



GROUP SPECIFICATION

Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Requirements and interface specification for Physical Infrastructure Management

Disclaimer

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

Reference

RGS/NFV-IFA053ed541

Keywordsanalysis, information model, management, MANO,
model, PIM, profile**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

Important notice

The present document can be downloaded from the
[ETSI Search & Browse Standards](#) application.

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on [ETSI deliver](#) repository.

Users should be aware that the present document may be revised or have its status changed,
this information is available in the [Milestones listing](#).

If you find errors in the present document, please send your comments to
the relevant service listed under [Committee Support Staff](#).

If you find a security vulnerability in the present document, please report it through our
[Coordinated Vulnerability Disclosure \(CVD\)](#) program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.

The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2026.
All rights reserved.

Contents

Intellectual Property Rights	5
Foreword.....	5
Modal verbs terminology.....	5
1 Scope	6
2 References	6
2.1 Normative references	6
2.2 Informative references.....	6
3 Definition of terms, symbols and abbreviations.....	8
3.1 Terms.....	8
3.2 Symbols.....	8
3.3 Abbreviations	9
4 Overview	9
4.1 Introduction	9
4.2 Background of Physical Infrastructure Management in NFV	9
4.2.1 Physical Infrastructure Management in enabling energy efficiency in NFV	9
4.2.2 Physical Infrastructure Management in CIS cluster management	10
4.2.3 Physical Infrastructure Management in virtualised RAN	11
4.2.4 Physical Infrastructure Management in enabling network connectivity integration and operationalization in NFV	11
4.3 Input from other organizations	12
4.3.1 Physical Infrastructure Management in Linux® Foundation Anuket	12
4.3.2 DMTF® Redfish®	12
5 Architectural Framework for Physical Infrastructure Management.....	13
5.1 Overview	13
5.2 PIM function and PIM services.....	13
6 Requirements for Physical Infrastructure Management	13
6.1 Introduction	13
6.2 Resource provisioning and lifecycle management	14
6.3 Inventory management	14
6.4 Topology management	16
6.5 Performance, fault and log management	17
7 Service requirements related to physical infrastructure management	19
7.1 Introduction	19
7.2 General service requirements	19
7.3 Physical resource provisioning and lifecycle management service interface requirements	19
7.4 Physical resource inventory management service interface requirements.....	20
7.5 Physical resource topology management service interface requirements.....	20
7.6 Physical resource performance management service interface requirements.....	20
7.7 Physical resource fault management service interface requirements	21
7.8 Physical resource log management service interface requirements.....	21
8 Information Modelling for Physical Infrastructure Management.....	22
8.1 Introduction	22
8.2 Relationship between information elements.....	22
8.3 Information elements.....	23
8.3.1 PhysicalResourceInventory information element	23
8.3.2 NfviPopInfo information element.....	23
8.3.3 ResourceZone information element.....	23
8.3.4 PhysicalResourcePool information element	23
8.3.5 PhysicalResourceInfo information element	24
Annex A (informative): Architectural Analysis of Physical Infrastructure Management	25

A.1	Overview and background.....	25
A.2	Placement options for PIM functionality	25
A.2.1	Introduction	25
A.2.2	Option 1: PIM functionality as part of the VIM.....	25
A.2.3	Option 2: PIM functionality as a new function in NFV-MANO.....	25
A.2.4	Option 3: PIM functionality distributed into different entities	26
A.2.5	Analysis.....	26
A.3	Architecture design considerations	28
A.3.1	Introduction	28
A.3.2	Reference point approach.....	28
A.3.3	SBA Approach	29
A.3.4	Analysis.....	29
A.4	Conclusion.....	30
Annex B (informative): Analysis of existing solutions		32
B.1	Introduction	32
B.2	OpenStack® Ironic	32
B.2.1	Overview	32
B.2.2	Features	32
B.3	Metal3.....	33
B.3.1	Overview	33
B.3.2	Features	34
B.4	Redfish®	35
B.4.1	Overview	35
B.4.2	Services and Models.....	35
B.4.3	Implementations	36
B.4.4	Provisioning and lifecycle management of hardware resources.....	36
B.4.5	Composition Service	37
B.4.6	Telemetry Service.....	37
B.4.7	Log Service	37
B.4.8	Event Service.....	38
B.5	Analysis.....	38
Annex C (informative): PlantUML source code for the information model.....		40
C.1	Source code	40
Annex D (informative): Change history		41
History		44

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "*Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards*", which is available from the ETSI Secretariat. Latest updates are available on the [ETSI IPR online database](#).

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

Trademarks

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

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

Foreword

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

Modal verbs terminology

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

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

1 Scope

The present document specifies requirements for physical infrastructure management in the NFV-MANO framework and service interface specifications for managing physical infrastructure with respect to life cycle, inventory, topology, log, configuration, performance and fault management of physical compute, storage and networking resources. The present document also specifies an information model for describing NFVI physical resources to support relevant interface requirements.

In addition, architectural analysis of physical infrastructure management within the NFV-MANO framework and analysis of existing solutions for management of physical resources are included as part of informative annexes to the present document.

2 References

2.1 Normative references

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

Referenced documents which are not found to be publicly available in the expected location might be found in the [ETSI docbox](#).

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

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

- [1] [IETF RFC 5424](#): "The Syslog Protocol".
- [2] [ETSI GS NFV-IFA 027](#): "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Performance Measurements Specification".
- [3] [ETSI GS NFV-IFA 045](#): "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Faults and alarms modelling specification".

2.2 Informative references

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

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

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

- [i.1] ETSI GR NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".
- [i.2] ETSI GR NFV-EVE 021: "Network Functions Virtualisation (NFV) Release 5; Evolution and Ecosystem; Report on energy efficiency aspects for NFV".
- [i.3] ETSI GR NFV-IFA 046: "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on NFV support for virtualisation of RAN".
- [i.4] ETSI GS NFV-IFA 036: "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Requirements for service interfaces and object model for container cluster management and orchestration specification".

- [i.5] ETSI GS NFV-IFA 040: "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Requirements for service interfaces and object model for OS container management and orchestration specification".
- [i.6] ETSI GR NFV-IFA 035: "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on network connectivity integration and operationalization for NFV".
- [i.7] [Anuket Reference Model for Cloud Infrastructure \(RM\)](#).
- [i.8] [GSM Association™, Official Document NG.126](#): "Cloud Infrastructure Reference Model, Version 4.0 (30 November 2023)".

NOTE 1: This document is created by Linux® Foundation Anuket, see [i.7], and is also published by GSMA Networks Group.

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

- [i.9] [DMTF® DSP0266](#): "Redfish Specification".
- [i.10] [DMTF® DSP0268](#): "Redfish Data Model Specification".
- [i.11] O-RAN.WG6.O2-GA&P-R003-v06.00: "O-RAN Working Group 6; O2 Interface General Aspects and Principles".
- [i.12] ETSI GS NFV 006: "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Architectural Framework Specification".
- [i.13] ETSI GR NFV-IFA 039: "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on Service Based Architecture (SBA) design".
- [i.14] [OpenStack® Ironic supported documentation \(23.2 series\)](#).

NOTE: The OpenStack® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. ETSI is not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

- [i.15] [OpenStack® Bare Metal API reference documentation](#).
- [i.16] [Metal³ user guide](#).
- [i.17] [Kubernetes®: "The Cluster API book"](#).
- [i.18] Void.
- [i.19] [DMTF Redfish® Developers Hub](#).
- [i.20] [DMTF Redfish® Interop Validator](#).
- [i.21] [Open Compute Project's Profiles](#).
- [i.22] [Open Process Automation Forum's Profiles](#).
- [i.23] [Sushy - Python library to communicate with Redfish based systems](#).
- [i.24] [DMTF® DSP2066](#): "Redfish Fabrics White Paper".
- [i.25] [DMTF® DSP2051](#): "Redfish Telemetry White Paper".
- [i.26] [DMTF® DSP2050](#): "Redfish Composability White Paper".
- [i.27] [DMTF® DSP2062](#): "Redfish Firmware Update White Paper".
- [i.28] [DMTF® DSP2059](#): "Redfish Certificate Management White Paper".
- [i.29] [DMTF® DSP2065](#): "Redfish Message Registry Guide".
- [i.30] Void.

- [i.31] O-RAN.WG6.O2IMS-INTERFACE-R003-v05.00: "O-RAN Working Group 6; O2ims Interface Specification".
- [i.32] O-RAN.WG6.O-CLOUD-IM.0-R003-v01.00: "O-RAN Working Group 6 (Cloudification and Orchestration Workgroup); O-Cloud Information Model".

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].

Network Functions Virtualisation Infrastructure (NFVI): totality of all hardware and software components that build up the environment in which VNFs are deployed

NOTE 1: The above definition is the same as the one defined in ETSI GR NFV 003 [i.1].

NOTE 2: There can be different infrastructure layers (or grouping sets) in the NFVI, i.e. physical infrastructure layer, virtualised infrastructure layer, and CIS infrastructure layer, which together build up the environment for the deployment of VNFs.

NFVI resource: building block in the NFVI

NOTE 1: NFVI resource can be physical resource, virtualised resource or CIS resource.

NOTE 2: Virtualised resources include virtual compute, virtual network and virtual storage resources.

EXAMPLE 1: Example of virtual compute resource is a Virtual Machine (VM). Example of virtual storage resource is a volume or object based virtual storage.

EXAMPLE 2: Examples of CIS resource are CIS instance, CISM instance, and CIS cluster node.

physical infrastructure: set of physical resources comprising the NFVI

Physical Infrastructure Management (PIM): management of physical resources and associated firmware/software elements

Physical Infrastructure Management (PIM) function: function within NFV-MANO which is responsible for management of physical resources

physical resource: hardware component that makes up the physical infrastructure of the NFVI

NOTE: Physical resource can be physical compute, physical network or physical storage, and have respective firmware/software enabling its operation and usage.

EXAMPLE: Servers, switches, routers, storage systems.

physical resource inventory: collection of all physical assets within NFVI

physical resource pool: logical grouping of physical resources within NFVI

NOTE 1: A physical resource pool can be solely based on a certain resource type (e.g. compute, storage, networking) or include a combination of them, and can span zero, one or multiple resource zones.

NOTE 2: A physical resource can be part of none, one or more than one physical resource pool.

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:

BM	Bare Metal
BMC	Baseboard Management Controller
CAPM3	Cluster API Provider Metal3
CSDL	Common Schema Definition Language
DCIM	DataCentre Infrastructure Management
DMTF	Distributed Management Task Force
FRU	Field Replaceable Unit
IPMI	Intelligent Platform Management Interface
OAM&P	Operations, Administration, Management and Provisioning
PIM	Physical Infrastructure Management
RAID	Redundant Array of Independent Disks
RAN	Radio Access Network
UUID	Universally Unique Identifier
VIF	Virtual InterFace
vRAN	virtualised RAN

4 Overview

4.1 Introduction

The present document introduces the concept of physical infrastructure management in the NFV framework and addresses management aspects, such as inventory management, resource provisioning, FM/PM monitoring, etc., of NFVI physical resources.

Management of physical resources plays an important role in enabling new NFV use cases and deployment scenarios. These include, but are not limited to, the following:

- enable energy efficient NFV deployments;
- support virtualised RAN (vRAN) use cases in NFV;
- support infrastructure management for bare-metal CIS clusters;
- enable network connectivity integration and operationalization.

Clause 4 summarizes the role of the Physical Infrastructure Management (PIM) function in NFV based on the relevant use cases and scenarios in ETSI GR NFV-EVE 021 [i.2], ETSI GR NFV-IFA 046 [i.3], ETSI GS NFV-IFA 036 [i.4] and ETSI GR NFV-IFA 035 [i.6].

4.2 Background of Physical Infrastructure Management in NFV

4.2.1 Physical Infrastructure Management in enabling energy efficiency in NFV

Several aspects of NFV that have an impact on power consumption and energy efficiency are investigated in ETSI GR NFV-EVE 021 [i.2]. In addition to documenting multiple use cases related to energy efficiency in NFV, ETSI GR NFV-EVE 021 [i.2] identifies relevant key issues and proposes potential solutions to address those issues.

In the context of energy efficiency use cases described in ETSI GR NFV-EVE 021 [i.2], the PIM manages the physical resources in the NFVI considering, but not limited to, the following aspects:

- **Physical resource management:** The PIM manages the overall physical resources in the NFVI, such as physical compute, storage and network resources. It maintains resource topology and inventory of its managed entities, indicating their supported capability and energy efficiency related metrics. The PIM is also capable of providing notifications in the events of changes in the state of its managed entities, e.g. notifications related to power state restoration of the physical resources, etc. The PIM is also envisioned to manage pool of physical resources based on different parameters, such as power consumption metrics, power management policies, etc.
- **Physical resource configuration management:** The PIM manages the power state and capabilities of physical resources in the NFVI such as changing power states (power on/power off) of hardware devices, adjusting input power supply, controlling Voltage/Frequency (V/F) ratio of CPUs, etc.
- **Power consumption monitoring:** The PIM acquires power consumption data of its managed physical resources by interacting with the relevant Baseboard Management Controllers (BMCs). The PIM also supports providing power consumption information and energy related metrics of physical resources to other NFV-MANO entities via its service interface(s).

4.2.2 Physical Infrastructure Management in CIS cluster management

For the deployment of containerized VNFs in NFV, a group of Container Infrastructure Service (CIS) cluster nodes is provided from the NFVI, and the corresponding compute resources are allocated by NFV-MANO. The Container Infrastructure Service Management (CISM) function, described in ETSI GS NFV-IFA 040 [i.5] is configured to use these CIS cluster nodes hosting CIS instance functionality for the deployment and management of containerized workloads. The composition of all CIS cluster nodes, hosting one or multiple CISM instances and CIS instances, form together a CIS cluster. ETSI GS NFV-IFA 036 [i.4] specifies requirements on CIS Cluster Management (CCM) services. CIS cluster nodes can be realized by either Virtual Machines (VMs) or by bare-metal servers or a combination of both, i.e. hybrid CIS clusters.

The CCM function, responsible for managing CIS clusters, allocates infrastructure resources for CIS cluster nodes. In case of bare-metal CIS clusters, the CCM allocates bare-metal servers for the CIS cluster by interacting with an infrastructure management function that manages available hardware resources and can assign bare-metal servers to the consumers. This infrastructure management functionality can be provided by the PIM.

In the context of CIS cluster management specified in ETSI GS NFV-IFA 036 [i.4], the PIM manages the physical resources in the NFVI considering, but not limited to, the following aspects:

- **Enable allocation of bare-metal servers to CIS clusters:** The PIM manages physical resources (e.g. bare-metal servers, storage servers, etc.) and can assign them to pool of resources that can in turn be assigned to a CIS cluster upon receiving request from the CCM, as described in clauses 4.2.5.3 and B.2 of ETSI GS NFV-IFA 036 [i.4].
- **Physical resource lifecycle management:** The PIM manages the lifecycle of physical resources in the NFVI, such as physical compute, storage and network resources. It offers resource management operations for managing bare-metal servers via its exposed (service) interfaces, similar to VIM in case of VM-based CIS clusters. See clause 4.2.9 of ETSI GS NFV-IFA 036 [i.4].
- **Physical resource fault and performance management:** The PIM offers performance information of bare-metal resources to the CCM, which in turn maps this information to the respective CIS cluster related object instances as described in clause 4.2.12 of ETSI GS NFV-IFA 036 [i.4]. The PIM also provides alarms related to physical resources that can be collected and mapped by the CCM for alarm management of CIS clusters.

4.2.3 Physical Infrastructure Management in virtualised RAN

NFV-MANO support for the virtualisation of the RAN is studied in ETSI GR NFV-IFA 046 [i.3], where several challenges are reported regarding the compatibility levels between NFV-MANO and the virtualisation framework developed by the O-RAN Alliance. In the O-RAN framework, infrastructure resource management is the responsibility of the Infrastructure Management Services (IMS) residing in O-Cloud [i.11]. The IMS is responsible for management of the O-Cloud resources and the software which is used to manage those resources. Similar to the NFVI case, in O-RAN, an O-Cloud Node Cluster can be of different types (i.e. Bare Metal Container Cluster, VM-based Container Cluster, VM Cluster, OS Cluster). For example, VM-based Container Cluster can be mapped to a CIS cluster based on VMs and a Bare Metal Container Cluster can be mapped to a bare metal CIS cluster. In principle, IMS provides collectively the functionality of virtualised infrastructure management, physical infrastructure management and CIS cluster nodes management, which in NFV is provided by the VIM, PIM and CCM respectively.

In support of virtualised RAN use cases described in ETSI GR NFV-IFA 046 [i.3], the PIM manages the physical resources in the NFVI considering, but not limited to, the following aspects:

- Physical resource management (as described in clause 4.2.1).
- Physical resource configuration management (as described in clause 4.2.1).
- Power consumption monitoring (as described in clause 4.2.1).
- Allocation of bare-metal servers to O-Cloud Node Clusters (similarly for CIS clusters as described in clause 4.2.2).
- Physical resource lifecycle management (as described in clause 4.2.2).
- Physical resource fault and performance management (as described in clause 4.2.2).

Information modelling for physical infrastructure management considered by PIM is specified in clause 8 of the present document. Information modelling considered by O-RAN Alliance for the elements comprising the O-Cloud is specified in O-RAN.WG6.O-CLOUD-IM.0-R003-v01.00 [i.32]. IMS services exposed over the O2IMS interface, which is specified in O-RAN.WG6.O2IMS-INTERFACE-R003-v05.00 [i.31], are related to infrastructure inventory, infrastructure monitoring and infrastructure lifecycle management. A detailed profiling of the two referenced specifications from O-RAN is not addressed in the present document.

4.2.4 Physical Infrastructure Management in enabling network connectivity integration and operationalization in NFV

ETSI GR NFV-IFA 035 [i.6] studies the integration and operationalization aspects of network connectivity in NFV. Solutions proposed in the ETSI GR NFV-IFA 035 [i.6] consider physical infrastructure management as an enabler for Operations, Administration, Management and Provisioning (OAM&P) of network devices in NFVI.

In support of network connectivity integration and operationalization described in ETSI GR NFV-IFA 035 [i.6], the PIM manages the physical resources in the NFVI considering, but not limited to, the following aspects:

- **Device Management:** The PIM manages the network devices on different levels, such as NFVI node level, NFVI-PoP network fabric level, etc. For the case of NFVI node level devices, the PIM interacts with one or multiple BMCs to manage Network Interface Cards (NICs)/adapters of NFVI nodes. At the NFVI-PoP level, the PIM manages devices forming the NFVI-PoP network fabric, such as switches, ports, and NFVI-PoP gateway devices, such as Customer Edge (CE) devices.
- **Relationship with SDN controller and other entities:** For NFVI-PoP network fabric and gateway management, the PIM performs device-level provisioning of underlay network devices, while an SDN controller can perform network level provisioning. Other NFV-MANO entities, such as VIM, CCM, etc. can interact with the PIM for the management of NFVI-PoP network elements.

4.3 Input from other organizations

4.3.1 Physical Infrastructure Management in Linux® Foundation Anuket

Linux® Foundation Anuket describes their view on infrastructure and its management in their Reference Model (RM) for Cloud Infrastructure [i.7], which is also available through GSMA [i.8]. The infrastructure model covers Virtual Infrastructure Layer, Hardware Infrastructure Layer and Workload Layer.

NOTE 1: Anuket uses the terms hardware infrastructure and hardware resource in their specifications related to physical infrastructure management. In the context of the present document, hardware resource and hardware infrastructure refer to physical resource and physical infrastructure respectively.

The model of hardware infrastructure resources distinguishes between cloud infrastructure (compute, storage and network resources) and hardware acceleration resources. The hardware infrastructure manager is responsible for equipment management and enables:

- provision, manage, monitor and delete hardware resources;
- physical hardware resource discovery, monitoring and topology;
- manage hardware infrastructure telemetry and log collection service.

NOTE 2: When Anuket mentions deleting hardware resources, this refers to the deletion of logical resources.

Anuket references Redfish® standard by DMTF®. More information on DMTF Redfish® standard is provided in clauses 4.3.2 and B.4.

4.3.2 DMTF® Redfish®

The Redfish® standard specifies a HTTP RESTful interface for exposing system manageability of infrastructure platforms. The Redfish® standard consists of an interface specification (see DMTF® DSP0266 [i.9]) and model specification (see DMTF® DSP0268 [i.10]). The Redfish® model is utilized for managing datacentre platforms (compute, storage and networking) as described in Open Compute Project's Profiles [i.21], DCIM and process automation components as described in the Open Process Automation Forum's Profiles [i.22]. The model primarily covers managing the physical and functional aspect of multiple systems.

The Redfish® model is expressed in three schema formats: OpenAPI, json-schema and Common Schema Definition Language (CSDL). This enables the tool-chains available for those schema formats.

The Redfish® Service which implements the Redfish® interface can:

- manage the physical hierarchy, and functional aspect of multiple systems, in addition to their connectivity (fabrics);
- manage the management hierarchy, which represents the layers of management (for example from datacentre manager to rack manager to node manager to component manager);
- manage sessions, events, accounts, telemetry and logs;
- manage the composition and aggregation of systems;
- manage certificates;
- update firmware, reset the system, configure BIOS and network ports.

Further information on DMTF Redfish® is provided in clause B.4.

5 Architectural Framework for Physical Infrastructure Management

5.1 Overview

Within NFV-MANO, the PIM function is responsible for the physical infrastructure management and provides its capabilities through a set of management services.

The services for the management of physical infrastructure are exposed via management service interfaces produced by the PIM function. These management service interfaces can be consumed by:

- other NFV-MANO functional entities; and/or
- consumers outside NFV-MANO (e.g. OSS/BSS).

NOTE: For the sake of conciseness, the 'PIM function' is also referred to as 'the PIM' throughout the present document.

5.2 PIM function and PIM services

The PIM provides multiple management services for physical infrastructure management. These services (among others) include physical resource provisioning and lifecycle management service, physical resource inventory management service, physical resource performance management service. The full set of physical infrastructure management services is specified in clause 7. These management services are exposed via their respective service interfaces produced by the PIM.

Figure 5.2-1 illustrates a logical representation of the PIM function exposing its various management services via respective management service interfaces.

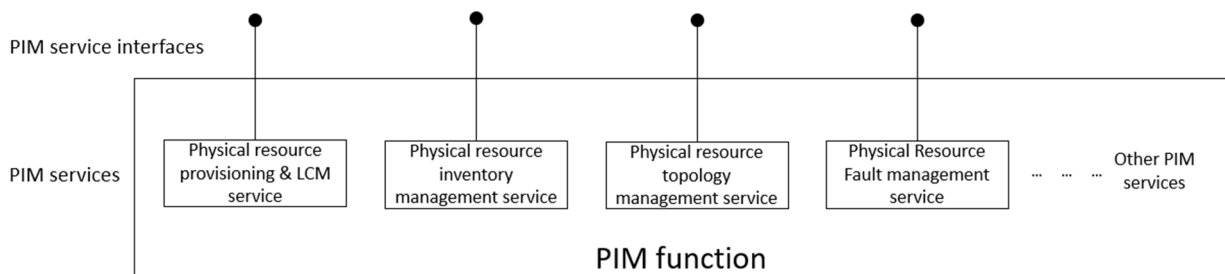


Figure 5.2-1: Logical representation of PIM function and its services

NOTE: Not all the PIM services are shown in figure 5.2-1. The management services that the PIM provides are specified in clause 7 of the present document.

6 Requirements for Physical Infrastructure Management

6.1 Introduction

Clause 6 defines requirements related to different aspects of physical infrastructure management in NFV framework. These aspects include, but are not limited to, the following:

- resource provisioning and lifecycle management;
- inventory management;

- topology management;
- performance, fault and log management.

6.2 Resource provisioning and lifecycle management

Resource provisioning of physical resources refers to setting up, preparing, and if required, configuring physical compute, storage, and network resources in the NFVI. The provisioned resources can then be used by NFV-MANO components, e.g. VIM, CISM, CCM, to support orchestration and management of virtualised resources, containerized workloads, and CIS clusters, supporting the deployment and execution of NSs and VNFs. Resource provisioning includes setting up systems comprising of different physical resources for provisioning composable infrastructure resources. One example of such composable physical resources is a computer system comprising of physical compute, network, storage and/or acceleration resources, that can then be provisioned as an individual physical resource for further use.

Regarding physical resource provisioning, physical infrastructure management shall support capabilities to:

- compose physical resources from disaggregated resources available in the physical resource pools;
- prepare physical resources for deployment (e.g. powering on, configuring network and boot settings for bare metal servers);
- maintain information about provisioning state of physical resources (e.g. during different stages of provisioning);
- reprovision physical resources (e.g. reconfiguring the physical resources for other usages or setups).

Lifecycle management of physical (compute, network, and storage) resources, as managed objects, refers to creating logical instances for the managed objects, updating their information, and deleting them when they are no longer in use. Examples of updating information of such managed objects are setting a new provision state of a physical resource, changing power state of a physical resource, etc.

Regarding lifecycle management of physical resources, physical infrastructure management shall support capabilities to:

- keep and store information about the provisioned physical resources;
- provide physical resource information to consumers, and enable the filtering of resource data based on the request from the consumer;
- create, update and delete physical resource information elements upon request by a consumer, including the creation/retrieval/update/deletion of properties and attributes (e.g. volumes, interfaces, assigned traits, etc.) associated with an individual resource during its lifecycle;
- update the state of the physical resource based on internal change in state or on-demand state change requests of the resource (e.g. change in power state of a resource, change in provision state due to maintenance or failure, etc.);
- provide notifications to consumer for the subscribed events.

Information about a physical resource managed by PIM is modelled in the PhysicalResourceInfo information element specified in clause 8 of the present document.

6.3 Inventory management

Inventory of physical infrastructure concerns information of Field Replaceable Unit (FRU) hardware and associated firmware and software assets within the NFVI.

Inventory of physical resources shall include information about:

- node racks infrastructure, such as racks, shelves and slots;
- compute hardware and its firmware/software;

- storage hardware and its firmware/software;
- network hardware and its firmware/software; and
- firmware/software for the virtualisation of the physical resources.

Figure 6.3-1 illustrates the scope of NFVI "physical" resources for which inventory information is considered.

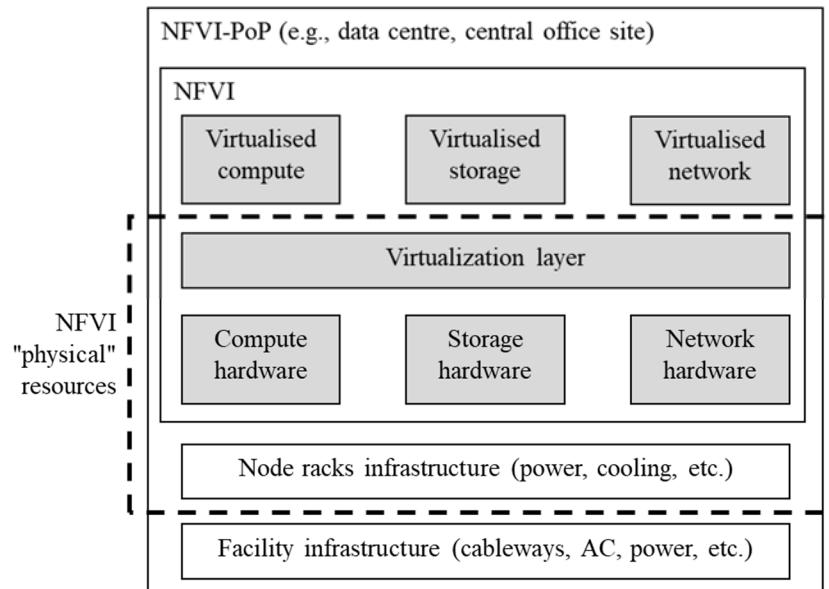


Figure 6.3-1: NFVI "physical" resources for inventory information

The inventory information can be organized either hierarchically (e.g. components are enclosed within other components thus creating a containment relationship) or in a flat manner. Figure 6.3-2 illustrates an example of the kinds of inventory information that can be kept about the resources and the relationship with physical infrastructure management.

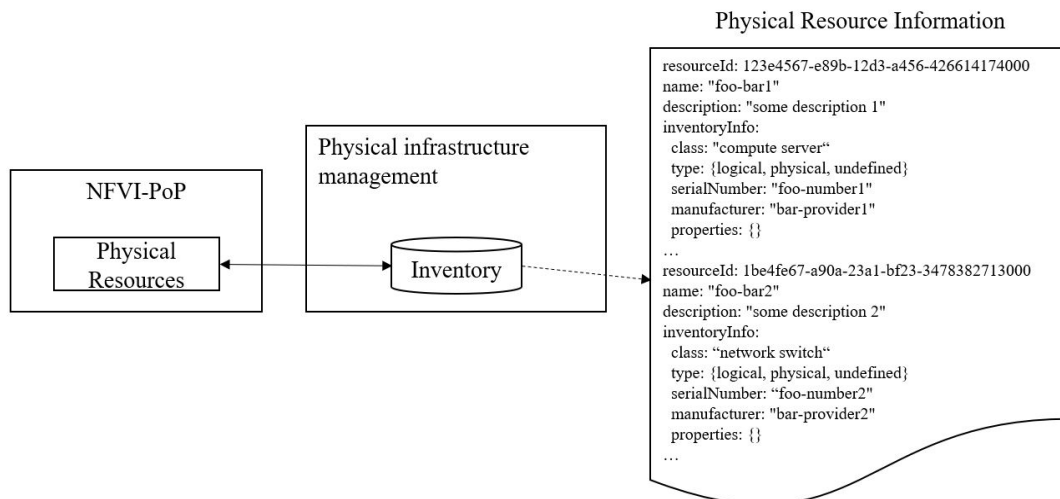


Figure 6.3-2: Example of inventory information

Figure 6.3-2 only provides an example of the kind of information available in the relevant information element, i.e. PhysicalResourceInfo. Only the inventory related information of the physical resource is highlighted in the above example. The whole set of information contained in the PhysicalResourceInfo information element is specified in clause 8 of the present document.

Regarding inventory management, the physical infrastructure management shall be capable to:

- keep and store information about the resource;
- provide the inventory information to consumers, and enable the filtering of inventory data based on the request from the consumer;
- update inventory information upon request by a consumer;
- update inventory information based on the discovery of resources from the NFVI.

Inventory information about a physical resource shall be included in the PhysicalResourceInfo information element specified in clause 8 of the present document.

6.4 Topology management

Topology management of physical resources concerns the management of the relationship between the physical resources that conform the NFVI. Topology information includes:

- physical topology: it concerns the relationship between physical resources from a physical perspective; and
- logical topology: it concerns the relationship between physical resources from a logical perspective.

Logical topology information is necessary to convey topological relationship among physical resources that are not uniquely confined to physical constraints. For instance, logical topology information can be used to define logical grouping of resources, such as physical resource pools, resource aggregates, resource zones, etc. Some of the variables to determine physical resource pool and resource zones can consider physical constraints such as physical location in the NFVI-PoP, power supply dependencies, etc., but not only. From a topology management perspective, the following scopes shall be considered by the physical infrastructure management:

- Physical resource pool: logical grouping of physical resource. The physical resource pool can be based solely on certain type of resource (e.g. compute, storage, network) or include a combination of them. A physical resource can be part of none, one or more than one physical resource pool. This is the same concept as defined in ETSI GR NFV 003 [i.1]. See also definition in clause 3.1.
- Resource zone: logical grouping of physical resources according to some physical isolation, redundancy capability or certain administrative policy for the NFVI. For instance, physical isolation can be achieved by using separate power supplies or considering resources spread on different physical building sites. This is the same concept as defined in ETSI GR NFV 003 [i.1].

Examples of use of physical resource pools include the grouping of:

- Resources with similar resource characteristics, e.g. set of compute resources with acceleration capabilities.
- Resources with same resource configuration, e.g. CPU set at certain power state.
- Resources with identified specific usage purpose, e.g. to be used by certain tenant.

Furthermore, physical resource pools can be composed (or nested), i.e. a resource pool can represent the composition of one or more other physical resource pools, or a resource pool can be nested into another resource pool. For instance, a physical resource pool can group compute physical resources with certain property, while another physical resource pool can group storage physical resources with another certain property, and the combination of both physical resource pools can be in turn grouped into a composed physical resource pool to represent the group of compute and storage physical resources used for allocating resources to a specific CIS cluster (see clause 4.2.5.3 of ETSI GS NFV-IFA 036 [i.4]).

Figure 6.4-1 illustrates an example of the concepts of resource zones and physical resource pools.

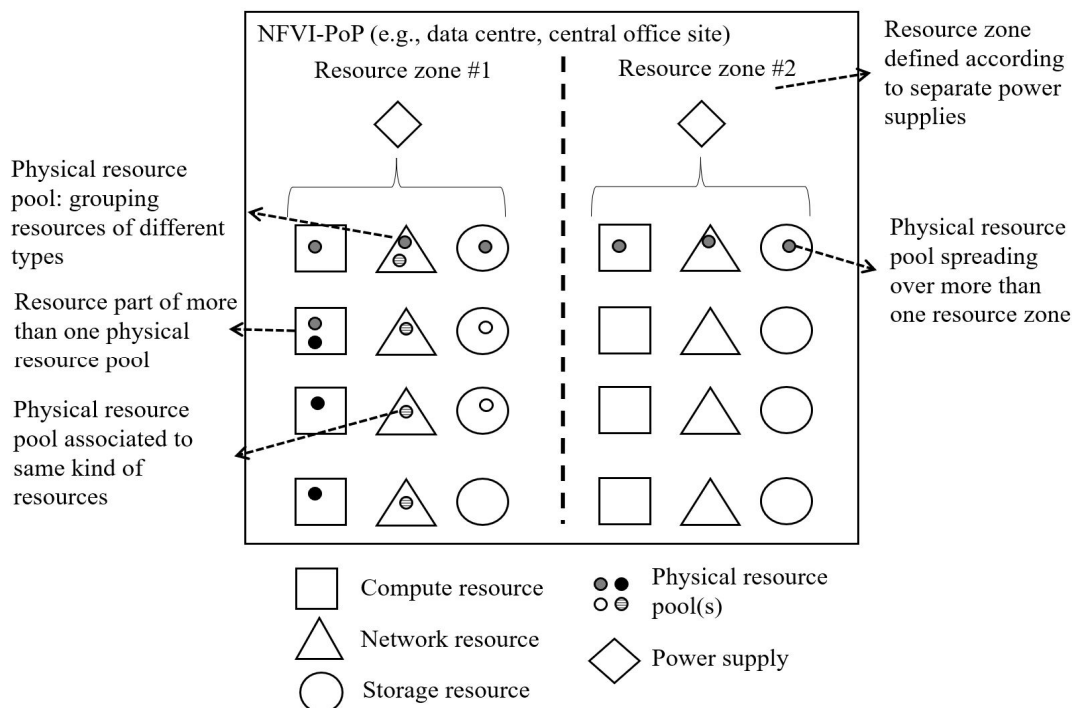


Figure 6.4-1: Example of resource zones and physical resource pools

The managed topology information shall support identifying the type of resource relationship from a physical perspective, according to:

- Containment: the resource is contained within another resource, e.g. a CPU is part of a bare metal compute server, or a compute server is within certain rack. This relationship also includes the case when some artifact resource is "installed" in another resource, e.g. a firmware installed in a bare metal compute server.
- Connection: this includes network connectivity and power connectivity, e.g. the resource is connected at L1 level with some other resource, or the resource gets its power supply from certain connected power distribution unit.

Regarding topology management, the physical infrastructure management shall be capable to:

- keep and store information about the physical and logical topology of physical resources in the NFVI;
- provide topology information to consumers, and enable the filtering of topology data based on the request from the consumer;
- create, update and delete logical topology information upon request by a consumer, including the creation/update/delete of physical resource pools and association of properties to the physical resource pools for topology and other management aspects (e.g. provisioning, monitoring);
- update the topology information based on the discovery and update of resources within the NFVI.

6.5 Performance, fault and log management

The physical resources, as objects managed by the physical infrastructure management, can have associated performance metrics and be subject to faults, for which alarms can be associated with and be exposed. Furthermore, physical resources and/or their management systems can generate logs.

Regarding performance management, the physical infrastructure management shall support capabilities to:

- manage the collection of performance data (i.e. measurements), e.g. by setting up PM jobs;
- provide to consumers the performance data, e.g. via streaming or performance file reporting; and

- notify to consumers about the availability of performance data, about performance degradation, and any other events related to the performance management.

Performance data associated to a physical resource is related to the properties and type of resource. The physical infrastructure management shall be capable to monitor and collect measurements of metrics that either are defined as counters (i.e. a cumulative metric representing a single monotonically increasing counter, whose value can only increase or reset to zero) or gauge (i.e. a metric representing a single numerical value that can either go up or down). Depending on the type of physical resource, metrics about average usage and peak usage over defined time periods shall be made available. For all physical resources that consume power/energy, measurements about the power consumptions shall also be provided. For collection of energy efficiency related KPIs, physical infrastructure management shall be capable to assess the productivity (or lack thereof) of a physical resource in a given time period. This refers to a physical resource being in "active" or "idle" state within a given time interval, where "active" state represents the state when productive work is being performed by the physical resource, and "idle" state refers to the state when the physical resource is not performing any productive work.

EXAMPLE: In the context of a physical resource of type storage, active state can refer to the state when the physical resource is performing data read/write operations (Input/Output operations) and/or data transfer operations. Alternatively, idle state implies that no operations are being performed by the physical resource.

Performance measurements related to physical resources managed by the PIM are specified in ETSI GS NFV-IFA 027 [2]. The PIM produces these performance measurements and makes them available to the consumers.

Regarding fault management, the physical infrastructure management shall support capabilities to:

- collect failure information and provide information to consumers;
- generate alarms based on fault information about physical resources;
- provide alarms information to consumers, via consumer issued query requests;
- notify about alarms to consumers, and any other events related to fault management; and
- manage the alarms, such as enabling to clear (automatically or on-demand) and acknowledge alarms.

Provided alarms shall conform to a standard specification, and their information include: identification of the alarm, associated managed object, root cause information, timing information about when the alarm is raised, changed, cleared, and acknowledged; information about when the fault event was observed; information about the fault properties, including perceived severity, type of event, type of fault, probable cause and fault details; and when the information about correlated alarms is available.

Alarms related to physical resources, as defined in clause 7.9 of ETSI GS NFV-IFA 045 [3], are provided by the PIM and are visible to the PIM Consumer.

Finally, regarding log management, the physical infrastructure management shall support capabilities to:

- collect logs about physical resources;
- manage the jobs/means of logging about physical resources, such as setting up the log levels (as defined in clause 6.2.1, table 2 of IETF RFC 5424 [1]) and the criteria of what information to log, and when to start and stop the log collection;
- provide information about the log management, such as responding to query requests of information about the log collection;
- notify about events related to log management, such as the availability of certain logs; and
- provide information about where and how to retrieve the logs.

7 Service requirements related to physical infrastructure management

7.1 Introduction

Clause 7 in the present document specifies the set of requirements applicable to interfaces exposing physical resource management services provided by the PIM function.

7.2 General service requirements

Table 7.2-1 specifies requirements applicable to the services provided by the PIM.

Table 7.2-1: PIM service requirements

Identifier	Requirement
PimSvc.001	The PIM shall provide a physical resource provisioning and lifecycle management service.
PimSvc.002	The PIM shall provide a physical resource inventory management service.
PimSvc.003	The PIM shall provide a physical resource topology management service.
PimSvc.004	The PIM shall provide a physical resource performance management service.
PimSvc.005	The PIM shall provide a physical resource fault management service.
PimSvc.006	The PIM shall provide a physical resource log management service.

7.3 Physical resource provisioning and lifecycle management service interface requirements

Table 7.3-1 specifies requirements applicable to the physical resource provisioning and lifecycle management service interface produced by the PIM.

Table 7.3-1: Physical resource provisioning and lifecycle management service interface requirements

Identifier	Requirement
PimPrLcmMgt.001	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support provisioning of physical resources (see notes 3 and 4).
PimPrLcmMgt.002	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support providing information about the managed physical resources to consumers (see note 3).
PimPrLcmMgt.003	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support lifecycle management operations related to managed physical resources (see note 1).
PimPrLcmMgt.004	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support modifying information about the managed physical resources (see notes 2 and 3).
PimPrLcmMgt.005	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support managing subscriptions to the notifications related to provisioning and lifecycle management of physical resources.
PimPrLcmMgt.006	The physical resource provisioning and lifecycle management service interface produced by the PIM shall support sending notifications to the subscribed consumers in the event of changes related to provisioning state and lifecycle of managed physical resources.
NOTE 1: The lifecycle management operations include creation, update, querying and deletion of corresponding managed objects related to physical resources that are managed by the PIM.	
NOTE 2: These modifications include, but are not limited to, creation, update, deletion of properties and attributes (e.g. volumes, interfaces, traits, etc.) associated with an individual physical resource during its lifecycle.	
NOTE 3: The properties and other information about physical resources are specified in clause 8.3.5.	
NOTE 4: Provisioning of physical resources includes setting up and provisioning of composed physical resources as described in clause 6.2.	

7.4 Physical resource inventory management service interface requirements

Table 7.4-1 specifies the requirements applicable to the interface of the physical resources inventory management service produced by the PIM.

Table 7.4-1: Physical resource inventory management service interface requirements

Identifier	Requirement
PimPrImMgt.001	The physical resource inventory management service interface produced by the PIM shall support providing inventory information related to physical resources in response to consumer-issued query requests.
PimPrImMgt.002	The physical resource inventory management service interface produced by the PIM shall support sending notifications regarding the changes of inventory information, such as discovery, update and removal of physical resources.
PimPrImMgt.003	The physical resource inventory management service interface produced by the PIM shall support managing subscriptions for notifications related to the inventory management of physical resources.
PimPrImMgt.004	The physical resource inventory management service interface produced by the PIM shall support updating inventory information upon request by a consumer.
NOTE: The physical resource inventory information is specified in clause 8.3.1.	

7.5 Physical resource topology management service interface requirements

Table 7.5-1 specifies the requirements applicable to the interface of the physical resources topology management service produced by the PIM.

Table 7.5-1: Physical resource topology management service interface requirements

Identifier	Requirement
PimPrTmMgt.001	The physical resource topology management service interface produced by the PIM shall support providing topology information related to physical resources in response to consumer-issued query requests (see notes 1 and 3).
PimPrTmMgt.002	The physical resource topology management service interface produced by the PIM shall support sending notifications regarding the changes of topology information, such as the creation, update and removal of topological relationships among physical resources (see notes 1 and 3).
PimPrTmMgt.003	The physical resource topology management service interface produced by the PIM shall support managing subscriptions for notifications related to the changes of topology information of physical resources.
PimPrTmMgt.004	The physical resource topology management service interface produced by the PIM shall support managing logical topology information (see note 2).
NOTE 1: The topology information includes the physical topology information and the logical topology information., as described in clause 6.4 and further specified in clause 8.	
NOTE 2: Management of logical topology information includes the creation, update, deletion of physical resource pools, and association of properties to the physical resource pools for topology.	
NOTE 3: Topology information of physical resources is specified in clause 8 by means of associations between physical resources (see clause 8.3.5), physical resource pools (see clause 8.3.4) and resource zones (see clause 8.3.3).	

7.6 Physical resource performance management service interface requirements

Table 7.6-1 specifies the requirements applicable to the interface of the physical resource performance management service produced by the PIM.

Table 7.6-1: Physical resource performance management service interface requirements

Identifier	Requirement
PimPrPmMgt.001	The physical resource performance management service interface produced by the PIM shall support reporting performance information related to physical resources (see note 2).
PimPrPmMgt.002	The physical resource performance management service interface produced by the PIM shall support sending notifications regarding the availability of performance data, as well as other events related to performance management on physical resources.
PimPrPmMgt.003	The physical resource performance management service interface produced by the PIM shall support controlling the collection and reporting of performance information in response to the request (see note 1).
PimPrPmMgt.004	The physical resource performance management service interface produced by the PIM shall support managing subscriptions to the notifications related to performance management of physical resources.
NOTE 1: The performance management includes creation, update, query and deletion of PM jobs and thresholds, and reporting of performance data as specified in clause 6.5.	
NOTE 2: Performance measurements related to physical resources that are produced by the PIM are specified in clause 7.6 of ETSI GS NFV-IFA 027 [2].	

7.7 Physical resource fault management service interface requirements

Table 7.7-1 specifies the requirements applicable to the interface of the physical resources fault management service produced by the PIM.

Table 7.7-1: Physical resource fault management service interface requirements

Identifier	Requirement
PimPrFmMgt.001	The physical resource fault management service interface produced by the PIM shall support providing alarm information related to physical resources in response to the consumer-issued query requests (see notes 1 and 3).
PimPrFmMgt.002	The physical resource fault management service interface produced by the PIM shall support sending notifications regarding alarm information, as well as other events related to fault management on physical resources, e.g. alarm clearing (see note 1).
PimPrFmMgt.003	The physical resource fault management service interface produced by the PIM shall support managing subscriptions to the notifications related to fault management of physical resources.
PimPrFmMgt.004	The physical resource fault management service interface produced by the PIM shall support managing alarms, such as clear or acknowledge alarms (see note 2).
NOTE 1: The alarm information shall include unambiguous information to distinguish the alarms, e.g. alarm identifier, the physical resources causing the alarm and the alarm cause, as specified in clause 6.5.	
NOTE 2: The relationship between fault management and alarm management is specified in ETSI GS NFV-IFA 045 [3].	
NOTE 3: Valid values of perceived severity, event type, fault type, probable cause and fault details applicable to specific alarms produced by PIM are specified as "Perceived severity", "Event type", "Alarm definition identifier", "Probable cause" and "Fault details" values of the Alarm applicable to a managed physical resource object in clause 7.9 of ETSI GS NFV-IFA 045 [3].	

7.8 Physical resource log management service interface requirements

Table 7.8-1 specifies the requirements applicable to the interface of the physical resource log management service produced by the PIM.

Table 7.8-1: Physical resource log management service interface requirements

Identifier	Requirement
PimPrLmMgt.001	The physical resource log management service interface produced by the PIM shall support sending notifications regarding the availability of logs, as well as other events related to log management on physical resources.
PimPrLmMgt.002	The physical resource log management service interface produced by the PIM shall support managing subscriptions to the notifications related to log management of physical resources.
PimPrLmMgt.003	The physical resource log management service interface produced by the PIM shall support controlling the collection of log information related to physical resources (see note).
PimPrLmMgt.004	The physical resource log management service interface produced by the PIM shall support querying log information of physical resources.
NOTE:	The log management includes controlling logs, i.e. creation, update, query and deletion of logging jobs. Logging jobs can be used for configuring log levels, configuring criteria of what information related to physical resources to log, starting and stopping the log collection, etc. as specified in clause 6.5.

8 Information Modelling for Physical Infrastructure Management

8.1 Introduction

Clause 8 of the present document specifies an abstract information model of NFVI physical resources.

8.2 Relationship between information elements

Figure 8.2-1 illustrates the relationship between different information elements specified in the present document, in the form of a UML class diagram.

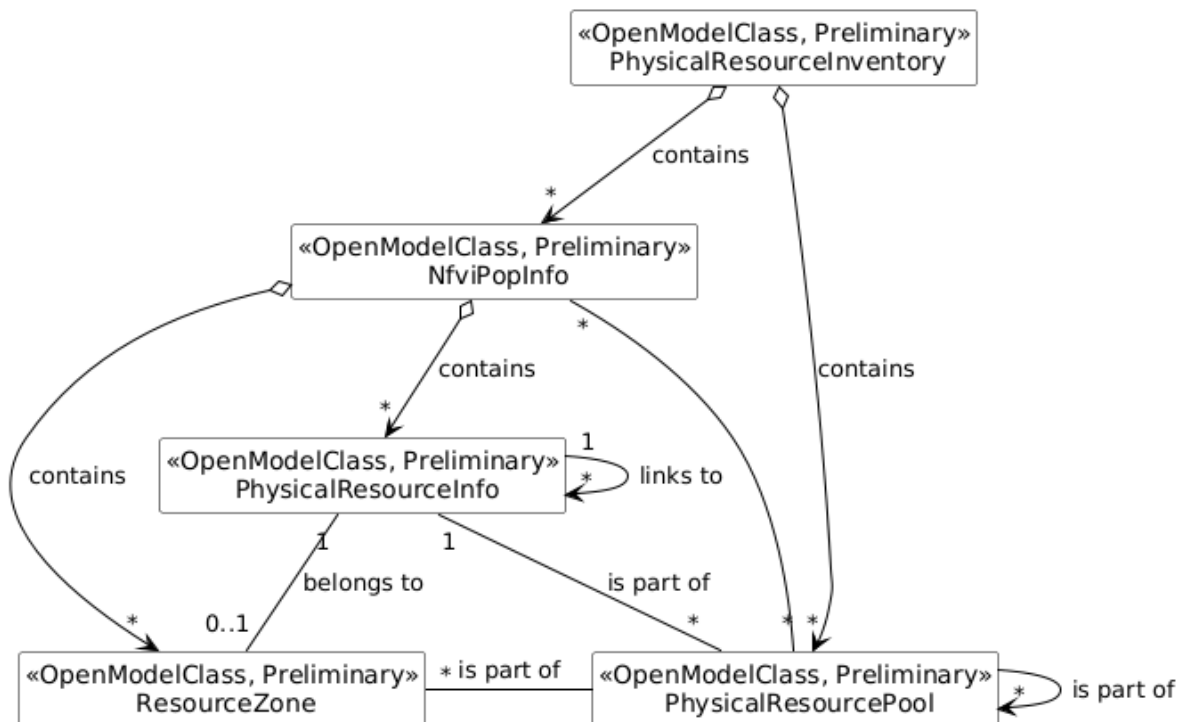


Figure 8.2-1: Relationship between different information elements

8.3 Information elements

8.3.1 PhysicalResourceInventory information element

The PhysicalResourceInventory information element contains information about the overall inventory of physical resources within the NFVI-PoPs and the topology related information of physical resources within the NFVI according to the information regarding resource zones, physical resource pools and the association between physical resources.

The PhysicalResourceInventory shall provide the following information:

- set of NFVI-PoPs conforming the inventoried physical infrastructure (more detailed information provided by NfviPopInfo information element); and
- set of physical resource pools across one or multiple NFVI-PoPs (more detailed information provided by PhysicalResourcePool information element).

8.3.2 NfviPopInfo information element

The NfviPopInfo information element contains information about the physical resources and resource zones contained within the NFVI-PoP.

The NfviPopInfo shall provide the following information:

- identifier of the NFVI-PoP;
- set of physical resources contained in the corresponding NFVI-PoP (more detailed information provided by PhysicalResourceInfo information element); and
- set of resource zones associated to NFVI-PoP (more detailed information provided by ResourceZone information element).

8.3.3 ResourceZone information element

The ResourceZone information element contains information about the physical resources belonging to a resource zone within an NFVI-PoP.

The ResourceZone shall provide the following information:

- identifier of the resource zone;
- description to identify the logical grouping determining the resource zone, e.g. power-source driven, administrative domain driven, etc.; and
- identifiers of physical resource pools within the resource zone.

8.3.4 PhysicalResourcePool information element

The PhysicalResourcePool information element contains information about physical resources belonging to one or multiple NFVI-PoPs that comprise a physical resource pool.

The PhysicalResourcePool shall provide the following information:

- identifier of the physical resource pool;
- description to identify the logical grouping represented by the physical resource pool, e.g. grouping resources of a certain type or with certain properties;
- information about the resource zones that the corresponding physical resource pool spreads across;

- information about the NFVI-PoPs and their resources within the corresponding physical resource pool:
 - identifiers of NFVI-PoPs that the physical resources belong to;
 - identifiers of physical resources; and
- identifiers of the nested resource pools within the corresponding resource pool.

8.3.5 PhysicalResourceInfo information element

The PhysicalResourceInfo information element contains information about an individual physical resource within the NFVI-PoP.

The PhysicalResourceInfo shall provide the following information:

- unique resource identifier for the physical resource (e.g. UUID of the physical resource);
- power state of the resource (e.g. power on, power off, suspended, etc.), in case the physical resource has power control/management capabilities;
- state information of the resource depending on the kind/class of the resource (e.g. provisioning state such as available, configured, active, operational state such as enabled, disabled, administrative state, such as locked, etc.);
- inventory related information about the resource:
 - type (whether it is logical, such as some firmware or software, or physical) and class of resource (compute, storage, network, node racks infrastructure, network fabric, or other);
 - identification information unique to the resource (e.g. a serial number);
 - provider or manufacturer of the resource; and
 - properties of the resource, including information about the capabilities of the resource as a managed object, such as power management capabilities, state management capabilities, etc.
- configuration information depending on the kind/class of the resource (e.g. boot settings for a bare metal server, RAID configurations of a storage disk, etc.);
- specific information about the physical resource depending on the kind/class of the resource, e.g. characteristics related to a compute node, network switch, hypervisor, etc.:
 - software related information of the resource and versioning (e.g. about firmware, hypervisor, OS, supported drivers, etc.);
 - compute related information (e.g. CPU, memory, number of sockets, supported P/C states);
 - storage related information (e.g. information about logical disks, volumes, storage capacity, etc.);
 - network related information (e.g. physical ports, virtual interfaces, MAC addresses, network OS, etc.).
- references to other logical and physical resources determining topological relationship between such resources, i.e. to represent containment (e.g. installation, resource composition) and/or connectivity (e.g. infrastructure network and energy related) topology;
- metadata, providing additional information describing the physical resource.

Annex A (informative): Architectural Analysis of Physical Infrastructure Management

A.1 Overview and background

This annex analyses architectural aspects of the PIM in the NFV-MANO framework building on the preliminary architectural analysis provided in other documents, e.g. ETSI GR NFV-IFA 046 [i.3]. Furthermore, general considerations from architecture design point of view and potential options for placing PIM functionality in the current NFV-MANO architectural framework are evaluated in this annex.

A.2 Placement options for PIM functionality

A.2.1 Introduction

Several solutions considering placement of PIM functionality in NFV are discussed in detail in clause 5.2.1.2.3 of ETSI GR NFV-IFA 046 [i.3]. Detailed analysis using pros and cons-based evaluation is also provided for the documented solutions. Following are the possible solutions identified in ETSI GR NFV-IFA 046 [i.3]:

- SOL-A3-1: Extending the VIM to offer PIM functionality.
- SOL-A3-2: PIM functionality in the OSS.
- SOL-A3-3: PIM functionality outside the OSS and outside NFV-MANO.
- SOL-A3-4: PIM functionality inside NFV-MANO.
- SOL-A3-5: PIM functionality as part of the NFVI.

Based on the analysis and evaluation, only solutions SOL-A3-1 and SOL-A3-4 are recommended in ETSI GR NFV-IFA 046 [i.3]. The two recommended solutions are further analysed in this annex as potential options to place PIM functionality in NFV-MANO framework. Solution SOL-A3-1 is termed as Option #1 and SOL-A3-4 is termed as Option #2 hereinafter.

While the solutions listed above and described in of ETSI GR NFV-IFA 046 [i.3] consider PIM provided as a single function, it can also be distributed into different functional entities which is analysed in Option #3.

A.2.2 Option 1: PIM functionality as part of the VIM

This option proposes extending VIM functionality to cover areas of physical infrastructure management, such as provisioning and monitoring of physical resources. For example, addition of new interfaces exposed by the VIM related to physical infrastructure management.

A.2.3 Option 2: PIM functionality as a new function in NFV-MANO

This option proposes to add the full PIM functionality into one single new function named "PIM" in NFV-MANO, which is responsible for holistic management of NFVI physical resources, such as inventory management, provisioning, monitoring, FM/PM, etc. Under this option, it is assumed that the PIM will be separate from the VIM (see an exemplary illustration of NFV-MANO framework with PIM as a separate functional entity in figure A.3.2-1). The VIM already covers some of the management aspects of physical resources in the NFVI, e.g. compute host reservation, NFVI capacity management, etc. With this option, distribution of responsibilities related to management of NFVI resources between the VIM and the PIM is to be considered.

A.2.4 Option 3: PIM functionality distributed into different entities

Besides the two options in clauses A.2.2 and A.2.3 derived from recommendations in ETSI GR NFV-IFA 046 [i.3], there is another placement option, in which PIM functionality is distributed into different NFV-MANO functional entities.

As mentioned in the scope clause of the present document, PIM functionality mainly includes physical infrastructure management with respect to two parts: provisioning and lifecycle management of compute/storage/networking resources, and FM/PM of compute/storage/networking resources. To simplify the representation in the present document, the former part of PIM functionality is abbreviated as BM-LCM (bare metal lifecycle management), and the latter part of PIM functionality is abbreviated as BM-FM/PM (bare metal fault management and performance management).

This placement option is proposed based on the observation that some industry mainstream solutions on bare metal provisioning (known as part of BM-LCM) adopt solutions like OpenStack® Ironic, for the case of BM deployments. The whole or part of VIM functionality in NFV-MANO can be provided by a set of OpenStack services. The rationale behind this placement option is that OpenStack opensource implementation includes the VIM functionality and part of PIM functionality (BM-LCM) as specified in ETSI NFV standards. Using this principle, this option places BM-LCM in the VIM and BM-FM/PM in a new NFV-MANO functional entity, the PIM. Figure A.2.4-1 illustrates this placement option for distributing PIM functionality into different NFV-MANO functional entities, the VIM and the PIM.

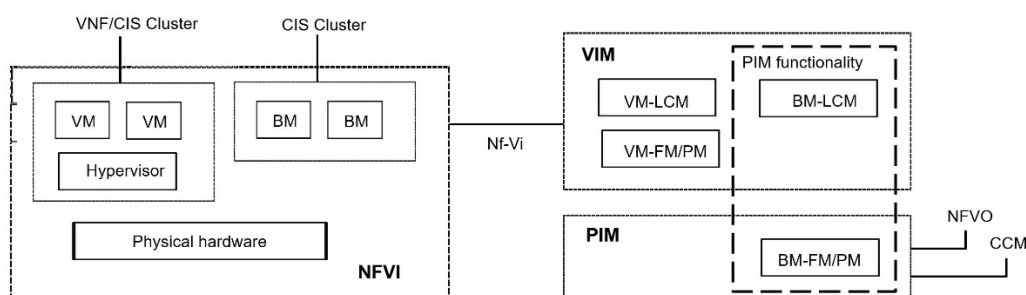


Figure A.2.4-1: PIM functionality distributed in different entities

A.2.5 Analysis

Table A.2.5-1 provides evaluation of placement options for PIM functionality based on the pros and cons documented in ETSI GR NFV-IFA 046 [i.3] as well as some additional pros and cons that aim to cover a broader scope for PIM functionality in NFV.

Table A.2.5-1: Evaluation of PIM functionality placement options

Options	Pros	Cons
Option 1: PIM functionality in the VIM	IFA046: <ul style="list-style-type: none"> • Leverages currently defined physical infrastructure management capabilities already present in current referenced VIM specifications. • No new reference point or interface is defined; resource management operations are over existing communication channels (Nf-Vi towards the NFVI and the VIM to the CCM). • All infrastructure (virtual and physical) management aspects are clearly identified towards a single entity thus enabling consolidation of infrastructure management. • Correlation of physical and virtual infrastructure can be performed without having to involve other management entities. Additional: <ul style="list-style-type: none"> • None. 	IFA046: <ul style="list-style-type: none"> • Additional complexity introduced in the VIM because of extending further beyond its current physical infrastructure management functionality. See note. Additional: <ul style="list-style-type: none"> • No clear delineation of management boundaries for physical and virtualised resources in NFV-MANO as the VIM is responsible for managing both virtual and physical resources. • Limits interactions between the CISM and the PIM related to management of physical resources, for example storage management and provisioning as NFV-MANO does not address interactions between CISM and VIM.
Option 2: PIM functionality as a new function in NFV-MANO	IFA046: <ul style="list-style-type: none"> • Able to handle scenarios with dynamic resource provisioning. • Possible direct mapping of the PIM to O-Cloud IMS operations. • Physical Infrastructure Management information can be used, referenced, and correlated at multiple levels within NFV-MANO, e.g. physical inventory information can be considered for end-to-end resources orchestration by NFV-MANO. Additional: <ul style="list-style-type: none"> • Granular control and monitoring of physical resources in the NFVI to support new NFV use cases, e.g. energy efficiency. • Direct interactions between the PIM and NFV-MANO entities for dynamic provisioning and control of physical resources. 	IFA046: <ul style="list-style-type: none"> • New reference points to be defined between the PIM and the VIM and between the PIM and the CCM. • Resource management requests are possible from two sources either the OSS or NFV-MANO entities, which could cause inconsistencies in state management. Additional: <ul style="list-style-type: none"> • No holistic management of virtualised and physical resource, as offered by a single entity. • Synchronization needed between the PIM and consumers, e.g. the VIM, for correlating physical to virtualised/cluster resources. • Distribution of responsibilities related to management of NFVI resources between the VIM and the PIM needs to be considered.
Option 3: PIM functionality distributed in different entities	<ul style="list-style-type: none"> • Leverages currently defined physical infrastructure management capabilities already present in current referenced VIM specifications. 	<ul style="list-style-type: none"> • Additional complexity introduced in the VIM because of extending further beyond its current physical infrastructure management functionality (e.g. lifecycle management of physical resources). • New reference points (e.g. FCAPS) to be defined between the PIM and the VIM and between the PIM and the CCM. • Synchronization needed between the PIM and consumers, e.g. the VIM, the CCM, for correlating physical to virtualised/CIS cluster resources.
NOTE: ETSI GR NFV-IFA 046 [i.3] also considers some additional cons for this option in the context of O-RAN specifications.		

A.3 Architecture design considerations

A.3.1 Introduction

This clause discusses architecture design considerations for the PIM in terms of its integration in NFV-MANO and analyses different approaches to specifying the PIM and its interfaces within NFV-MANO framework. These architecture design considerations are discussed independent of the placement options described in clause A.2. Nevertheless, the placement options will have an impact on how the PIM interfaces can be specified in NFV-MANO. Relation between the placement options and architecture design approaches is described in clause A.3.4.

A.3.2 Reference point approach

NFV-MANO architectural framework in ETSI GS NFV 006 [i.12] specifies NFV-MANO Functional Blocks (FB), i.e. the NFVO, the VNFM, the VIM, the WIM (as part of NFV-MANO or external to it) and the reference points for interaction between those FBs. To facilitate interactions with other entities, NFV-MANO FBs expose standard interfaces that are supported over relevant reference points. Furthermore, concept of service interfaces has been introduced since Release 4 of ETSI NFV, where multiple functions, especially those offering container support e.g. the CISM, the CCM, expose their service interfaces to other NFV-MANO and non-NFV-MANO entities.

To integrate PIM functionality in NFV-MANO, one option could be to add the PIM as a new functional entity in NFV-MANO framework (as described in clause A.2). It is already established in prior studies, e.g. ETSI GR NFV-IFA 046 [i.3], ETSI GR NFV-EVE 021 [i.2], that the PIM can interact with multiple NFV-MANO entities to support new features and use cases summarized in clause 4 of the present document. For example, interactions between the PIM and the VIM, the PIM and the CCM, the PIM and the NFVO are already identified to fulfil physical infrastructure management in NFV. Therefore, using the reference point approach, multiple reference points could be established within NFV-MANO between the PIM and other NFV-MANO entities. An exemplary illustration of NFV-MANO framework with the PIM as a new functional entity is shown in figure A.3.2-1.

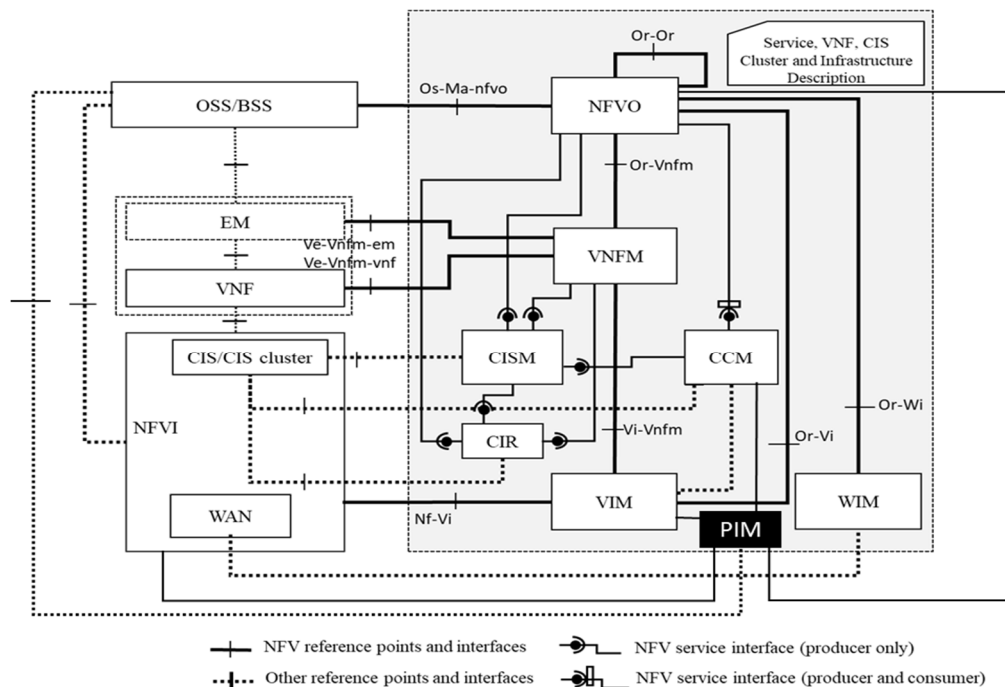


Figure A.3.2-1: The PIM function inside NFV-MANO as an example placement option

Figure A.3.2-1 is based on the one of the architectural placement options for the PIM discussed in ETSI GR NFV-IFA 046 [i.3]. The figure shows multiple points of interaction between the PIM and the NFVO, the PIM and the VIM, the PIM and the CCM, etc., over which the PIM could expose its interfaces.

A.3.3 SBA Approach

ETSI GR NFV-IFA 039 [i.13] outlines SBA principles for NFV-MANO that enable the NFV-MANO services to be defined in a consumer-independent fashion. By decoupling the interfaces and their functionalities from specific reference points, SBA, in its pure form, also enables flexible placement of NFV-MANO services being offered by the service producers. These services are exposed by the producer over one or multiple service interfaces independently of how the consumers intend to use those services. SBA approach can be considered, if NFV-MANO architecture implements one of SBA transformation targets, according to recommendations described in ETSI GR NFV-IFA 039 [i.13].

NOTE: The approach described in this clause is based on the general SBA principles for NFV described in ETSI GR NFV-IFA 039 [i.13]. A specific approach can be applied using one of the SBA transformation targets described in the same document.

Using the SBA approach, PIM functionality can be provided by a logical function (as service producer) offering different physical infrastructure management services via service interfaces. These services can be consumed by service consumers, such as other services offered by NFV-MANO entities, such as the NFVO, the VIM, the CISM, etc., as well as non-NFV-MANO entities, e.g. the OSS/BSS, etc. Figure A.3.3-1 illustrates an example of this concept.

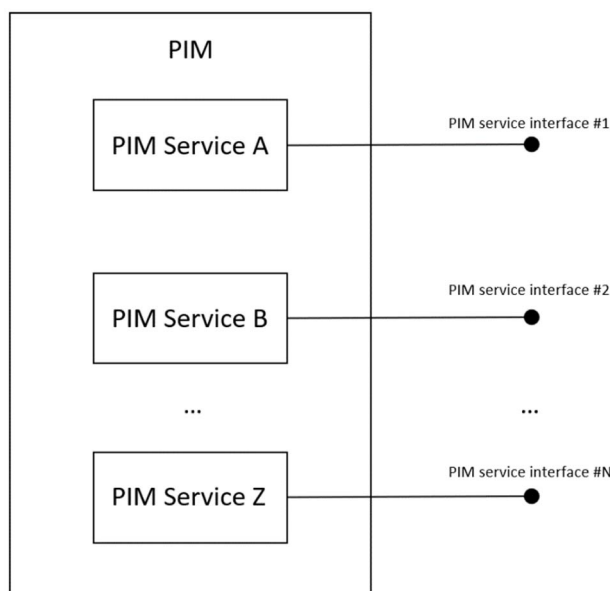


Figure A.3.3-1: Example of PIM services and service interfaces

Specifying PIM services according to the principles of SBA provides flexible placement options for PIM functionality in NFV-MANO framework as discussed in clause A.2.

A.3.4 Analysis

Pros and cons-based analysis of the two approaches to PIM architecture design is summarized in table A.3.4-1. Relationship between architecture design approaches and placement options for PIM functionality is also highlighted.

NOTE: The analysis provided in table A.3.4-1 does not analyse the overall SBA approach for architecture design. The analysis for applying SBA to the whole NFV-MANO architecture is already documented in ETSI GR NFV-IFA 039 [i.13]. Analysis in table A.3.4-1 is limited to design of the PIM services and service interfaces.

Table A.3.4-1: Evaluation of architecture design approaches for the PIM

Approach	Pros	Cons	Relation with placement options
Reference Point approach	Uniformity with the way existing interfaces exposed by the VIM and other NFV-MANO FBs are defined in NFV-MANO.	A set of same interfaces exposed by the PIM could be duplicated over multiple reference points in NFV-MANO. From specification point of view, interface specifications can be duplicated in multiple documents, increasing complexity for future enhancements in the specs. See note.	<p>Option #1 PIM functionality in the VIM:</p> <ul style="list-style-type: none"> Suitable for this placement option. New interfaces can be added as VIM exposed interfaces over the relevant reference points. <p>Option #2 PIM functionality as a new function in NFV-MANO:</p> <ul style="list-style-type: none"> Suitable for this placement option. New reference points can be defined between the PIM and other NFV-MANO entities, specifying interfaces supported over each reference point. <p>Option #3 PIM functionality distributed into different entities:</p> <ul style="list-style-type: none"> Suitable for this placement option. Existing reference points related to the VIM can be enhanced to add new interfaces if required. New reference points can be introduced between the PIM and other NFV-MANO entities, specifying interfaces supported over each reference point.
SBA approach	Using SBA approach, the PIM service interfaces can be specified in a consumer-agnostic manner.	This approach can be considered inconsistent with how the interfaces related to resource (both virtual and physical) management are specified in NFV-MANO, e.g. the VIM exposed interfaces.	<p>Option #1 PIM functionality in the VIM:</p> <ul style="list-style-type: none"> From specifications point of view, there may be a mismatch between the approaches as VIM interfaces specifications are not specified using the SBA approach. From implementation standpoint, the discussion is irrelevant. An implementation of the VIM, the NFV-MANO functional block, can choose to offer PIM services according to NFV-MANO PIM service interface specifications (see clause A.4). <p>Option #2 PIM functionality as a new function in NFV-MANO:</p> <ul style="list-style-type: none"> Suitable. New PIM function in NFV-MANO can offer its functionalities using service interfaces, similar to other NFV-MANO functions, e.g. the CISM, the CCM, the MDAF, etc. <p>Option #3 PIM functionality distributed into different entities:</p> <ul style="list-style-type: none"> Same as 'Option #1 PIM functionality in the VIM'.
NOTE: To avoid this con, all interfaces related to the PIM can be specified in one place and the relevant interfaces can be referenced elsewhere if needed.			

A.4 Conclusion

This annex concludes the analysis related to the placement of PIM functionality within NFV-MANO as well as the interface design analysis for the PIM exposed interfaces.

It is concluded to follow the 'Option 2: PIM in NFV-MANO' principle described in this annex. Using this principle, the PIM is specified as a logical function within the NFV-MANO framework in clause 5 of the present document. This placement is purely from a logical and architectural point of view.

It is concluded to follow the 'SBA approach' described in this annex, i.e. the functionality offered by the PIM is specified in the form of services, e.g. physical resource provisioning and lifecycle management service, physical resource inventory management service, etc. These services are offered by the PIM via service interfaces. Reference points between the PIM and any other NFV-MANO entities have not been identified or specified in the present document.

NOTE: If some reference points need to be specified for interactions between the PIM and other NFV-MANO entities, those reference points can reuse relevant interface specifications specified in the present document.

Annex B (informative): Analysis of existing solutions

B.1 Introduction

This annex provides capability analysis of existing industry solutions related to physical infrastructure management.

B.2 OpenStack® Ironic

B.2.1 Overview

The OpenStack® Bare Metal service, also known as Ironic, supports provisioning and management of physical machines and servers. The service can be used independently or within an OpenStack-based cloud environment. If integrated within an OpenStack cloud, the Bare Metal service offers the OpenStack Compute service a uniform way to manage heterogeneous physical machines and servers by abstracting hardware-specific drivers.

One of the components of the Bare Metal service, 'ironic-conductor', manages hardware nodes by interacting with respective device drivers. Ironic conductor instances can provision physical nodes using the Intelligent Platform Management Interface (IPMI) or vendor-specific protocols. Another component of the Bare Metal service, 'ironic-api', is used to enrol hardware which is to be managed by the service. The RESTful API also exposes a list of ironic-conductor instances and corresponding drivers supported by them.

NOTE: Details about concepts, architecture, components of OpenStack's Bare Metal service are available in the official documentation of OpenStack® Ironic [i.14].

B.2.2 Features

The features and capabilities of OpenStack Bare Metal service for provisioning and management of bare metal nodes can be summarized as below:

- **Resource provisioning/deprovisioning:** Provisioning of bare metal resources, i.e. nodes, can include the following: preparing nodes (reconfiguring, etc.) for enrolment, configuring drivers, boot and network settings for the nodes, adding scheduling information (e.g. resource class, traits, etc.) and finally enrolling the nodes with the Bare Metal service so that the newly provisioned nodes are available for deployment of workload(s). Once the nodes are 'available' for deployment, user images can be deployed. The Bare Metal service also offers cleaning and tearing down of nodes (removing volumes, configurations, etc.) as part of the deprovisioning process.
- **Nodes management:** Nodes management refers to creating, updating, and deleting ironic nodes. This can be performed via the Bare Metal API. Additionally, the Bare Metal service offers management tasks associated with individual nodes, such as, setting maintenance flag for a node, setting boot device, changing node's boot mode/power state/provision state, etc.
- **Network management:** The Bare Metal service offers management of node's network resources, such as Network Interface Cards (NICs) via the use of ports (physical level ports with MAC addresses) and port groups. Attaching/detaching Virtual InterFaces (VIFs) to/from a node is also possible via the Bare Metal API.
- **Storage management:** The Bare Metal service offers management of storage resources for ironic nodes, such as managing node's logical disks, RAID configurations, storage interfaces, etc. There can also be external resources, e.g. from a Block Storage service, such as OpenStack Cinder. Remote volumes can be connected to Ironic nodes as either targets or connectors via the Bare Metal API.
- **Allocation and scheduling management:** The Bare Metal service can make use of resource classes, traits and other parameters to maintain logical grouping of nodes, using the 'allocations' resource of the Bare Metal API. Allocations can be used to identify the (set of) node(s) for customized scheduling.

- **Notifications:** The Bare Metal service can also generate notifications over a message bus to indicate various events, such as, change in power state of a node. These notifications can be consumed by external entities like a monitoring service. Notifications can come from the Bare Metal API (ironic-api) as a result of CRUD operations, or from ironic-conductor(s) as a result of state changes in the nodes, e.g. during transitions in power states/provision states.
- **Hardware inspection:** The Bare Metal service supports discovery of inventory, features and capabilities of hardware nodes. This information is collected by the inspecting agents and processes deployed on managed nodes and sent to the Bare Metal service.
- **Metric collection:** The Bare Metal service supports collection of hardware metrics exposed by respective hardware managers/drivers. For example, the Redfish driver supports sending hardware metrics to the Bare Metal service which can be exposed by the Bare Metal service as notifications. Bare Metal service also exposes IPMI metrics as notifications which can be consumed by other OpenStack services, like Telemetry (OpenStack project name: Ceilometer).
- **Logging:** Like other OpenStack services, the Bare Metal service generates logs pertaining to various severity levels and enables collection of logs related to bare metal deployments. Clients can configure different parameters like when to collect the logs, location for log collection, expiration timer for stored logs, etc. In addition to deployment related logs, the Bare Metal service generates other logs like ironic-conductor logs.
- **Fault management:** Faults detected by Ironic in the bare metal nodes can be retrieved in the 'Node' resource managed by the Bare Metal service. The Node resource also contains information about the last error that occurred in a node.

NOTE 1: Details about features and capabilities of OpenStack's Bare Metal service are available in the official documentation of OpenStack Ironic [i.14].

NOTE 2: Details about API versions and endpoints of Bare Metal API are available in the official reference documentation of the Bare Metal API [i.15].

NOTE 3: It is unclear whether fault related events that occur in bare metal nodes are exposed as notifications by the Bare Metal service.

B.3 Metal3

B.3.1 Overview

The Metal3 (pronounced as Metal-Kubed) project is an open-source project that provides a set of tools for managing bare-metal infrastructure used by the Kubernetes [i.17] cluster. Metal3 is itself a Kubernetes application which runs on the Kubernetes cluster, and uses Kubernetes resources and APIs as its interface.

As shown in figure B.3.1-1, Metal3 is one of the providers for the Kubernetes [i.17] sub-project Cluster API [i.17]. Cluster API provides infrastructure agnostic Kubernetes cluster lifecycle management, and Metal3 brings the bare metal implementation capabilities for it.

Metal3 is integrated with the components of Ironic [i.14] from the OpenStack ecosystem, for booting and installing bare metal machines. Metal3 handles the installation of Ironic as a standalone component, which means there is no need to bring along the rest components of OpenStack other than Ironic.




Figure B.3.1-1: Metal3 as a provider for the Kubernetes [i.17] sub-project Cluster API

NOTE: Details about concepts, architecture, components of Metal3 are available in the official documentation of Metal3 [i.16].

B.3.2 Features

The features and capabilities of Metal3 for provisioning and management of bare metal nodes are composed of three main components as shown in figure B.3.1-1:

- Cluster API Provider Metal3, which enables the management of physical servers using the Kubernetes [i.17] sub-project Cluster API.

The Cluster API [i.17] itself is shared across multiple cloud infrastructure providers. Cluster API Provider Metal3 is one of the providers for Cluster API and enables Cluster API consumers to deploy a Cluster API based cluster on top of bare metal infrastructure using Metal3.

- Bare Metal Operator, which automates the provisioning and management of bare-metal servers using the open-source OpenStack Ironic project [i.14].

The Bare Metal Operator is a custom Kubernetes [i.17] controller that deploys bare metal host servers using OpenStack Ironic. The Bare Metal Operator is responsible for the following:

- Inspect the host's hardware details and report them on the corresponding BareMetalHost object, which includes information about CPUs, RAM, disks, NICs, etc.
- Provision hosts with a desired image.
- Clean a host's disk contents before or after provisioning.

The managed resource object of Bare Metal Operator is represented in Kubernetes [i.17] as BareMetalHost. The BareMetalHost object contains information about the server as shown below.

- Known server properties: Fields such as *bootMACAddress* of the server and are known in advance for Metal3 consumers to use this infrastructure.
- Unknown server properties: Fields such as CPU and disk are properties of the server and are discovered by Ironic.
- Consumer supplied fields: Fields such as image are supplied by Metal3 consumers to dictate boot image for the server.
- Dynamic fields: Fields such as IP are dynamically assigned to the server at run time.
- Ironic, the core component in Metal3, which is responsible for the provisioning and management of bare metal servers.

The features and capabilities of Ironic are summarized in clause B.2.2 in the present document.

NOTE 1: The Metal3 project [i.16] does not provide bare metal infrastructure fault management and performance management capabilities.

NOTE 2: For provisioning and management of new kind of NFVI resources, to assess whether Ironic integrated with the other OpenStack components or the standalone Ironic service (used in Metal3) have the same capabilities is not in the scope of the present document.

B.4 Redfish®

B.4.1 Overview

DMTF® DSP0266 [i.9] specifies a HTTP RESTful interface that a remote client can use to manage a conformant Redfish service. The interface specification defines the RESTful behaviour of the resources. The model specification defines the structure of the HTTP resources.

Redfish's model expands as other standards bodies seek to use Redfish® for their management domains. DMTF® has received extension requests for storage services, fabric management, DataCentre Infrastructure Management (DCIM), industrial IoT, and process automation organizations.

Redfish® enables management domains to specify a subset of the Redfish® model that needs to be implemented for their management domain. This prescription for conformance is specified in a machine-readable file, a profile. The DMTF® supports an open-source validator [i.20] which can be used to test an implementation for conformance to a profile. There are profiles owned by the Open Compute Project [i.21], the Open Process Automation Forum [i.22] and OpenStack.

Both OpenStack and Metal3 can manage platforms which expose Redfish® via the Sushy repository [i.23].

B.4.2 Services and Models

Redfish® provides mechanisms to manage various aspects for manageability. These mechanisms are exposed in the general resource model or encapsulated in a service resource:

- **Logical, Physical, Manager model:** The Redfish® model contains a logical model, a physical model and a manager model. The latter two models model the physical and management hierarchies, respectively. The logical model (or functional model) covers compute, storage and fabric functions. An example of logical model is collection of hardware components that make up 'Systems' in Redfish, e.g. processors, disks, NICs, memory components making up the 'ComputerSystem' collection. The physical model represents the physical view of 'Systems', e.g. (collection of) resources contained in a 'Chassis'. The manager model contains collection of managers (e.g. BMCs) that are managing available (collection of) resources contained within the 'Chassis', making up 'Systems'.
- **Fabrics model:** The fabric model enables the user to manage various fabrics using the same management pattern. This includes PCIe, CXL, Ethernet, SANs. DMTF® works with the OpenFabric Alliance on this model, which is described in Redfish Fabrics White Paper [i.24].
- **OS, VMs and Containers model:** The OS/VM/Container model was created from a contribution by the Open Process Automation Forum for managing software entities above the hardware platform. Currently, the model provides only basic manageability.
- **Telemetry Service:** The telemetry service, as described in Redfish Telemetry White Paper [i.25] provides a mechanism to describe metrics, specify metric reports and triggers.
- **Composition Service:** The composition service enables a user to request a composition be created from the resources managed by a Redfish® service. The requests can be made explicitly, by specifying constraints or specifying intent. The composition service is described in Redfish Composability White Paper [i.26].
- **Firmware Update Service:** The firmware update service supports updating, by pushing or pulling, one or more firmware images. More information on firmware update service and data model is available in Redfish Firmware Update White Paper [i.27].

- **Certificate Service:** The certificate management service describes certificates for devices, services, and other resources. It also provides interface for clients to manage certificates. More information on the Redfish certificate data model as well as the common workflows to manage certificates is provided in Redfish Certificate Management White Paper [i.28].
- **Aggregation Service:** The aggregation service provides a mechanism for performing actions on an ad-hoc or persistent group of resources. Access points define the properties needed to access the entity being aggregated and connection methods describe the protocol or other semantics of the connection.
- **Account, Session, Event, Task, Job Services:** These services are specified to support the remote client. The account service is the basis for providing privilege. The session service shows each session connected to the Redfish® service. The event service enables a user to subscribe to asynchronous messages. The task service enables the user to track asynchronous actions. The job service enables the user to schedule and execute operations as jobs. Actions supported by different Redfish resources can be scheduled and executed in particular order using a job, or collection of jobs.
- **Message Registries:** Messages emitted by a Redfish® service are specified in Message Registries. Details for each message registry and their messages are available in the Redfish Message Registry Guide [i.29]. Each entry specifies the string and fields of the message.
- **Discovery:** Redfish supports discovery of managed devices that support Redfish via Simple Service Discovery Protocol (SSDP). Changes in the status of Redfish managed devices like related to availability can also be configured to be discovered by the Redfish service.

Some relevant Redfish components and services are detailed further in subsequent clauses.

B.4.3 Implementations

DMTF® supports an open-source repository of tooling to enable Redfish® client development, which is available at DMTF Redfish® Developers Hub [i.19]. Open-source implementations of the Redfish® service are available from several sources. Table B.4.3-1 shows a few of these open-source implementations. Some repositories have tooling that generates code for Redfish® resources from a mockup or schema file.

Table B.4.3-1: Open-Source implementations of a Redfish® service

Implementation	Language	Code-generator uses	Repo owner
OpenBMC	C++ (firmware)		Linux® Foundation
Sunfish	Python (SW agent)	mockup, CSDL/JSON-schema	OpenFabric Alliance
Redfish® Service Framework	Java™ (SW agent)	mockup, OpenAPI	Arizona State University
PSME	C++ (SW agent)		OCP
Device Manager	GoLang (SW agent)		OCP

B.4.4 Provisioning and lifecycle management of hardware resources

Redfish supports managing different kinds of hardware resources via its RESTful interface, which is specified in the Redfish specifications [i.9]. The Redfish data model, specified in [i.10], describes and models various kinds of hardware systems as Redfish resources, which can be managed by performing CRUD operations via the Redfish interface. Clients can invoke different HTTP requests (POST, GET, PUT, PATCH and DELETE) towards the RESTful interface to manage the lifecycle of Redfish resources, e.g. ComputerSystem.

The Redfish interface also enables configuring or updating a particular Redfish resource directly via appropriate HTTP methods (e.g. PUT, PATCH). Furthermore, some Redfish resources support management 'actions' that can be invoked via the Redfish interface by making appropriate POST requests towards the desired actions supported by those resources. For example, the ComputerSystem resource supports the 'Reset' action, the Bios resource supports the 'ResetBios' action.

Other solutions for provisioning and management of infrastructure for cloud deployments like OpenStack Ironic and Metal3 can also make use of Redfish service for managing underlying hardware. These solutions can interact with the Redfish service using the Redfish RESTful interface for managing relevant Redfish resources. For example, OpenStack Ironic makes use of the Sushy library [i.23] for communicating with the Redfish service for hardware management.

B.4.5 Composition Service

The Redfish composition service, as described in detail in the Redfish Composability White Paper [i.26], enables clients to create and manage composed logical systems comprising of disaggregated hardware resources. Redfish via its data model enables describing composable hardware, as well as manage such compositions using the Redfish interface.

The Redfish composition service makes use of 'resource blocks' and 'resource zones' to inventory the overall hardware resources that can be made part of different compositions. The service describes 'active' and 'free' pools to differentiate between resources that are already part of compositions and resources that are available for new compositions respectively.

Resource blocks contain resources that can be used as building blocks to form different compositions. A resource block can be comprised of different kinds of hardware resources and their combinations. For example, a resource block with resource block type 'storage' can contain one or more drives, another resource block containing CPU-memory complex and a network controller can be of type 'compute' and 'network'. A resource block of type 'ComputerSystem' can contain one or multiple whole computer systems. These resources in the resource blocks are consumed only via compositions created and managed by the Redfish composition service. Resource zones are used to enable restrictions during composition formations. A resource zone can contain multiple resource blocks that can be part of same compositions.

The two main types of composition approaches supported by the Redfish composition service are specific composition and constrained composition. Specific composition enables clients to compose and manage logical systems using specific resource blocks and resource zones. Constrained compositions delegate composition of systems to the Redfish service itself and clients can just provide their composition requests in a declarative fashion. The Redfish composition service assigns requested resources from appropriate resource zones and resource blocks to compose and provision desired composed system.

B.4.6 Telemetry Service

The Redfish telemetry service, described in Redfish Telemetry White Paper [i.25], enables Redfish clients to obtain information about different metrics related to managed resources, e.g. performance metrics related to a processor, power consumption metrics for a chassis, etc. Different kinds of metrics can be defined along with their several characteristics using 'metric definitions' in the Redfish telemetry model.

Redfish telemetry service also supports the ability to generate metric reports which can be sent as notifications via Redfish event service, logged using the Redfish log service or both. The metric reports can be specified using 'metric report definitions' in the Redfish telemetry model.

Telemetry service also enables specifying triggers or thresholds associated to metric properties, which can result in one or more actions. For example, a trigger can be sent as an alert via the event service and/or logged into the logs by the log service.

B.4.7 Log Service

Redfish offers log service to capture logs associated with Redfish resources (e.g. Chassis), devices (e.g. PCIe device), resource blocks (i.e. collection of resources), models (e.g. Manager), systems and their components (e.g. memory device of a Computer System) and services (e.g. telemetry service, job service), etc.

Via the use of LogService schema in the Redfish data model specified in [i.10], log service enables collecting log entries as well as properties related to log collection e.g. time duration of log collection, size of logs, logging purpose, etc. The format of log entries collected by the log service is defined by the LogEntry schema in the Redfish data model.

B.4.8 Event Service

Redfish's event service is used for managing event subscription and sending notifications to event destinations, i.e. subscribers related to subscribed events. Subscriptions can be created for events related to Redfish resources and/or events collected in Redfish message registries. There are some message registries dedicated to events related to services such as log, telemetry, jobs. Several alarms related messages are also produced by the Redfish service and are available in relevant message registries. Subscriptions can be made on the basis of message registries to get only the messages from select registries.

B.5 Analysis

Table B.5-1 provides a high-level mapping between management services offered by the PIM and the features and capabilities of different solutions discussed in this annex. Detailed comparison of PIM functional requirements and service interface specifications with capabilities of relevant open-source solutions and de-facto standards is left for Stage 3 specifications.

Table B.5-1: Mapping of PIM management services with solution features and capabilities

PIM Management Services	Relevant features and capabilities of solutions		
	OpenStack Ironic	Metal3	Redfish
Physical resource provisioning and lifecycle management service	<ul style="list-style-type: none"> Node provisioning and lifecycle management of bare metal nodes Configuration management and composition management of bare metal nodes, e.g. managing network/storage attachments of bare metal nodes Notifications related to lifecycle state changes in bare metal nodes 	<ul style="list-style-type: none"> Provisioning, lifecycle and configuration management of bare metal hosts (see note 1) 	<ul style="list-style-type: none"> Lifecycle and configuration management of Redfish resources via Redfish interface Provisioning of composed systems via Composition service Notifications related to Redfish resources generated by the Event service
Physical resource inventory management service	<ul style="list-style-type: none"> Discovery of physical node's hardware resources and capabilities via hardware inspection Maintaining inventory information of nodes 	<ul style="list-style-type: none"> Discovery of hardware capabilities and inventory information of bare metal hosts (see note 1) 	<ul style="list-style-type: none"> Discovery of available hardware devices that can be managed by Redfish service Comprehensive schema to represent managed hardware devices using different models, i.e. logical, physical and manager model

PIM Management Services	Relevant features and capabilities of solutions		
	OpenStack Ironic	Metal3	Redfish
Physical resource topology management service	<ul style="list-style-type: none"> Node allocations for creating logical grouping of nodes based on parameters such as resource classes, traits and labels Maintaining information about linkage, i.e. connection/containment relationship between different physical resources in a Node, e.g. compute, networking, storage resources in a node 	<ul style="list-style-type: none"> Support provided by Cluster API Provider Metal3 (CAPM3) for creation of 'Machine' and 'MachinePool' objects of ClusterAPI [i.17] (see note 2) 	<ul style="list-style-type: none"> Support for 'resource zones' and 'resource blocks' for creating pools and topological relationships between hardware resources Comprehensive schema to represent managed hardware devices using different models, i.e. logical, physical and manager model Maintaining information about linkage, i.e. connection/containment relationship between different Redfish resources, e.g. logical systems (such as ComputerSystem) contained within physical systems (such as Chassis), compute, networking, storage resources making up ComputerSystem (collections), etc.
Physical resource performance management service	<ul style="list-style-type: none"> Hardware metrics collection from Redfish driver Exposing metrics collected from host machine's IPMI sensors and meters (e.g. power, temperature, CPU utilization, etc.) 	<ul style="list-style-type: none"> No direct correspondence (see note 3) 	<ul style="list-style-type: none"> Performance metrics related to different Redfish-managed resources via Telemetry service Notifications related to availability of metric reports
Physical resource fault management service	<ul style="list-style-type: none"> Maintaining information about faults and errors linked to bare metal nodes 	<ul style="list-style-type: none"> No direct correspondence (see notes 3 and 4) 	<ul style="list-style-type: none"> Alarms related messages produced by the Redfish Event service, stored in relevant message registries Notifications related to alarms against subscriptions based on select message registries
Physical resource log management service	<ul style="list-style-type: none"> Different logs generated and collected by the Bare Metal service 	<ul style="list-style-type: none"> No direct correspondence (see note 3) 	<ul style="list-style-type: none"> Logs generated by the Redfish Log Service associated with different Redfish resources Log collection control offered by the Log service, e.g. time duration of log collection, size of logs, logging purpose, etc.
NOTE 1: Metal3 uses Ironic for these features.			
NOTE 2: These logical groupings related to machines and machine pools are provided by the Kubernetes [i.17] Cluster API and not Metal3 itself.			
NOTE 3: Metal3 might be able to use the relevant information collected by Ironic. However, it is unclear from Metal3 documentation.			
NOTE 4: The CAPM3 might provide fault and failure related information to Cluster API to be included in the 'status' object of Cluster API's 'Machine' resource.			

Annex C (informative): PlantUML source code for the information model

C.1 Source code

```

@startuml
hide circle
hide methods
hide members

class "<<OpenModelClass, Preliminary>>\n PhysicalResourceInventory" as PhysicalResourceInventory{}
class "<<OpenModelClass, Preliminary>>\n NfviPopInfo" as NfviPopInfo{}
class "<<OpenModelClass, Preliminary>>\n PhysicalResourcePool" as PhysicalResourcePool{}
class "<<OpenModelClass, Preliminary>>\n ResourceZone" as ResourceZone{}
class "<<OpenModelClass, Preliminary>>\n PhysicalResourceInfo" as PhysicalResourceInfo{}

skinparam class {
    AttributeIconSize 0
    BackgroundColor white
    BorderColor black
    ArrowColor black
}
skinparam Shadowing false
skinparam Monochrome true
skinparam ClassBackgroundColor White
skinparam NoteBackgroundColor White

'skinparam linetype polyline
'skinparam linetype ortho

PhysicalResourceInventory o--> "*" NfviPopInfo: contains
PhysicalResourceInventory o--> "*" PhysicalResourcePool: contains

NfviPopInfo o--> "*" PhysicalResourceInfo: contains
NfviPopInfo o--> "*" ResourceZone: contains
NfviPopInfo "*" -- "*" PhysicalResourcePool

PhysicalResourceInfo "1" --> "*" PhysicalResourceInfo: links to
PhysicalResourceInfo "1" -- "0..1" ResourceZone: belongs to
PhysicalResourceInfo "1" -- "*" PhysicalResourcePool: is part of

PhysicalResourcePool --> "*" PhysicalResourcePool: is part of
PhysicalResourcePool -left- "*" ResourceZone: is part of

@enduml

```

Annex D (informative): Change history

Date	Version	Information about changes
May 2023	V0.0.1	Initial skeleton of the GS as provided in NFVIFA(23)000313r1 (approved in NFVIFA#332-F2F).
May 2023	V0.0.2	Updates the Scope clause by implementing contribution: NFVIFA(23)000352r1 - IFA053 Scope (approved in NFVIFA#335).
June 2023	V0.1.0	First major drafting implementing technical contributions: <ul style="list-style-type: none"> NFVIFA(23)000421r1 - IFA053 Clause 4 Introduction NFVIFA(23)000449r2 - IFA053 Clause 4 PIM for Green NFV
September 2023	V0.2.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(23)000539r1 - IFA053 Clause 4 PIM for CIS cluster management NFVIFA(23)000621r2 - IFA053 3.1 General terminologies
September 2023	V0.3.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(23)000628r3 - IFA053 Annex A PIM Placement Options NFVIFA(23)000644r1 - IFA053 Clause 4 PIM for NFV connectivity NFVIFA(23)000654r1 - IFA053 Clause 4.3 Input from other organizations NFVIFA(23)000655r1 - FEATxx IFA053 Role of PIM in vRAN NFVIFA(23)000659r2 - IFA053 Annex A Interface Design - conventional approach NFVIFA(23)000661r2 - IFA053 Annex A Interface Design Analysis
November 2023	V0.4.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(23)000660r2 - IFA053 Annex A Interface design SBA approach NFVIFA(23)000681r3 - IFA053 Annex A.2.3 Option of distributing PIM functionality into entities Rapporteur actions: <ul style="list-style-type: none"> Corrected typo 'Manger' and corrected in-text reference to figure A.2.4-1 in NFVIFA(23)000681r3 Use correct clause numbers while referring to Options#1 and Options#2 in NFVIFA(23)000681r3 Update subclause numbers after A.2.4 Corrected some old numbering mistakes in A.2.5
December 2023	V0.5.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(23)000704r2 - FEAT33 IFA053 Clause 6 Introduction NFVIFA(23)000728 - FEAT33 IFA053 Annex B OpenStack Ironic Overview
January 2024	V0.6.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(23)000823r1 - FEAT33 IFA053 Annex B OpenStack Ironic Features NFVIFA(23)000798 - FEAT33 IFA053 Annex B Metal3 Overview NFVIFA(23)000770r3 - FEAT33 FEAT29 IFA053 Resource Provisioning and LCM NFVIFA(23)000745r1 - FEAT33 FEAT29 IFA053 Clause 6 Inventory NFVIFA(23)000746r1 - FEAT33 FEAT29 IFA053 Clause 6 Topology NFVIFA(23)000747 - FEAT33 FEAT29 IFA053 Clause 6 Monitoring Rapporteur actions: <ul style="list-style-type: none"> Added a heading B.3 for Metal3 analysis in Annex B to accommodate content of contribution #798 Minor editorial: to add "of" in the last paragraph of contribution #798 Improve description of the link provided for metal3 in informative references Aligned the description of bullet items in clause 6.1 according to the order and names of following clauses Deleted the EN in contribution #770r3 after implementing the concerned action
February 2024	V0.7.0	Implements following contributions: <ul style="list-style-type: none"> NFVIFA(24)000003r2 - FEAT33 FEAT29 IFA053 PIM Annex B Metal3 Features NFVIFA(24)000005r2 - FEAT33 IFA053 Clause 1 Scope update NFVIFA(24)000083r1 - IFA053 A.2.5 Analysis on Option 3 NFVIFA(24)000098r1 - IFA053 6.5 Add physical resource performance management service interface requirements

Date	Version	Information about changes
March 2024	V0.8.0	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000182 - FEAT33 IFA053 Clause 7 General service requirements • NFVIFA(24)000124r1 - FEAT33 IFA053 Clause 7 PIM LCM service interface requirements • NFVIFA(24)000121r2 - IFA053 Add physical resources inventory management service interface requirements • NFVIFA(24)000122r2 - IFA053 Add physical resources topology management service interface requirements • NFVIFA(24)000120r3 - IFA053 Add physical resource fault management service interface requirements • NFVIFA(24)000118r1 - IFA053 6.5 Add physical resources log management service interface requirements • NFVIFA(24)000105r4 - Redfish Text for IFA053 • NFVIFA(24)000145r1 - FEAT33 IFA053 Annex A.3 SBA Analysis • NFVIFA(24)000137r2 - FEAT33 IFA053 Clause 5 Introduction of PIM logical function Rapporteur actions: <ul style="list-style-type: none"> - Ordered clause 7 service interface specifications according to their order in table 7.2-1. - Deleted placeholder editor's note in clause 4.3.3 for Redfish text.
May 2024	V0.9.0	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000218r1 - FEAT33 IFA053 Clause 8 Information model diagram • NFVIFA(24)000190 - FEAT33 IFA053 Clause 8 Multiple information elements • NFVIFA(24)000191r1 - FEAT33 IFA053 Clause 8 PhysicalResourceInfo information element • NFVIFA(24)000010r6 - IFA053 Alignment of terms on PIM • NFVIFA(24)000244 - IFA053 7.1 Introduction
May 2024	V0.10.0	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000295r1 - FEAT33 IFA053 Annex A Conclusion • NFVIFA(24)000296r1 - FEAT33 IFA053 Annex A Alignment of terms and concepts related to the PIM • NFVIFA(24)000302r1 - FEAT33 IFA053 Removal of placeholder ENs • NFVIFA(24)000303 - IFA053 Reference updates • NFVIFA(24)000304r1 - IFA053 Term alignments on physical resource • NFVIFA(24)000305r1 - IFA053 Scope improvements Rapporteur actions: <ul style="list-style-type: none"> - Addition of the word management for better readability of the first sentence in NFVIFA(24)000305r1
May 2024	V0.11.0	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000306r1 - IFA053 8.1 Introduction • NFVIFA(24)000309r1 - FEAT33 IFA053 Clause 3 Definitions • NFVIFA(24)000319r1 - FEAT33 IFA053 Clause 6 Editorial adjustments • NFVIFA(24)000320 - FEAT33 IFA053 Clause 7 Editorial adjustments • NFVIFA(24)000323r1 - FEAT33 IFA053 Clause 4 Review Contribution • NFVIFA(24)000324r1 - FEAT33 IFA053 Annex B Review Contribution • NFVIFA(24)000326 - FEAT33 IFA053 Review contribution on alignment of physical resource pool

Date	Version	Information about changes
June 2024	V0.12.0	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000315r3 - FEAT33 IFA053 Review • NFVIFA(24)000316 - FEAT33 IFA053 Update Anuket and DMTF • NFVIFA(24)000317r1 - FEAT33 IFA053 ODIM • NFVIFA(24)000334r1 - FEAT33 IFA053 Resolution of remaining ENs • NFVIFA(24)000337r1 - FEAT33 IFA053 Editorial review of IFA053 v0.11.0 • NFVIFA(24)000338r1 - FEAT33 IFA053 Review Clause 7 Linking interfaces to IM • NFVIFA(24)000339r1 - FEAT33 IFA053 Review Clause 6.5 Addressing EN measurements alarms • NFVIFA(24)000340r1 - FEAT33 IFA053 Review topology information model and related concepts • NFVIFA(24)000341r1 - FEAT33 IFA053 Multiple clauses General technical review • NFVIFA(24)000346r1 - FEAT33 IFA053 EN resolution in Annex B • NFVIFA(24)000336r2 - FEAT33 IFA053 Review Addressing ENs Rapporteur actions: <ul style="list-style-type: none"> • While implementing NFVIFA(24)000316: creation of MS Word links for easier access, minor rephrasing for the pointers to other clauses, addition of appropriate name while referencing [i.20] in text • Simplified the name of [i.10] in clause 2.2 • Remove duplicate mention of references in text while implementing #337r2 in clause 4.3.2 • Added IFA027 and IFA045 in informative references due to their usage in #339r1
August 2024	V5.1.2	Initial version of IFA053ed521 based on the published version v5.1.1
August 2024	V5.1.3	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000463 - IFA053ed521 Annex B Redfish - provisioning aspects and composition service • NFVIFA(24)000453 - FEAT33 IFA053ed521 Annex B Additional description of Redfish features and capabilities Rapporteur actions: <ul style="list-style-type: none"> - Adjust clause numbers and fix reference numbers while implementing the new contributions
September 2024	V5.1.4	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(24)000497r1 - IFA053ed521 Annex B Analysis of solutions against PIM services • NFVIFA(24)000498r1 - IFA053ed521 Provisioning and composition updates • NFVIFA(24)000530 - FEAT33 IFA053ed521 A.3.2 Small Improvement
October 2024	V5.1.5	Implements following contribution: <ul style="list-style-type: none"> • NFVIFA(24)000596 - FEAT19a IFA053ed521 Adding container networking recommendation
April 2025	V5.2.2	Initial version of IFA053ed531 based on the published version v5.2.1
May 2025	V5.2.3	Implements following contribution: <ul style="list-style-type: none"> • NFVIFA(25)000103r2 - IFA053ed531 Update references to IFA027 and IFA045 Rapporteur actions: <ul style="list-style-type: none"> • Reverse the order of normative references for IFA027 and IFA045 while implementing #103r2
September 2025	V5.3.2	Initial version of IFA053ed541 based on the published version v5.3.1
September 2025	V5.3.3	Implements following contributions: <ul style="list-style-type: none"> • NFVIFA(25)000226r1 - FEAT33 IFA053ed541 Clarification about active and idle states • NFVIFA(25)000227 - IFA053ed541 Information Model Updates

History

Version	Date	Status
V5.1.1	August 2024	Publication
V5.2.1	November 2024	Publication
V5.3.1	August 2025	Publication
V5.4.1	February 2026	Publication