See note" in the missing reference in action #4 of PIM. - Table 5.9.6-1: removed extra space in the NOTE 3.
July 2023	0.14.0	Implementation of contributions agreed at EVE#249 and EVE#246: - NFVEVE(23)000103r2: FEAT29 EVE021 Clause 7.3 Evaluation of solutions - NFVEVE(23)000104r1: FEAT29 EVE021 Clause 8 Recommendations - NFVEVE(23)000090: FEAT29 EVE021 Clause 7.1 Further update summary of solutions (5th round)

Date	Version	Information about changes
August 2023	0.15.0	Implementation of contributions agreed/approved at EVE#251, EVE#252, EVE#253 and EVE#254 as part of the EVE021 WG final review: - NFVEVE(23)000146r2: FEAT29 EVE021 Editorial review (used as the baseline for this present version). - NFVEVE(23)000116: FEAT29 EVE021 EN resolution vCPU power characteristics and VNF power profiles mapping - NFVEVE(23)000122: FEAT29 EVE021 Clause 7.1 Final update summary of solutions (6th round) - NFVEVE(23)000127: FEAT29 EVE021 Review removing empty clauses - NFVEVE(23)000128: FEAT29 EVE021 Review addressing editorial editor's notes - NFVEVE(23)000126r1: FEAT29 EVE021 Clause 8 Recommendations regarding solution #23 - NFVEVE(23)000133r1: FEAT29 EVE021 EN resolution of power consumption correlation with Sol#23 - NFVEVE(23)000143r1: FEAT29 EVE021 Multiple clauses Addressing remaining EN - NFVEVE(23)000147: FEAT29 EVE021 Clause 8.3 Addressing EN recommendation of VNF power management - NFVEVE(23)000148: FEAT29 EVE021 Clause 8.7 Final clean-up of EN on recommendations - NFVEVE(23)000149r4: FEAT29 EVE021 Multiple clauses Small technical fixes Rapporteur actions: - Several clauses: additional editorial review comments: change "concerns to" to either "concerns" or "is about", and remove the "on" in the "influence on". - Further renumbering of the recommendations to convey the new recommendations added by other contributions, in addition to the baseline renumbering proposed by contribution 146r2. - Clause 7.2.15: undo the deletion of "NFV power metrics" in the first paragraph, as wrong deletion was shown in 146r2. - Clause 8.7: renumber the notes. - Clause 4.3.5: added the "see note" to the "Scope" column, and the note at the end of the table, like in other clauses, as proposed by contribution 149r4. - Deleted the "Annex: Bibliography"

History

Document history		
V5.1.1	September 2023	Publication